

OXFORD

THE FOUNDATIONS OF BEHAVIORAL ECONOMIC ANALYSIS



SANJIT DHAMI

Praise for *The Foundations of Behavioral Economic Analysis*

“The publication of this book is a landmark occasion for the field of behavioural economics. Until now there has been no comprehensive survey of the field suitable for graduate students. Professor Dhami has thoroughly and rigorously filled that gap. The book will be placed in a handy place in my office since I plan to consult it regularly.”

Richard H. Thaler, Charles R. Walgreen Distinguished Service Professor of Economics and Behavioral Science, University of Chicago

“*Foundations of Behavioral Economic Analysis* offers a fascinating mix of theory and evidence and is the most comprehensive synthesis of behavioral economics at an advanced level. It will be very useful for advanced researchers as well as for graduate students in behavioral economics and beyond.”

Ernst Fehr, Professor of Economics, University of Zurich

“This book is a tour de force, a literal encyclopedia of behavioral economics. Its extraordinary breadth and depth, spanning all aspects from psychological foundations to the most recent advances and seamlessly integrating theory with experiments, will make it the must-have reference for anyone interested in this field, and more generally in where economics is headed. It will quickly become the standard textbook for all graduate courses in behavioral economics, and a much-thumbed companion for all researchers working at the frontier.”

Roland Benabou, Theodore A. Wells '29 Professor of Economics and Public Affairs, Princeton University

“In *Foundations of Behavioral Economic Analysis*, Sanjit Dhami offers the first summary and exposition of research in this rapidly growing and increasingly influential subfield. The coverage is comprehensive, extending even to the recent subtopics of behavioural welfare economics and neuroeconomics. The book is distinguished by its detailed yet readable coverage of theory and evidence and its balanced discussion of the philosophical and methodological differences and similarities between ‘behavioural’ and neoclassical approaches to microeconomics. Select undergraduates, graduate students, and interested scholars will all gain from this masterful book.”

Vincent P. Crawford, Drummond Professor of Political Economy, University of Oxford, and Research Professor, University of California, San Diego

“Economic theory in the twentieth century developed an extremely powerful repertoire of analytical techniques for studying human behavior, but labored under the rather bizarre misconception that the postulates of rational choice were sufficient to characterize economic behavior. Behavioral economic from the late twentieth century to the present demonstrated the explanatory power of hitching these analytical techniques to empirical data gleaned from laboratory and field experiments. The result has radically transformed economics as a scientific discipline, and the best is surely yet to come. Sanjit Dhami has performed a monumental task in consolidating this research and explaining the results in a rigorous yet accessible manner, while highlighting major controversies and sketching the central research questions facing us today.”

Herbert Gintis, Santa Fe Institute

“For someone, like myself, who started by being ignorant of the richness of the conversation within behavioral economics on a variety of issues, this magisterial volume is the ideal introduction, at once lucid and sophisticated.”

Abhijit V. Banerjee, Ford Foundation International Professor of Economics, Massachusetts Institute of Technology.

“In the development of any field there comes a moment where the results already established must be synthesized, explained and consolidated both for those in the field and those outside. In this amazing volume Sanjit Dhami has done just that and far more. This book will serve as an encyclopedic must-have reference for anyone seeking to do work in this field or just curious about it. The coverage is exhaustive and the exposition extremely clear and at a level suitable for advanced undergraduates, graduate students, and professionals. This is truly an achievement.”

Andrew Schotter Professor of Economics, New York University and Director, Center for Experimental Social Science

“Displaying wit and wisdom, in *Foundations of Behavioral Economic Analysis* Professor Dhami conveys both the substance and the excitement of the burgeoning field of behavioral economics. This remarkable volume will serve as a reference for practitioners and a compelling entry-point for the curious.”

George Loewenstein, Herbert A. Simon Professor of Economics and Psychology, Carnegie Mellon University

“This is a unique and truly remarkable achievement. It is a magnificent overview of behavioral economics, by far the best there is, and it should define the field for at least a generation. But it is much more than that. It is also a brilliant set of original discussions, with pathbreaking thinking on every important topic. An invaluable resource for policymakers, students, and professors—and if they want to try something really special, for everyone else.”

Cass R. Sunstein, coauthor of Nudge and Founder and Director of the Program on Behavioral Economics and Public Policy, Harvard Law School

“The expansion of behavioral economics during the past twenty years has been remarkable, much of it concerning strategic interaction and using tools from game theory. Sanjit Dhami’s amazing book summarizes and even defines the field, broadly as well as in depth. His coverage of theory as well as of experiments is superb. *Foundations of Behavioral Economic Analysis* will be an indispensable resource for students and scholars who wish to understand where the action is.”

Martin Dufwenberg, Karl & Stevie Eller Professor and Director of the Institute for Behavioral Economics at the University Arizona

“Sanjit Dhami’s *Foundations of Behavioral Economic Analysis* is a major and most impressive achievement. It provides an exhaustive account and a masterful synthesis of the state of the art after more than three decades of behavioral economics. For many years to come it will be an indispensable reference for researchers in economics and psychology, and it is bound to become the standard text in graduate and advanced undergraduate courses on behavioural and experimental economics.”

Klaus M. Schmidt, Professor of Economics, University of Munich

“This is the most complete and stimulating book on behavioral economics. With elegance and unprecedented elaborateness, it ties together a wealth of experimental findings, rigorous theoretical insights, and exciting applications across all relevant fields of behavioral research. Sanjit Dhami’s work has been shaped by numerous comments of the leaders in the field. Now,

in the years to come, it will be the standard that shapes how the next generation of students and researchers think about behavior and its science."

Axel Ockenfels, University of Cologne, Speaker of the Cologne Excellence Center of Social and Economic Behavior

"This book covers all relevant theoretical aspects of behavioral economics in great depth. A great strength is its comprehensiveness: it covers the whole field. The book thus is unique in bringing to the fore the unity and diversity of the behavioral approach. The material is well-organized and accessible to a wide audience. It is invaluable to anyone teaching or studying any topic in behavioral economics, showing how the topic fits into the whole."

Peter Wakker, Professor of Economics, Erasmus University Rotterdam

"*Foundations of Behavioral Economic Analysis* will be a central textbook for behavioral economics. One key feature is its appealing focus on the interplay between theory and evidence. For researchers, it will be a great source of information, puzzles, and challenges for the many years to come. It is a major achievement."

Xavier Gabaix, Martin J. Gruber Professor of Finance, New York University, Stern School

The Foundations of Behavioral Economic Analysis

The Foundations of Behavioral Economic Analysis

SANJIT DHAMI

OXFORD
UNIVERSITY PRESS

OXFORD

UNIVERSITY PRESS

Great Clarendon Street, Oxford, OX2 6DP,
United Kingdom

Oxford University Press is a department of the University of Oxford.
It furthers the University's objective of excellence in research, scholarship,
and education by publishing worldwide. Oxford is a registered trade mark of
Oxford University Press in the UK and in certain other countries

© Sanjit Dhani 2016

The moral rights of the author have been asserted

First Edition published in 2016

Impression: 1

All rights reserved. No part of this publication may be reproduced, stored in
a retrieval system, or transmitted, in any form or by any means, without the
prior permission in writing of Oxford University Press, or as expressly permitted
by law, by licence or under terms agreed with the appropriate reprographics
rights organization. Enquiries concerning reproduction outside the scope of the
above should be sent to the Rights Department, Oxford University Press, at the
address above

You must not circulate this work in any other form
and you must impose this same condition on any acquirer

Published in the United States of America by Oxford University Press
198 Madison Avenue, New York, NY 10016, United States of America

British Library Cataloguing in Publication Data

Data available

Library of Congress Control Number: 2015958965

ISBN 978-0-19-871552-8 (hardback)

978-0-19-871553-5 (paperback)

Printed in Great Britain by
Bell & Bain Ltd., Glasgow

Links to third party websites are provided by Oxford in good faith and
for information only. Oxford disclaims any responsibility for the materials
contained in any third party website referenced in this work

To my Parents, Shammi and Sahaj.

PREFACE AND ACKNOWLEDGMENTS

Neoclassical economics is a logically consistent and parsimonious framework of analysis that is based on a relatively small set of core assumptions, and it offers clear, testable, predictions. However, extensive and growing empirical evidence reveals human behavior that is difficult to reconcile within the typical neoclassical models. There has been a parallel development in rigorous theoretical models that explains better the emerging stylized facts on human behavior. These models have borrowed insights from psychology, sociology, anthropology, neuroscience, and evolutionary biology. Yet, these models maintain a distinct economic identity in terms of their approach, rigor, and parsimony. Collectively, these models form the subject matter of behavioral economics, which is possibly the fastest growing and most promising area in economics.

This book is an account of behavioral economics that starts with the basics and takes the reader to the research frontiers in the subject. Depending on how one chooses to use it, the book is suitable for courses at the advanced undergraduate, postgraduate, and research level in economics, and the related social sciences, including, but not restricted to, psychology, management, finance, political science, and sociology. The book should also serve as an essential reference book for anyone generally interested in behavioral economics at any level, and also serve to stimulate the interests of non-specialist academics, specialist academics who are looking for a bird's-eye view of the entire field, and policymakers looking for policy applications of behavioral economics. It would be desirable to assign this book as background reading to courses in economic theory. The book is also, in my view, the minimum subject matter that anyone who writes behavioral economics as their research interest, should be deeply familiar with.

In November 2003, two months after I joined the department of economics at the University of Leicester, I chanced upon an invitation to attend a talk by a colleague, Ali al-Nowaihi, on the subject of *prospect theory*. Ali, a mathematician by training, an economist by profession, and a keen student of the philosophy of science, put forward a Popperian view to evaluate economic theories. He argued that *expected utility theory* was decisively rejected by the evidence, and prospect theory was the most satisfactory decision theory currently available. As a purely neoclassically trained economist, I was troubled by the claims, but also extremely skeptical. For a start, prospect theory sounded like a strange name for a theory, and the evidence was largely “experimental,” a data source, that I knew little about. As my defensive instincts started to kick in, I wondered if prospect theory really was so important, then surely my graduate courses, many taught by leading decision theorists, would have found some reason to mention it. Nor was there any mention of such a theory in conversations with colleagues at the two British universities where I had taught so far, or at seminars or conferences that I had attended.

However, rather than just dismiss Ali, a very likeable and respected figure in the department, I decided to put his seemingly extreme views to the test. One of my majors was in public economics, so I decided to conduct a prospect theory analysis of tax evasion in the hope of explaining the *tax evasion puzzles*, which had been outstanding for three decades (details in Part 1). There was already some preliminary work in this area that Ali had mentioned in passing, but none of the papers explained all the puzzles in one fell swoop, using all components of prospect theory. It took me just a few weeks to work out the results. To my utter amazement, prospect

theory explained the qualitative and quantitative tax evasion puzzles. By contrast, the predictions based on an expected utility analysis were wrong by a factor of up to 100. This led to my first joint publication with Ali, with whom I have spent many years of fruitful collaboration since then.

This initial, and successful, encounter with prospect theory convinced me that I needed to explore behavioral economics in greater depth. Yet, around 2004, there was no definitive graduate text on behavioral economics. To be sure, there were many excellent sets of collected readings, and several insightful surveys and commentaries on selected aspects of behavioral economics that I eagerly read. In particular, while there were many excellent discussions of the experimental evidence, a full treatment of behavioral economic theory and its applications was missing. One could always pursue the journal articles, but the literature was already enormous, rapidly expanding, and scattered, which made it difficult to spot the links between the various models or to clearly visualize how the various pieces of the jigsaw fitted together. This book was motivated initially by the lack of a serious graduate book on the entire subject matter of behavioral economics, my desire to master behavioral economics, and to support my growing research agenda with Ali. In due course, and as the full range of the subject matter gradually dawned upon me, the scope of the book naturally became more ambitious and daring.

I strive to strike a balance between behavioral economic theory, the experimental evidence, and applications of behavioral economics. The choice of theoretical models in this book is dictated, first and foremost, by their ability to explain the empirical evidence. In some cases, where no decisive empirical evidence is available, I make a judgment on which models are more promising than others, although I give a wide berth to most models.

The main prerequisite for the book is training in the first two to three years of a reasonably good British or North American undergraduate degree in economics, or its equivalent. Any further concepts and techniques are introduced in the book, where needed. A prior course in behavioral economics is not a prerequisite for the book.

The book is divided into nine parts that cover decision making under risk, uncertainty, and ambiguity; other-regarding preferences; behavioral time discounting; models of behavioral game theory and learning; role of emotions in decision making; models of bounded rationality; judgment heuristics and mental accounting; behavioral welfare economics; and neuroeconomics. The book also considers a range of applications of the theory to most areas in economics that include microeconomics, contract theory, macroeconomics, industrial organization, labor economics, development economics, public economics, political economy, and finance. A set of exercises at the end of each part, except the part on neuroeconomics, serves to enhance the reader's understanding of the subject.

Behavioral economics is now a mainstream area in economics. One just has to look at the growing and large number of journal publications and Ph.D. theses every year; the Nobel Prizes to Herbert Simon, Daniel Kahneman, Robert Shiller, Alvin Roth, Vernon Smith, and George Akerlof; the John Bates Clarke medal to Matthew Rabin; the growing importance of behavioral economics among policymakers, as witnessed by the 2015 World Bank Development Report, and the formation of the behavioral insights team in the UK; and the choice of Richard Thaler as the incoming President of the American Economic Association.

It is fair to say that no self-respecting economics department can now afford to omit a course in behavioral economics from its undergraduate or graduate curriculum; indeed, doing so would be grossly unjust to its students and a retrogressive step. Nor can any academic economist, who wishes to retain professional honesty and a balanced opinion on the subject, afford to be unfamiliar with the subject matter of behavioral economics; I am often amused by the ignorance and

arrogance of many who pass judgment on behavioral economics with supreme confidence, yet appear to have little understanding of it.

This book has taken more than ten years to write, and my debts are deep and profound. My first and foremost debt and gratitude is to my loving family without which this book could not have been written. To my parents, Manohar and Baljeet, for their unconditional lifelong love and support, and instilling in me the core values of honesty, commitment, and hard work. To my wife, Shammi, and my son, Sahaj, for their patience, sacrifice, unflinching support, and constant encouragement. When I started writing this book, Sahaj was in primary school, and in the month of its first publication, he could be packing his bags to join a university. I do not recommend this as the best template to encourage your son to write any books in the future. However, there are close parallels between Sahaj's educational journey from primary school to university, with my own journey in behavioral economics.

I owe a deep intellectual debt to my long-time coauthor and friend, Ali al-Nowaihi. I first learnt about prospect theory from him. I also owe my appreciation of methodology and the philosophy of science entirely to him. He has undertaken a larger burden of our joint research in the last few years, allowing me to be immersed in the book. For all these reasons, he is very much a coauthor of the book in spirit.

I am extremely grateful to many academics and Ph.D. students who unselfishly and generously contributed their time and efforts to reading drafts of various parts of the book. The participation of so many leading behavioral economists in the making of this book is unprecedented and has really made it into a public project for which I shall always be very grateful. Herbert Gintis, Martin Dufwenberg, and Vincent Crawford deserve special mention for being so very gracious with their inputs into most parts of the book, and very quickly responding to my queries.

Many others also played a critical role in the writing of this book and commented on material closer to their areas of interest, and/or offered valuable encouragement and advice. In particular, I wish to thank Mohammed Abdellaoui, Ali al-Nowaihi, Dan Ariely, Douglas Barrett, Björn Bartling, Karna Basu, Kaushik Basu, Pierpaolo Battigalli, Roland Benabou, Florian Biermann, Gary Bolton, Subir Bose, David Colander, Andrew Colman, Patricio Dalton, Alexandra Dias, Florian Englmaier, Armin Falk, Ernst Fehr, Urs Fischbacher, Xavier Gabaix, Sayantan Ghosal, Uri Gneezy, Werner Guth, Shaun Hargreaves Heap, Fabian Herweg, Karla Hoff, Philippe Jehiel, David Laibson, George Loewenstein, Michel Marechal, Friederike Mengel, Joshua Miller, Axel Ockenfels, Amnon Rapoport, Ludovic Renou, Alvin Roth, Klaus Schmidt, Andrei Shleifer, Dennis Snower, Joe Stiglitz, Cass Sunstein, Richard Thaler, Jean-Robert Tyran, Peter Wakker, Klaus Waelde, Eyal Winter, and Peyton Young. I owe a profound intellectual debt to many others who did not read the book manuscript but whose work has greatly inspired me. These include Daniel Kahneman, Amos Tversky, Colin Camerer, Matthew Rabin, Herbert Simon, Robert Shiller, and George Akerlof. I am also very grateful to two successive Heads of the economics department at Leicester, Steve Hall and Chris Wallace, who tried to free up as much of my time as possible for writing the book.

I would like to specially acknowledge the enormous amount of work put in by two extremely conscientious and able Ph.D. students, Teimuraz Gogsadze and Junaid Arshad. They closely read and commented on successive drafts of the manuscript at all stages, offered very useful advice, and served as excellent sounding boards for new ideas. Jingyi Mao came up with a very nice cover for the book in a burst of creativity, for which I am very grateful. Other Ph.D. students who carefully read and commented on selected parts of the manuscript include: Ala Avoyan, Nino Dognohadze, Sneha Gaddam, Narges Hajimoladarvish, Emma Manifold, Jingyi Mao, Alexandros Rigos, David Tsirekidze, Yongli Wang, Mengxing Wei, and Mariam Zaldastanishvili.

I would be remiss not to thank the large number of other researchers whose work has made this book possible. I must also sincerely apologize to authors who feel that their work has been inadequately cited or not given the importance they feel that it deserves. To such authors, I say, omission of your papers does not mean that I necessarily viewed your papers as unimportant. In mitigation, I do not intend my book to be a survey of all the experimental results on all topics in behavioral economics; there are already excellent sources with this objective. And, quite possibly, I was simply unaware of your important work, which is in keeping with the evidence on limited attention and bounded rationality that plays an important role in this book.

I am very grateful to the team at Oxford University Press who have done an excellent job at all stages of this book. In particular, I would like to thank Adam Swallow, the commissioning editor for economics and finance at OUP for his patience, good cheer, organizational skills, and sound advice. Scott Parris, the economics editor at the US office of OUP, who retired just as this book was about to come out, was the first to spot the importance of this project. He offered very valuable advice and encouragement throughout the writing stage and played a key role in my decision to go with OUP. I must also thank Niko Pfund, the President of Oxford University Press USA, for his continued interest in the manuscript over several years, despite his many other responsibilities. The production and marketing teams at OUP were a pleasure to work with. Jon Billam took on the challenge of copy-editing an unusually large book with great enthusiasm. I am also very grateful to Emma Slaughter, the production editor for the book; Kim Stringer, the indexer; Kim Allen, the proofreader; Carla Hodge-Degler who took over as production editor from Emma; and to Leigh-Ann Bard, the marketing manager for the book.

CONTENTS

<i>List of Figures</i>	xxi
<i>List of Tables</i>	xxx

Introduction	1
1 The antecedents of behavioral economics	3
2 On methodology in economics	5
3 The experimental method in economics	10
4 Approach and organization of the book	21
5 Five theoretical approaches in behavioral economics	25
6 Five examples of behavioral evidence	48
Appendix A: The random lottery incentive mechanism	59
Appendix B: In lieu of a problem set	59
References	65

PART 1 Behavioral Economics of Risk, Uncertainty, and Ambiguity

Introduction to Part 1	79
1 The Evidence on Human Choice under Risk and Uncertainty	83
1.1 Introduction	83
1.2 The elements of classical decision theory	84
1.3 Subjective expected utility theory (SEU)	89
1.4 Eliciting the utility function under EU	92
1.5 Violations of expected utility theory	95
2 Behavioral Models of Decision Making	110
2.1 Introduction	110
2.2 Probability weighting functions	115
2.3 Rank dependent utility theory (RDU)	124
2.4 Prospect theory (PT)	130
2.5 Elicitation of utility and probability weighting functions in PT	140
2.6 The axiomatic foundations of PT	148
2.7 Third generation PT and stochastic reference points	152
2.8 Stochastic reference points in PT under rational expectations	153
2.9 Limitations of PT	172
2.10 A selection of other behavioral theories	181
2.11 Human behavior for extreme probability events	191
2.12 Risk preferences and time preferences	199
3 Applications of Behavioral Decision Theory	213
3.1 Introduction	213

3.2	The endowment effect and exchange asymmetries	217
3.3	Prospect theory preferences in primates	235
3.4	Myopic loss aversion	239
3.5	Why do people pay taxes?	242
3.6	Explanation of Rabin's paradox using PT	248
3.7	Goals as reference points	250
3.8	Why is it so hard to find a taxi on a rainy day in New York?	256
3.9	Some implications of inverse S-shaped weighting functions	263
3.10	Contracts as reference points	264
3.11	Moral hazard, loss aversion, and optimal contracts	272
3.12	Renegotiation, long-term contracts, and loss aversion	278
4	Human Behavior under Ambiguity	284
4.1	Introduction	284
4.2	A problem for SEU: The Ellsberg paradox	286
4.3	Neoclassical models of ambiguity	288
4.4	Behavioral models of ambiguity	301
	Exercises for Part 1	310
	References for Part 1	316
 PART 2 Other-Regarding Preferences		
Introduction to Part 2		339
5	The Evidence on Human Sociality	344
5.1	Introduction	344
5.2	Ultimatum and dictator games	349
5.3	Gift exchange and trust games	357
5.4	Public goods games	370
5.5	How representative is the lab evidence?	386
5.6	Indirect reciprocity	395
6	Models of Other-Regarding Preferences	398
6.1	Introduction	398
6.2	The Fehr–Schmidt model	401
6.3	The ERC model	410
6.4	Fairness and stochastic dominance	412
6.5	Behavioral political economy	417
6.6	Fairness, general equilibrium, and welfare	426
6.7	Evidence on models of social preference and reciprocity	431
6.8	A discussion of other promising models	444
7	Human Morality and Social Identity	453
7.1	Introduction	453
7.2	Human virtues	455
7.3	Social identity	465

8	Incentives and Other-Regarding Preferences	478
8.1	Introduction	478
8.2	Moral hazard and other-regarding preferences	483
8.3	Incomplete contracts under other-regarding preferences	495
8.4	Reciprocity and long-term contracts	510
8.5	Extrinsic and intrinsic motivation: theoretical framework	527
8.6	Extrinsic and intrinsic motivation: empirical evidence	540
	Exercises for Part 2	562
	References for Part 2	566

PART 3 Behavioral Time Discounting

	Introduction to Part 3	583
9	The Evidence on Temporal Human Choice	586
9.1	Introduction	586
9.2	The discounted utility model	587
9.3	The exponential discounted utility model (EDU)	588
9.4	Anomalies of the EDU model	594
10	Behavioral Models of Time Discounting	605
10.1	Introduction	605
10.2	Explaining the sign and magnitude effects	607
10.3	Explaining the common difference effect	611
10.4	Attribute-based models	627
10.5	The reference time theory (RT)	632
11	Applications of Present-Biased Preferences	644
11.1	Introduction	644
11.2	Two frameworks for modeling present-biased preferences	646
11.3	Optimal consumption under hyperbolic discounting	651
11.4	Endogenous retirement decisions	664
11.5	Procrastination and preproperation	669
11.6	Interdependence, multiple activities, and competition	677
11.7	Optimal sin taxes	679
11.8	Investment and leisure goods	683
11.9	Other applications of hyperbolic discounting	688
	Exercises for Part 3	692
	References for Part 3	700

PART 4 Behavioral Game Theory

	Introduction to Part 4	709
12	The Evidence on Strategic Human Choice	713
12.1	Introduction	713

12.2 Iterated deletion of dominated strategies	724
12.3 Mixed strategy Nash equilibria	738
12.4 Coordination games	758
12.5 Bargaining games	786
12.6 Asymmetric information, signaling, and cheap talk	819
12.7 Public signals and correlated equilibria	843
12.8 Strategic complements and strategic substitutes	852
12.9 A digression on competitive equilibrium experiments	863
13 Models of Behavioral Game Theory	875
13.1 Introduction	875
13.2 Quantal response equilibrium (QRE)	881
13.3 Level- k and cognitive hierarchy models	894
13.4 Applications of level- k and CH models	904
13.5 Psychological game theory	925
13.6 Correlated equilibrium and social norms	962
13.7 Other behavioral models of how people play games	966
13.8 Behavioral economics of auctions	988
Exercise for Part 4	1003
References for Part 4	1013
 PART 5 Behavioral Models of Learning	
Introduction to Part 5	1033
14 Evolutionary Game Theory	1037
14.1 Introduction	1037
14.2 Some preliminaries	1041
14.3 Evolutionary stable strategies in two-player games	1043
14.4 Relation of ESS with other equilibrium concepts	1047
14.5 Replicator dynamics	1048
14.6 Applications of evolutionary game theory	1053
14.7 Playing the field	1063
14.8 Gene-culture coevolution	1063
14.9 The evolution of altruism and reciprocity	1073
14.10 Social learning and replicator-like dynamics	1085
15 Models of Learning	1092
15.1 Introduction	1092
15.2 Reinforcement learning	1097
15.3 Belief-based models of learning	1106
15.4 The experience weighted attraction model of learning	1120
15.5 Learning models: two points of view	1131
15.6 The class of learning direction theories	1132
15.7 Rule-based learning	1142
15.8 Multiple games, complexity, and categorization	1148
15.9 Possibility and impossibility of rational learning	1153

16 Stochastic Social Dynamics	1158
16.1 Introduction	1158
16.2 Fixing ideas in stochastic dynamics	1159
16.3 Perturbed Markov dynamics in 2×2 games	1166
16.4 Memory, conventions, and risk dominance in 3×3 games	1174
16.5 Social networks	1182
Appendix A: Ordinary differential equations	1186
Appendix B: A note on choice under stochastic utility	1193
Appendix C: A brief user's guide to Markov Chains	1196
Exercises for Part 5	1205
References for Part 5	1210

PART 6 Emotions

Introduction to Part 6	1223
17 Emotions and Human Behavior	1226
17.1 Introduction	1226
17.2 Visceral influences and the rationality of emotions	1229
17.3 Cue-conditioned behavior and habit formation	1235
17.4 Anticipation and delay under certainty	1242
17.5 Fear and anxiety under uncertainty	1245
17.6 Projection bias	1250
17.7 Temptation preferences: a revealed preference approach	1258
17.8 Temptation and conflicts between commitment and flexibility	1266
17.9 Happiness economics	1270
18 Interactions between Emotions and Cognition	1288
18.1 Introduction	1288
18.2 Emotions and a two-modules view of the brain	1290
18.3 A dual selves model with costly commitment	1298
18.4 Information asymmetries between emotions and cognition	1304
18.5 Strategic ignorance, confidence, and self-esteem	1308
Exercises for Part 6	1322
References for Part 6	1328

PART 7 Bounded Rationality

Introduction to Part 7	1339
19 Judgment Heuristics	1342
19.1 Introduction	1342
19.2 The law of small numbers	1350
19.3 Conjunction fallacy	1361
19.4 The availability heuristic	1365

19.5	The affect heuristic	1367
19.6	Anchoring and adjustment	1370
19.7	Base rate neglect and conservatism	1375
19.8	Hindsight bias	1383
19.9	Confirmation bias	1389
19.10	Other judgment heuristics	1395
19.11	Dual process models and judgment heuristics	1398
19.12	Coarse thinking and persuasion	1401
19.13	Mental models	1407
19.14	Herbert Simon's approach to bounded rationality	1409
19.15	The great rationality debate	1426
19.16	Shrouded attributes	1431
19.17	Limited attention	1436
19.18	Do experts exhibit biases?	1449
20	Mental Accounting	1451
20.1	Introduction	1451
20.2	Framing and mental accounting	1455
20.3	Prospect theory and mental accounting	1456
20.4	The behavioral life-cycle model	1459
20.5	The red and black of mental accounting	1463
20.6	Choice bracketing	1469
20.7	Coherent arbitrariness	1471
20.8	Sunk costs and mental accounting	1473
20.9	Some other mental accounting phenomena	1476
21	Bounded Rationality in Financial Markets	1485
21.1	Introduction	1485
21.2	The efficient markets hypothesis	1488
21.3	Noise trader risk	1502
21.4	The limits to arbitrage	1509
21.5	Gradual flow of information	1516
21.6	Stock market underreaction and overreaction	1520
21.7	Behavioral corporate finance	1526
	Exercises for Part 7	1544
	References for Part 7	1552

PART 8 Behavioral Welfare Economics

Introduction to Part 8	1577
22 Behavioral Welfare Economics	1579
22.1 Introduction	1579
22.2 Fixing basic ideas	1583
22.3 Soft paternalism	1590
22.4 Regulation under imperfect self-information	1601
22.5 Choice and non-choice data: What is the scope of economics?	1603

22.6 Choice-based behavioral welfare economics	1611
22.7 Revealed preference under limited attention	1618
22.8 The contractarian approach	1623
22.9 Behavioral public finance and welfare	1625
Exercises for Part 8	1635
References for Part 8	1638

PART 9 Neuroeconomics

Introduction to Part 9	1643
23 Neuroeconomics	1644
23.1 Introduction	1644
23.2 A brief introduction to the brain	1648
23.3 An introduction to neuroeconomic techniques	1655
23.4 The neuroeconomics of risky decisions	1661
23.5 The neuroeconomics of social preferences	1669
23.6 The neuroeconomics of time preferences	1686
23.7 The neuroeconomics of strategic interaction	1689
23.8 Pharmacoeconomics: an application to the social effects of oxytocin	1692
References for Part 9	1695
<i>Appendix on Game Theory</i>	1703
<i>Index of Named Authors</i>	1725
<i>Subject Index</i>	1736

LIST OF FIGURES

1	The power form of the utility function under prospect theory.	27
2	A plot of the Prelec function for $\beta = 1$ and $\alpha = 0.5$.	28
3	A plot of the discount rate for the one-parameter hyperbolic discounting function with varying time delay.	38
4	Distribution of choices by different subject pools in a p-beauty contest game.	42
5	Choices of game theorist subjects in a p-beauty contest game.	43
6	Response of quality to price in a competitive equilibrium.	50
7	Response of quality to price for local and non-local dealers in the lab and in the field.	51
8	Public good contributions in the presence and absence of punishments.	52
1.1	Illustration of the common ratio effect in a probability triangle.	100
1.2	Concave indifference curves.	102
1.3	Convex indifference curves.	103
2.1	Relative risk premium for the Student and the Representative population.	117
2.2	Non-parametrically estimated probability weighting functions for the Student and Representative population based on aggregate data.	118
2.3	Plots of $w(p) = e^{-(\ln p)^{0.5}}$ and $w(p) = e^{-(\ln p)^2}$.	120
2.4	An illustration of optimism and pessimism as the shape of the probability weighting function varies.	126
2.5	Shapes of decision weights when there are multiple outcomes under RDU.	127
2.6	The utility function under prospect theory.	132
2.7	Indifference curves under RDU over the full range of probabilities (left panel) and with the boundaries truncated (right panel).	137
2.8	Non-parametric estimates of utility for gains and losses under PT and EU based on median data.	143
2.9	Parametric estimates of the power form of utility for gains and losses under PT based on median data.	144
2.10	The case of two choice sets.	158
2.11	Histogram of accumulated earnings (in euros) at which a subject stopped.	170
2.12	Decision tree for experiment 1 in Ericson and Fuster (2011).	171
2.13	Median relative risk premium (<i>RRP</i>) for different types of decision makers as the probability, p , is varied.	177
2.14	Stake dependent probability weights in the domain of losses and gains (upper panel) and 95% confidence bands based on the percentile bootstrap method (lower panel).	178
2.15	Two types of decision makers for the student and the representative population.	180
2.16	A hypothetical probability weighting function.	192
2.17	The composite Prelec weighting function (CPF).	196
2.18	Reduction of lotteries when both time and risk are present, in three different cases.	200
2.19	Average consumption, $\bar{c}_t$, in various conditions.	202

2.20	A plot of the estimated discount rate $\hat{\theta}$ (vertical axis) and the time delay t (horizontal axis).	210
2.21	A choice between a compound lottery and its equivalent simple lottery.	211
2.22	Estimated Prelec weighting functions.	212
3.1	Fraction of successful par and birdie putts by distance to the hole.	228
3.2	Crossing of indifference curves when there is an endowment effect.	231
3.3	Tests of capuchin monkey preferences.	236
3.4	Choices made by chimpanzees in different conditions.	238
3.5	Average bets of students and professionals.	240
3.6	Amount invested in each treatment.	241
3.7	Plot of (3.30) for the case of $\lambda = 2$.	249
3.8	An illustration of the equilibrium in the Crawford–Meng (2011) model.	262
3.9	Valuations and costs in two states of the world.	269
3.10	Valuations and costs in three states of the world.	269
3.11	Quality dependent valuation and cost data in two states of the world.	270
3.12	Sequence of moves in a bilateral trade model with renegotiation.	279
4.1	Source-dependent probability weighting functions.	307
4.2	Subject heterogeneity for three subjects expressed in terms of their source functions.	308
5.1	The relation between effort and rent in the gift exchange game.	358
5.2	The effect of the equity norm on effort levels.	361
5.3	Behavior of workers' effort levels in the library task.	362
5.4	The relation between reciprocity and tenure.	363
5.5	Productivity of workers in different wage conditions.	364
5.6	Average responses for the positive reciprocity questions (left panel) and negative reciprocity questions (right panel).	365
5.7	Measures of trust and trustworthiness in a meta study.	367
5.8	The trust game in Huck et al. (2012).	369
5.9	Evolution of average contributions in three different treatments in the presence and in the absence of punishments.	374
5.10	Explaining the pattern of contributions when there is heterogeneity in social preferences.	378
5.11	Choice of institutions and contributions.	379
5.12	Contest expenditures in the treatments without punishment.	381
5.13	Actual and equilibrium rent dissipation.	383
5.14	Prosocial and antisocial punishment in the public goods game (left panel) and the relation between antisocial punishment and contributions (right panel).	384
5.15	Mean contributions in the P treatment (left panel) and the N treatment (right panel).	385
5.16	Results for the single- and double-anonymous treatments.	388
5.17	Workers' performance response to the wage increase (vertical axis), conditional on their fairness perceptions.	389
5.18	Checking the external validity of lab data on social preferences.	391
5.19	Average time spent on monitoring by different behavioral types. The bars represent standard errors of the means.	393
5.20	Third party punishments in the dictator game.	397
6.1	The SSS and SFS economies.	425
6.2	The SSS, SSF, and FSF economies.	426

6.3	Modified ultimatum game.	435
6.4	Effect on b as a changes.	437
6.5	Normal form representation of a two-player, sequential, social dilemma game.	440
6.6	Cumulative distribution of the second player's cooperative decision (c1 to c6).	441
6.7	Two variations of the ultimatum game, indexed by the parameters x and y .	442
6.8	Average rejection rates in the UG^*x game.	443
6.9	Average rejection rates in the UG^*+y game.	443
7.1	Results for the baseline treatment.	458
7.2	Messages sent by player A in each state.	460
7.3	The battle of sexes (BOS) game and the prisoner's dilemma game (PD).	470
7.4	Cooperation rates and beliefs in the modified trust game. Bars represent standard errors of the means.	471
7.5	Group membership and punishment. Bars represent standard errors of the means.	472
8.1	Share of bonus contracts and incentive contracts over successive rounds.	506
8.2	Relation between the actual bonus paid by the principal and effort level of the agent.	507
8.3	A comparison of the length of the worker-firm relationship under the two treatments.	518
8.4	Average temporal wage in each treatment.	519
8.5	Average temporal effort in each treatment.	520
8.6	Average earnings per trade of the workers and firms in each treatment.	521
8.7	Comparison of mean effort levels under various treatments.	522
8.8	Payoffs of the principal and agent in treatment RANDOM (left panel) and treatment FIXED (right panel).	525
8.9	Share of employment contracts in different treatments (left panel). Frequency of power abuse and share of employment contracts in each treatment (right panel).	526
8.10	Diagrammatic representation of the condition in (8.77).	530
8.11	Loci of combinations of (u_I, u_E) on which (8.86) binds.	533
8.12	A trust game that differentiates between voluntary and involuntary trust.	536
8.13	Effort and rents in three cases.	541
8.14	Effect of intentions in the trust game.	542
8.15	Median choice of action by the agent over when the principal does, and does not, control.	543
8.16	Mean earnings as a percentage of maximum earnings in different games.	545
8.17	Mean earnings as a percentage of maximum earnings in different games.	545
8.18	Average profits and average effort levels conditional on the employees' reputations.	547
8.19	Share of contract-F in each treatment.	548
8.20	Buyer/seller profits as a function of buyer/seller choices.	549
8.21	Average number of parents who arrive late.	551
8.22	Percentage of blood donors under varying levels of incentives.	553
8.23	The effect of incentives on prosocial behavior in different treatments.	555
8.24	Total public goods contributions over successive rounds when the incentives are varied.	557
8.25	Number of weekly visits to the gym. Error bars reflect one standard error.	561
9.1	A plot of the discount rate in (9.30) for the one-parameter hyperbolic discounting function with varying time delay.	596
9.2	Sequential choices for the Ok and Masatlioglu (2007) example.	602

10.1	The discount factors for various discount functions: Exponential (continuous bold curve), quasi-hyperbolic (lighter step function), and generalized hyperbolic (dashed curve).	616
10.2	Sequence of outcome–time pairs.	626
10.3	The delay and the speedup functions.	639
11.1	Diagrammatic illustration of the equilibrium in the Diamond–Kőszegi model.	669
12.1	The extensive form and the strategically equivalent normal form in the baseline treatment.	725
12.2	An extensive form and the strategically equivalent normal form.	727
12.3	Checking for three steps of iterated deletion of dominated strategies.	728
12.4	Four-move and six-move centipede games.	729
12.5	A three-player centipede game.	731
12.6	A six-move centipede game.	732
12.7	Percentage of all-B play in various rounds.	737
12.8	O’Neill’s zero-sum game to test for an MSE.	739
12.9	The payoff matrix for games 1 and 3.	742
12.10	The payoff matrix for games 2 and 4.	742
12.11	A three-player matching pennies game.	744
12.12	The matching pennies and the unprofitable game.	745
12.13	Distribution of locations for pooled data.	746
12.14	The pursuer–evader game and the 4×4 game in Levitt et al. (2010).	755
12.15	Type-I game ($z = 600$) and Type-II game ($z = 500$).	759
12.16	Coordination game (100 treatment).	762
12.17	Game 2R, game R, and game 0.6R.	768
12.18	A simple coordination game (SCG).	771
12.19	Median choice in sessions 1–10 over 15 periods in the continental divide game.	773
12.20	Average minimum effort levels over rounds 11–20 for different treatments.	775
12.21	Three equivalent stag-hunt games.	777
12.22	Minimum action chosen under different treatments in Block-I games.	781
12.23	Time paths of equilibria.	782
12.24	Three coordination games.	783
12.25	Average minimum effort in different treatments and rounds.	784
12.26	Average minimum effort level in different treatments and rounds.	785
12.27	A coordination stage game.	786
12.28	Frequency of agreements and disagreements over the whole period (left panel) and in the last 30 seconds (right panel).	793
12.29	Opening offers to player 2 for cells 1, 2, 5, and 6 in Table 12.22.	797
12.30	Opening offers to player 2 for cells 3, 4, 7, and 8 in Table 12.22.	798
12.31	Predictions of the model (bold line) versus the actual data in Roth et al. (1991) (dotted line) for offers and rejections in ultimatum games.	800
12.32	Player 1’s offer of the pie to player 2.	802
12.33	Optimal linear bidding functions for sellers and buyers.	807
12.34	Optimal bidding strategies for buyers and sellers in experiment 1 (left panel) and in experiment 2 (right panel).	809
12.35	The trade and no-trade regions in a two-sided asymmetric information bargaining game.	810

12.36	Trade outcomes for the no-communication, written, and face to face treatments.	812
12.37	Games G1 and G2.	814
12.38	Description of a bargaining table for one of the 24 games in Isoni et al. (2013).	815
12.39	The MOUSELAB computer display for the three-round bargaining game.	817
12.40	The signaling game in Brandts and Holt (1992).	820
12.41	Results for Treatment I.	826
12.42	Results for Treatment II.	828
12.43	Contribution rates under asymmetric information.	838
12.44	Contribution rates under full information.	839
12.45	Relative frequencies of cooperation, coordination, and efficiency.	842
12.46	The set of possible payoffs in a correlated equilibrium.	845
12.47	Percentage of actions consistent with the recommended play in the various treatments.	847
12.48	Beliefs for different recommendations for periods 31–60.	848
12.49	Symmetric game of chicken in Duffy and Feltovich (2010).	848
12.50	Equilibria in the game of chicken in Duffy–Feltovich.	850
12.51	The best response functions and equilibria in various cases.	854
12.52	Average degree of cooperation in each round for complements and substitutes.	855
12.53	Average degree of cooperation for non-JPM and JPM pairs.	856
12.54	A partial representation of the Fehr–Tyran (2008) experimental setting.	857
12.55	Payoff matrix for type y players in the substitutes case for the post-shock period.	859
12.56	Average nominal price across time in the complements and substitutes treatment.	860
12.57	Actual average prices and average best replies in the complements and substitutes treatments (Periods 16–18).	861
12.58	A plot of the demand and supply curves in Table 12.34.	866
12.59	The time path of prices over successive rounds of the experiment.	867
12.60	Price dynamics.	870
12.61	Price dynamics with ZI-C and ZI-U robot traders and humans.	873
13.1	Two-player matching pennies game.	884
13.2	QRE and Nash equilibrium for the cases $x = 1$ and $x = 3$.	885
13.3	Three normal form games.	886
13.4	Predictions of the quantal response equilibrium for games A, B, and C for various values of λ .	887
13.5	The bilateral crises game.	889
13.6	Claims, over time, for various values of R .	893
13.7	Average claims, logit predictions, and the Nash prediction in the traveller's dilemma game.	894
13.8	Results for the p-beauty contest game when $p = 2/3$.	898
13.9	Results for the p-beauty contest using diverse subject pools with $p = 2/3$, $a = 0$, $b = 100$.	899
13.10	The non-incentivized Chicago skyscrapers game in Crawford et al. (2008) with permission from American Economic Association.	905
13.11	Results with incentivized experiments in Crawford et al. (2008) with permission from American Economic Association.	905
13.12	Battle of Sexes game in Crawford (2007); $a > 1$.	908

13.13	A stag-hunt game.	912
13.14	A two-player market entry game ($a > 1$).	916
13.15	MOUSELAB interface in one of the games in Costa-Gomes et al. (2001) with permission of The Econometric Society.	917
13.16	The bravery game.	926
13.17	A guilt game.	927
13.18	The confidence game in Geanakoplos et al. (1989).	931
13.19	A simple example of a sequential move psychological game.	933
13.20	Dependence of utilities on beliefs of others.	935
13.21	Sequential prisoner's dilemma game.	937
13.22	Correlation between contributions and first order beliefs in a public goods game.	944
13.23	The effect of c and μ_R on p .	949
13.24	The relation between contributions and second order beliefs in the public goods game.	951
13.25	A partnership game in the absence and presence of guilt.	952
13.26	The effect of communication/promises on partnerships.	954
13.27	Bubble plots of the relation between induced second order beliefs and contributions in the dictator game (left panel) and the amount sent back by trustees in the trust game (right panel).	955
13.28	The distribution of the coefficients on guesses that are significant at the 5% level.	959
13.29	A psychological game in which utility depends on plans that are off the equilibrium path of play.	961
13.30	A coordination game.	963
13.31	A sequential move game.	967
13.32	Games G1 and G2 in Huck et al. (2011).	968
13.33	The prisoner's dilemma game.	977
13.34	A coordination game.	984
13.35	Logit best responses for the game in Figure 13.34 when the payoff from (R, S) is $(-14, 4)$.	985
13.36	An experimental game from Colman et al. (2014).	987
13.37	Actual and fitted bids as a function of the signals.	997
4.E1	Percentage of (B, B) choices in various rounds.	1004
4.E2	Three simple strategic form games and the choices of players.	1005
4.E3	An asymmetric matching pennies game in Goeree et al. (2005)	1006
4.E4	The Far Pavilions Game.	1007
4.E5	A zero-sum game of hide-and-seek.	1007
4.E6	Payoffs in the n -player stag-hunt game. The column player represents the group average and the row player is any of the n players.	1008
4.E7	Game 4 from Binmore et al. (2001).	1009
4.E8	A hawk-dove game.	1009
4.E9	A Trust game in extensive form.	1010
4.E10	Two normal form games from Jehiel (2005).	1010
14.1	The two-player prisoner's dilemma game.	1042
14.2	Relation between replicator dynamics, evolutionary stable strategies, and Nash equilibria.	1051

14.3	Graph of the replicator dynamics for the PD game.	1054
14.4	Payoff matrix for the two-period tit-for-tat game.	1055
14.5	Dynamics of the tit-for-tat strategy in a repeated prisoner's dilemma game.	1057
14.6	Payoff matrix for the game of chicken.	1057
14.7	A plot of the dynamic equation of motion in the game of chicken.	1058
14.8	The rock-scissors-paper game.	1059
14.9	Orbits for the rock-scissors-paper game.	1061
14.10	A two-player coordination game.	1061
14.11	Coordination game in the two-population case.	1062
14.12	A prisoner's dilemma game.	1074
14.13	Prisoner's dilemma game, modified to take account of inclusive fitness.	1074
14.14	Prisoner's dilemma game, modified to take account of direct reciprocity.	1078
14.15	Indirect reciprocity and the prisoner's dilemma game.	1079
14.16	Fitness levels from each strategy in two different populations.	1080
14.17	A two-population, 2×2 strategic form game.	1085
15.1	Results of fitting the data from 5 simple classes of 2×2 games.	1104
15.2	A symmetric 2×2 coordination game.	1109
15.3	Non-convergence of probabilities under fictitious play.	1110
15.4	A symmetric two-player coordination game.	1111
15.5	A 2×2 stage game.	1112
15.6	A comparison of stated and fictitious play beliefs.	1113
15.7	Three stag-hunt games with varying levels of the optimization premium.	1116
15.8	Introducing complexity through the channel of asymmetric information.	1119
15.9	Aggregate behavior across all experiments over 40 rounds of play.	1120
15.10	The EWA cube.	1124
15.11	Original game (left) and transformed game (right) based on an aspiration level of 7 for the row player and 6 for the column player.	1135
15.12	Q values for various learning models. Stationary concepts depicted in dark bars and learning models are depicted in light bars.	1137
15.13	Q values of various models with and without loss aversion.	1138
15.14	A 3×3 coordination game.	1142
15.15	Choice frequencies over various rounds of play.	1143
15.16	A 5×5 strategic form game from Stahl (2000).	1147
15.17	Six normal form games used in the experiments of Grimm and Mengel (2012).	1150
15.18	Predicted shares of A choices by row player, based on logit regressions, in treatments M, MI, M+F, and M+G.	1152
15.19	A 2×2 coordination game.	1154
16.1	A 2×2 coordination game.	1161
16.2	Simulated plot of the transitions between the two states.	1161
16.3	A symmetric two-player coordination game.	1162
16.4	A 1-tree and 2-tree for the coordination game.	1164
16.5	Payoff matrix for the technology adoption game.	1165
16.6	1-trees (first row), 2-trees (second row), 3-trees (third row) for the technology adoption game.	1165
16.7	A symmetric, 2×2 normal form game.	1168

16.8	A 2×2 coordination game.	1170
16.9	A coordination game with a <i>better informed</i> player.	1173
16.10	A generic 3×3 coordination game.	1174
16.11	No strict risk dominant equilibrium.	1175
16.12	A 3×3 coordination game.	1175
16.13	Non-equivalence between risk dominance and stochastic stability	1177
16.14	Description of utilities from carrying gold (G) or silver (S).	1179
16.15	A 3×3 coordination game in which the two players are drawn from different populations.	1180
16.16	Hypothetical distribution of farms.	1185
5.A1	Stability of steady states.	1191
5.A2	Stability of the steady state in a first order non-linear differential equation.	1192
5.A3	j -trees for two recurrence classes with one state each.	1202
5.A4	1-tree (first row), 2-tree (second row), 3-tree (third row) for the case of three recurrence classes with one state each.	1203
5.E1	A 2×2 game.	1205
5.E2	The stag-hunt game.	1205
5.E3	A hawk-dove game.	1206
5.E4	Evolutionary stable strategies are not unbeatable.	1206
17.1	Skin conductance response (SCR) amplitude and heart rate (HR), 0–12 s after offer presentation onset.	1232
17.2	Schematic description of the “risk as feelings” thesis.	1235
17.3	Dynamics in the cue-dependent model.	1241
17.4	Maximum payment as a proportion of current value to obtain/avoid outcomes at selected time delays.	1243
17.5	Mean happiness and real household income for a cross section of Americans in 1994.	1274
17.6	Relation between mean happiness and real per capita GDP for the US between 1975 and 1997.	1275
17.7	The Easterlin paradox for Japan, over the period 1958–1991.	1276
17.8	Before-win and after-win happiness of people who unexpectedly win the lottery.	1277
17.9	Comparisons of own incomes with the incomes of others in a cross section of countries with different income levels.	1279
17.10	Relationship between mean life satisfaction and income, measured with Gallup World Poll data.	1280
17.11	The relationship between life satisfaction and income.	1281
17.12	The relationship between life satisfaction and per capita GDP for 95% of the world’s population.	1282
17.13	Relation between life satisfaction and income for the 25 largest countries.	1283
17.14	Non-parametric regressions of average health satisfaction against per capita GDP for seven age groups.	1285
17.15	The relationship between economic freedom, democracy, and well-being and happiness.	1286
18.1	Schematic representation of the dual process model.	1292
18.2	Willpower costs and optimal consumption.	1303

18.3	The optimal recall probability as β varies.	1319
19.1	Results for streaks of length 1–6 for all outside bets.	1354
19.2	Retail level data for Lotto Texas wins.	1356
19.3	Is the observed sequence with a given alternation rate more likely to have come from basketball or a coin toss?	1358
19.4	High price is salient (left panel) and low price is salient (right panel).	1372
19.5	Anchoring and insufficient adjustment.	1374
19.6	Determining the relative importance of the strength and the weight of evidence.	1382
19.7	A coordination game to study the role of aspiration adaptation.	1413
19.8	The problem of catching a ball.	1418
19.9	Classifying the degree of risk of suspected heart attack patients.	1418
19.10	Description of the take-the-best algorithm.	1421
19.11	Variations in demand due to a visible excise tax (left panel) and a shrouded sales tax (right panel).	1432
19.12	Predictions of a model of limited attention on used car values.	1442
19.13	Raw prices of cars based on data for 22 million cars sold during the sample period, 2002–2008.	1443
19.14	Determination of the optimal attention parameters as a function of the variance of the underlying parameter.	1446
20.1	Hedonic costs (light vertical lines) and hedonic benefits (bold shapes) for the two cases of prepayment and postpayment for vacations.	1467
20.2	The WTA of the three groups of students for increasing and decreasing durations of sounds.	1473
20.3	The relative optimality of food stamps and direct cash benefits.	1479
21.1	A plot of price and dividends against time.	1496
21.2	Cumulative average residuals for W and L portfolios of 35 stocks.	1497
21.3	Log deviations of prices from the predicted 1.5 : 1 ratio for Royal Dutch and Shell Transport.	1499
21.4	A basic description of the structure of the three-period model under professional arbitrage.	1509
21.5	Stability analysis in the Shleifer and Vishny (1997) model.	1515
21.6	Figure corresponding to the proof of Proposition 21.9.	1524
21.7	A decision tree depiction of the model of Gervais et al. (2011).	1535
7.E1	Average number of mazes solved.	1546
22.1	A 2×2 matching pennies game.	1607
22.2	Two possible preferences, $\succ_1$ and $\succ_2$, along with the corresponding attention filters, Γ_1 and Γ_2 .	1619
22.3	Tax incidence under limited attention.	1627
22.4	The excess burden of taxation under limited attention.	1630
23.1	A classification of the brain into four lobes: frontal, parietal, occipital, and temporal.	1649
23.2	A convenient classification of brain regions.	1650
23.3	A classification of the brain into three regions: forebrain, midbrain, and hindbrain.	1651
23.4	Structure of the basal ganglia and the approximate position of its various parts.	1652
23.5	Different layers of the cerebral cortex. © Anna Hine, Shutterstock	1652

23.6	The limbic system of the brain.	1653
23.7	Brodmann areas of the human brain.	1654
23.8	The structure of a neuron.	1654
23.9	A drawing showing the skull of Phineas Gage, after a tamping iron went through it.	1660
23.10	fMRI results for the gain and loss frames.	1664
23.11	The coefficient β in (23.2) (shown as the dots) for the left striatum (left panel) and the right striatum (right panel) for different probabilities.	1668
23.12	Panel A: Each choice involved choosing from a neutral or valenced slot machine. The decision and rewards time intervals are shown. Panel B: Probability distributions over rewards of the three slot machines, each represented by a different pattern. Panels C and D: Choices of subjects by disaggregated and aggregated treatment effects.	1670
23.13	Activation of various brain regions correlated with SV (first row), PE (second row), and R (third row).	1671
23.14	The behavioral evidence. Vertical bars denote standard errors of the mean.	1671
23.15	The fMRI evidence on activation of the ventral striatum for the contrast between own-transfers relative to partner-transfers for both groups.	1672
23.16	The fMRI evidence on activation of the vmPFC for the contrast between own-transfers relative to partner-transfers.	1672
23.17	Mean responses by player A when player B returns 0 MUs, and the actual payoff reduction of player B in various treatments.	1674
23.18	Peak activation in the caudate nucleus relative to the mean level in all four treatments.	1675
23.19	Different offers in an ultimatum game that signify varying levels of fairness.	1677
23.20	Outgroup effects captured by third party punishment of outgroup members who play opportunistically against ingroup members who cooperate.	1679
23.21	Ingroup effects captured by third party punishment of ingroup members who play opportunistically against outgroup members who cooperate.	1680
23.22	Mean rejection rates $\pm$ s.e.m. for different treatment groups and for different offers.	1681
23.23	Relative back transfers in the two treatment conditions with varying degrees of brain function disruption.	1682
23.24	Two different measures of social networks plotted against amygdala size for two age groups: young participants (black circles) and older participants (gray triangles).	1683
23.25	The behavioral data.	1688
23.26	Activation in the β areas and the δ areas in the brain as rewards with varying delays are presented to the subjects.	1688
23.27	Results in the p-beauty contest game. Behavioral results (panel A) and fMRI results (panel B). M is to be read as the probability p .	1691
23.28	Transfers in the trust experiments (panel a) and risk experiments (panel b).	1693
A.1	Best response functions for the battle of sexes game.	1709
A.2	Two examples of extensive form games.	1713
A.3	An illustration of sequential rationality with two examples.	1716
A.4	The signaling game in Brandts and Holt (1992).	1719

LIST OF TABLES

1.1	Drop in marginal utility for successive intervals in Rabin's calibration theorem.	105
2.1	Representative graphs of the Prelec function.	121
2.2	The fourfold classification of risk attitudes under prospect theory.	135
2.3	Estimates of loss aversion in various studies.	145
2.4	Estimates of the parameters of various probability weighting functions.	146
2.5	Various computations to explain the relevant empirical findings in different conditions.	204
3.1	Comparison of the predicted value of penalty rates under EU and PT.	248
3.2	Some simulations for optimal goals and effort levels.	254
3.3	A summary of the solution to the simple labor supply problem in Farber (2008).	258
3.4	Summary of the marginal rates of substitution between Y and h in various domains.	261
3.5	Summary of the results for rigid and flexible contracts.	272
4.1	Description of the four Savage acts in the three-colors experiment.	287
4.2	Results from the Fox and Tversky (1995) study. Standard errors in parantheses.	304
6.1	Parameter estimates of the Fehr and Schmidt (1999) model of inequality aversion.	403
6.2	Payoffs of three players in three different allocations.	432
6.3	Summary of predictions for player B.	437
7.1	Payoffs of the players in each treatment in Gneezy (2005).	456
8.1	The cost of effort function.	511
8.2	The relevant economic information in each state of the world.	522
9.1	Choices among temporal sequences of meals.	603
10.1	Examples of discount functions when there is a reference point for time.	637
10.2	The delay function, speedup function, and generating function corresponding to common discount functions.	640
11.1	Optimal consumption profiles for two values of the coefficient of relative risk aversion, 0.5 (columns 2–4) and 1.5 (columns 5–7).	655
12.1	Treatments 2–7.	726
12.2	Conditional probability of choosing action D.	730
12.3	Conditional probabilities of action D at each of the six nodes.	733
12.4	Experimental results.	739
12.5	Percentage of subjects choosing various strategies in the 4×4 and 6×6 games.	743
12.6	Results for the symmetric condition.	747
12.7	Frequencies of card choices in various treatments.	750
12.8	Results for the 2×2 pursuer–evader game.	756
12.9	Results for the 4×4 game.	757
12.10	Checking for primary and secondary salience.	760
12.11	Responses compiled from Table 2 of Mehta et al. (1994).	760
12.12	Payoff matrix for game-I.	764
12.13	Payoff matrix for game-III.	764
12.14	Payoff matrix for treatment A.	766
12.15	Experimental results.	768

12.16	Experimental results.	770
12.17	Payoff matrix for the continental divide game.	773
12.18	A median action game (payoffs in cents).	776
12.19	Distribution of fifth period group minimum in various minimum action games.	778
12.20	Frequency of Pareto optimal coordination in period 10.	780
12.21	Description of the four bargaining games.	790
12.22	Description of various bargaining games in Ochs and Roth (1989).	796
12.23	Type-contingent messages sent by senders and message-contingent actions of the receivers for the matchings in each of the parts, a, b, and c.	821
12.24	Type-contingent payoffs of the incumbent in the market entry game.	823
12.25	Payoffs of the entrant in the two treatments.	823
12.26	Payoff matrix in an entry deterrence game.	829
12.27	Predictions of the equilibrium actions of the weak monopolist and the entrant.	829
12.28	Type-contingent assets and end of period value of firms.	832
12.29	Equilibrium predictions of the signaling model in Cadsby et al. (1990).	833
12.30	Actual and predicted results of the signaling model in Cadsby et al. (1990).	834
12.31	Relative frequencies of cooperation, coordination, and efficiency.	841
12.32	A description of the various treatments in Cason and Sharma (2007).	846
12.33	Frequencies of observed outcomes at the aggregate level.	851
12.34	Buyers' valuations, sellers' costs, and market data.	865
12.35	Some of the data from a double auction experiment reported in Davis and Holt (1993, p. 129).	869
13.1	Game 8 from Costa-Gomes et al. (2001).	903
13.2	Equilibrium actions in a level- k model without communication.	909
13.3	Optimal choice of messages and actions in the two-stage game with communication.	910
13.4	Percentage of voters from each political party expecting their own candidate to win in the US presidential elections.	976
4.E1	Goeree and Holt's replication of the Beard-Beil (1994) findings.	1003
4.E2	A signaling game from Goeree and Holt (2001).	1004
14.1	Number of conjugal unions of each type.	1069
14.2	Number of offspring for each conjugal pair.	1069
15.1	A comparison of the results of alternative learning models.	1105
15.2	Summary of the treatments in Huck et al. (1999).	1113
15.3	Results averaged across 40 rounds of play. * denotes significantly different from the predicted value at 5% and ** denotes significantly different at 0.1% level.	1119
15.4	Hit rates and the best fits arising from an application of various learning models to different games.	1128
15.5	Estimates for the parameters of the sophisticated and adaptive EWA models. Standard errors in parentheses.	1130
15.6	Average defection period, $k = 1, 2, \dots, 10$, for experienced subjects.	1134
15.7	Fits of various learning models based on a comparison of the log likelihoods.	1148
15.8	Frequencies of various choices of actions in games 1 and 4.	1151
15.9	Sequence of actions and own-payoffs in a two-player coordination game.	1154
17.1	Responses to the income comparison question.	1279
19.1	Assessed probability that the next flip of a coin will be heads (H), given a sequence of three flips; each flip is either heads (H) or tails (T).	1353
19.2	Violations of the conjunction fallacy.	1363

19.3	Empirical evidence for the <i>availability</i> and <i>affect</i> heuristics.	1367
19.4	Percentage of responses consistent with Bayes' rule in different empirical studies.	1377
19.5	Probability that the coin is head-biased, conditional on the observations.	1380
19.6	Mean probabilities assigned to each of four outcomes, 1–4, in the four events, A–D.	1385
19.7	Median values across decision makers.	1390
19.8	Different usages of the System 1–System 2 terms.	1399
19.9	The cue values for each of the four cities.	1419
21.1	States and transition probabilities in the mean reverting model (M).	1521
21.2	States and transition probabilities in the trending model (T).	1521
21.3	Markov transition probabilities for the two states of the world.	1522
21.4	Simulations for the Barberis et al. (1998) model.	1525
7.E1	A summary of the composition of two different dinner sets.	1545
7.E2	Two different problems given to two groups of subjects in a between-groups design.	1545
22.1	Fraction of subjects who pay the fee in each treatment.	1608
8.E1	Payoffs in state 1.	1635
8.E2	Payoffs in state 2.	1636
8.E3	Payoffs for each pair of actions and psychological states.	1636

Introduction

The *neoclassical framework* in economics provides a coherent and internally consistent body of theory that offers rigorous, parsimonious, and falsifiable models of human behavior. Augmented with auxiliary assumptions, it is flexible enough to analyze a wide range of phenomena and its qualitative predictions often accord with one's gut feeling about many phenomena. For instance, under suitable assumptions, consumers buy less when the price is higher, and more when their incomes are higher; socially undesirable activity typically falls when the probability of detection and fines is higher; workers often respond positively to monetary incentives; in many cases, changes in tax rates and devaluations have had the predicted effects on economic activity; banning certain activities often drives them underground.

Throughout the book, I avoid the loaded term *standard economics* to refer to *neoclassical economics* because this might give the latter a certain empirical sanctity. In actual practice, the neoclassical framework includes, but is not restricted exclusively to, consistent preferences, subjective expected utility, Bayes' rule to update probabilities, self-regarding preferences, emotionless deliberation, exponential discounting, unlimited cognitive abilities, unlimited attention, unlimited willpower, and frame and context independence of preferences. Neoclassical economics is also typically underpinned by optimization based solution methods and an equilibrium approach.

In principle, the neoclassical framework is capable of relaxing many of its standard assumptions. For instance, it can allow for reference dependence preferences, social preferences, frame dependent preferences, and non-exponential models of discounting. However, these extensions are rare in actual practice, and when they are made, the neoclassical framework typically does not have fundamental new insights to offer. For instance, adding reference dependent preferences generates few, if any, insights in the absence of a theory about how human behavior differs in the domains of gains and losses relative to a reference point. Similarly, adding other-regarding preferences without attempting to fit such a model to the behavior of humans, particularly to the evidence from experimental games, offers little progress. For these reasons, my use of the term *neoclassical economics* in this book, is shorthand for *the typical practice in neoclassical economics*.

The intellectual developments in neoclassical economics that preoccupied some of the best brains in economics in the twentieth century are truly impressive. However, its empirical success in predicting and explaining human behavior is modest. Indeed, an impressive thorough and detailed body of experimental, neuroeconomic, and field evidence, based on several decades of work, raises serious concerns about the core assumptions and predictions of neoclassical models. This has been matched by impressive theoretical developments, drawing on insights from psychology, biology, anthropology, sociology and other social sciences, that has come to be known as *behavioral economics*. These models have had much greater empirical success relative to neoclassical models.

Increasing the explanatory power of neoclassical economics is very worthwhile in its own right, but Thaler (2015) adds another reason for studying behavioral economics in his inimitable style: “Behavioral economics is more interesting and more fun than regular economics. It is the undismal science.”

This book offers an account of formal behavioral economic theory, its applications, and a discussion of the underlying experimental and field evidence.¹

I do not feel it is necessary to enter into a very detailed discussion, often quite fruitless, about the definition and scope of behavioral economics. There is a danger that one may propose definitions that are either too broad and have ambiguous scope, or too narrow with limited scope; each of these outcomes would be unfair for a newly emerging field. Any falsifiable theory that replaces/modifies any of the core features of neoclassical economics, by alternatives that have a better empirical foundation in human behavior is a potential member of the class of behavioral economic theories, if it can pass stringent empirical tests.

The standard toolkit in neoclassical economics is adequate for the study of behavioral economics, which, in many cases, offers a conservative, yet empirically critical, extension to neoclassical models. Most behavioral models adopt the same underlying optimization framework, are typically underpinned by axiomatic foundations, are parsimonious, rigorous, falsifiable, and internally consistent. I use the word “rigorous” purely for its practical appeal to most neoclassical economists but I agree with the sentiments expressed by Gintis (2009, p. xviii): “The economic theorist’s overvaluation of rigor is a symptom of their undervaluation of explanatory power. The truth is its own justification and needs no help from rigor.”

Let us note three points upfront lest there be any confusion. First, behavioral economics is an enhancement of neoclassical economics to take account of more empirically supported evidence on human behavior, and not its antithesis. Second, there is no paradigmatic battle between behavioral economics and neoclassical economics. As in every science, we progress by taking account of evidence that suggests a refinement and improvement of existing models. In this case, the relevant improvement appears to have the steepest gradient in the direction of constructively incorporating insights from other social sciences, but most notably from psychology. Third, although the book provides a great deal of evidence that is contrary to many of the core assumptions of neoclassical economics, it is not about simply bashing inadequate neoclassical assumptions and predictions. Rather, it offers a constructive way forward by highlighting developments in behavioral economic theory and suggesting a research program in this direction. My personal hope is that behavioral economics ceases to exist as a separate field within economics, and this becomes the normal way in which we do economics.

A distinction is sometimes drawn between experimental economics and behavioral economics.² However, the activity of behavioral economists and experimental economists is complementary and collaborative. It is often difficult to spot the dividing line between their work. For instance, experimental economists not only test the predictions of economic models, but their results have often been critical in suggesting further developments in behavioral models. Behavioral theorists on the other hand, often suggest experiments that could test their proposed theories.

¹ For a non-technical treatment of behavioral economics, the reader can consult the extremely readable and witty account by Thaler (2015) that offers a much more detailed historical account of developments in behavioral economics from the 1970s onwards from a personal perspective.

² Loewenstein (1999) gives a nice discussion of the methods in each of these areas and offers the following definition (p. F25): “BEs [behavioral economists] are methodological eclectics. They define themselves, not on the basis of the research methods that they employ, but rather their application of psychological insights to economics. In recent published research, BEs are as likely to use field research as experimentation . . . EE [experimental economists] on the other hand, define themselves on the basis of their endorsement and use of experimentation as a research tool.”

The introduction is intended to be a substantial stand-alone chapter in its own right that is integral to the book. Its purpose is to provide essential material and insights that sets the stage for the rest of the book. Perhaps, it ought also to be assigned as essential background reading in any advanced microeconomics course. Section 1 traces some of the historical developments that led to modern behavioral economics. Section 2 considers a descriptive and prescriptive approach to how economists practice their craft. Section 3 considers the role of the experimental method in behavioral economics. Section 4 considers the approach adopted by the book, and offers brief comments about the organization of the book. Section 5 outlines five theoretical ideas in behavioral economics, locating them within the larger literature, and evaluating them against the empirical evidence. Section 6 considers the empirical evidence on five interesting questions asked by, or of, behavioral economics. There are two appendices. Appendix A outlines the random lottery incentive mechanism that lies at the heart of the modern experimental method in economics. Appendix B asks you to think of 50 questions as a problem set, but I deliberately give you very little structure at this stage in order to enable a free-spirited approach to the answers. Rigorous answers to these questions can be found in the book.

1 The antecedents of behavioral economics

While Adam Smith's justly celebrated book, *The Wealth of Nations*, is widely cited, his other book, *The Theory of Moral Sentiments*, has received less attention. *The Theory of Moral Sentiments* reads like an agenda for modern behavioral economics; it recognizes many behavioral phenomena such as loss aversion, altruism, emotions, willpower, and the planner–doer framework (Ashraf et al., 2005). Classical economists such as Jeremy Bentham wrote about the psychological underpinnings of utility and Francis Edgeworth wrote about social preferences (Camerer and Loewenstein, 2004). Bardsley et al. (2010) trace the beginnings of experimental economics to the classical economists such as David Hume, Stanley Jevons, and Francis Edgeworth; Jevon's marginal utility analysis derived its motivation from experimental observations about the relation between stimuli and sensations.

Two factors contributed to the gradual elimination of psychology from economics. First, around the turn of the twentieth century, there was “a distaste for the psychology of their period, as well as the hedonistic assumptions of Benthamite utility” (Camerer and Loewenstein, 2004). The second was the revealed preference approach popularized by Paul Samuelson that emphasized the observation of *choice behavior* rather than the psychological foundations for choice behavior (Bruni and Sugden, 2007). Glimcher and Fehr (2014, p. xviii) write: “It cannot be emphasized enough how much the revealed-preference view suppressed interest in the psychological nature of preferences, because clever axiomatic systems could be used to infer properties of unobservable preference from choice.”

Important, and path-breaking, developments in behavioral economics took place in the 1950s and 1960s that included: violations of the independence axiom of expected utility theory (Allais, 1953); violations of subjective expected utility (Ellsberg, 1961; Markowitz, 1952); demonstration of the importance of bounded rationality (Simon, 1978; Selten, 1998);³ and early work on quasi hyperbolic discounting (Phelps and Pollak, 1968). However, at that time, this work struggled to get the attention that it deserved.

³ Simon (1978) refers to Herbert Simon's Nobel lecture that traces the historical development of bounded rationality through the 1950s and 1960s. Selten (1998) is an English language version of a paper that appeared initially in German in 1962.

An important catalyst for the development of behavioral economics was the decline of the behavioralist school in psychology, and the emergence of cognitive psychology. Cognitive psychology emphasized the role of mental processes in the understanding of tasks involving decision making, perception, attention, memory, and problem solving. Some cognitive psychologists naturally turned their attention to testing their models against the neoclassical framework. The two most important cognitive psychologists in this category were Daniel Kahneman and Amos Tversky, whose work in the 1970s helped kick-start modern behavioral economics. Along with Richard Thaler, who was an economist by training, and was struggling to make sense of several anomalies in neoclassical economics from the mid 1970s onwards, they are some of the earliest modern behavioral economists.

Before I discuss some of the seminal ideas and insights of modern behavioral economics,⁴ in Sections 5 and 6, I take a detour into two topics without which it could be difficult to fully appreciate the behavioral framework.

The first digression is on the *role of methodology in economics* that I consider in Section 2. In my defense, I cannot offer a better justification for this digression than the following quote from Gintis (2009, p. xvi) that nicely captures the essence of the problem:

Economic theory has been particularly compromised by its neglect of the facts concerning human behavior. . . . I happened to be reading a popular introductory graduate text on quantum mechanics, as well as a leading graduate text in microeconomics. The physics text began with the anomaly of blackbody radiation, . . . The text continued, page after page, with new anomalies . . . and new, partially successful models explaining the anomalies. In about 1925, this culminated with Heisenberg's wave mechanics and Schrödinger's equation, which fully unified the field. By contrast, the microeconomics text, despite its beauty, did not contain a single fact in the whole thousand-page volume. Rather the authors built economic theory in axiomatic fashion, making assumptions on the basis of their intuitive plausibility, their incorporation of the "stylized facts" of everyday life, or their appeal to the principles of rational thought. . . . We will see that empirical evidence challenges some of the core assumptions in classical game theory and neoclassical economics.

The second topic is the *role of experimental evidence in economics* that I consider in Section 3. The justification for this section is the continued skepticism of many economists about experimental economics, which constitutes an important part of the evidence base for behavioral economics. The following quote attributed to the Nobel Prize winner Gary Becker from a magazine interview (Camerer, 2015, p. 250) is probably not unrepresentative: "One can get excellent suggestions from experiments, but economic theory is not about how people act in experiments, but how they act in markets. And those are very different things. That may be useful to get suggestions, but it is not a test of the theory. The theory is not about how people answer questions. It is a theory about how people actually choose in market situations."

What follows is a somewhat long introduction, but this is a somewhat long book too. In mitigation, the first one third of the introduction largely deals with background material that reflects the somewhat unsettled nature of economics. My hope is that if a second edition of this book is ever written, then there would be enough convergence of views on this material, so that I can safely omit it.

⁴ I must recommend to the reader two must-read non-technical, but intellectually rich and deep, accounts of some of the central ideas in behavioral economics and the modern history of the subject. The first is Kahneman (2011) that focuses mainly on judgment and decision making and the second is an incredibly entertaining personal account by Thaler (2015) that covers a wide range of topics such as fairness, mental accounting, and behavioral finance.

2 On methodology in economics

The actual practice of behavioral economics is influenced, directly or indirectly, by Popperian views on methodology (Popper, 1934, 1963). Popper begins by distinguishing between science and non-science. A scientific hypothesis must be falsifiable in the sense that it must specify the conditions under which the hypothesis can be rejected. Further, one can only refute theories but never prove that they are true. For instance, the observation of a million white swans is consistent with the hypothesis that “all swans are white” but does not prove that the hypothesis is true; for the very next observation could be a non-white swan.

The best recipe for the advancement of science, in the Popperian view, is to subject scientific hypotheses to stringent testing, i.e., expose the hypotheses to tests that are most likely to reject them. In the strict Popperian view, one observation that is contrary to a hypothesis rejects it. For instance, a single observation of a black swan rejects the hypothesis that all swans are white. Science progresses by advancing a new hypothesis that explains everything that a rejected hypothesis explained, but, in addition, it explains some new phenomenon that the rejected hypothesis could not. For an application of the Popperian position to economic contexts, see Blaug (1992), Hausman (1992), and Hands (2001).

One concern with the Popperian approach is that a test of a hypothesis is a joint test of the hypothesis and several auxiliary assumptions. Thus, a rejection may arise because the hypothesis is incorrect, or the auxiliary assumptions might have been rejected, or both; this is known as the *Duhem–Quine thesis* (DQT). For instance, in an experimental test that rejects mixed strategy Nash equilibrium, one might wonder if the rejection was caused by one of several confounding factors, such as an inappropriate subject pool, unclear experimental instructions, and inadequate incentives.

While the Popperian position is *prescriptive* (how should we best do science?), a *descriptive* view (how is science actually done?) was offered by Kuhn (1962). Kuhn noted that knowledge in science does not accumulate in a linear manner. He highlighted, instead, the role of periodic revolutions in science, or an abrupt transformation in the existing worldview, a *paradigm shift*. He distinguishes between three phases in the development of any science. In Prescience, there is no central paradigm, but there is an attempt to focus on a set of problems. In normal science, the longest of the three phases, there is the establishment of a central paradigm, great progress is made in answering many of the questions posed during Prescience, and much success is achieved in answering new questions. In a departure from the Popperian prescriptive position, in this phase, rejections of the paradigm are robustly challenged or ignored, and belief in the paradigm is unshakable. However, as anomalies gradually begin to accumulate, and reach a tipping point, a crises takes place in the paradigm. There is a sudden paradigm shift and a new paradigm that subsumes the old paradigm takes its place.

One prescriptive response to the DQT and to Kuhn’s descriptive ideas, while retaining a Popperian approach, was proposed by Lakatos (1970) under the name: *The methodology of scientific research programs* (MSRP). Lakatos distinguished between a set of non-expendable statements or assumptions, which is the *hard core of a research program*, and a set of expendable auxiliary assumptions. In a distinctly non-Popperian recommendation, but reminiscent of the normal science phase of Kuhn, the hard core is insulated from refutation; this also addresses the DQT. For instance, Newtonian physics has a hard core that comprises of the three laws of dynamics and a law of gravitation. Any refutation of the research program, in this phase, is then ascribed to a failure of the auxiliary assumptions, which are modified to explain the refutation.

One potential defense of this approach is that it allows for a period of time for the development of a new research program that can take account of the emerging refutations. However, a practical downside could be that proponents of a research program might engage in defensive methodology for far too long, and resist the development of a new research program that has a different hard core. To take account of this possibility, Lakatos termed a research program as *theoretically progressive*, if refinements that take place by altering auxiliary assumptions but not the hard-core, lead to the explanation of existing anomalies and to novel predictions. A research program is *empirically progressive*, if the novel predictions are not refuted. Adherence to a hard core is only admissible if research programs are theoretically and empirically progressive.

Eventually anomalies play the most important part in giving rise to new research programs; Lakatos noted that all theories are born into and die in a sea of anomalies.⁵ The reader may find below that the actual practice in behavioral/experimental economics appears to be closer to the Lakatosian view than the Popperian view.⁶ For instance, in decision theory, the hard core may be thought to comprise of completeness, transitivity, and first order stochastic dominance (Bardsley et al., 2010, p. 129). Indeed, neither expected utility theory nor the main behavioral alternatives such as rank dependent utility, theory of disappointment aversion, or prospect theory, are willing to relax the assumption of well-behaved preferences. This makes it difficult for most decision theories to explain framing effects, although prospect theory is potentially able to capture framing effects through changes in the reference point.

In light of this brief discussion on methodology, let us return to the “neglect of the facts concerning human behavior in economics” that Herbert Gintis highlights above. Why should such a situation have arisen? In order to understand this, consider the following three representative views, written by some of the leaders in neoclassical economics.

Dekel and Lipman (2010, p. 264) write: “Hence the choice of a model will depend on the purpose for which the model is used, the modeler’s intuition, and the modeler’s subjective judgment of plausibility One economist may reject another’s intuition, and, ultimately, the marketplace of ideas will make some judgments.”

Gilboa et al. (2014, p. F. 516) write: “In particular, we agree that: economic models are often viewed differently than models in the other sciences; economic theory seems to value generality and simplicity at the cost of accuracy; models are expected to convey a message much more than to describe a well-defined reality; these models are often akin to observations, or to gedankenexperiments; and the economic theorist is typically not required to clearly specify where his model might be applicable and how.”

Rubinstein (2006, p. 882) writes: “As in the case of fables, models in economic theory are derived from observations of the real world, but are not meant to be testable. As in the case of fables, models have limited scope. As in the case of a good fable, a good model can have an enormous influence on the real world, not by providing advice or by predicting the future, but rather by influencing culture. Yes, I do think we are simply the tellers of fables, but is that not wonderful?”

Modern economics has been heavily influenced by the *instrumental position* taken by Friedman (1953). Friedman argued that we should not judge economic theories by the realism of its assumptions but rather, by the accuracy of its predictions. He writes (p. 14): “Truly important and

⁵ Closer to home, economists would remember the influential *anomalies feature* that Richard Thaler wrote for the *Journal of Economic Perspectives* from 1987 to 2006. Indeed, in the very first piece, Thaler, keenly aware of methodological issues, quoted from Thomas Kuhn.

⁶ For a critique of the Lakatosian approach as applied to economics, see Hands (1991) and De Marchi and Blaug (1991).

significant hypotheses will be found to have ‘assumptions’ that are wildly inaccurate descriptive representations of reality, and, in general, the more significant the theory, the more unrealistic the assumptions To be important, therefore, a hypothesis must be descriptively false in its assumptions” And shortly thereafter (p. 15) he writes: “To put this point less paradoxically, the relevant question to ask about the ‘assumptions’ of a theory is not whether they are descriptively ‘realistic,’ for they never are, but whether they are sufficiently good approximations for the purpose in hand.”

A natural progression of Friedman’s position can be found in Gilboa et al. (2014, F. 514): “Why does economic theory engage in relatively heavy technical analysis, when its basic premises are so inaccurate? Given the various violations of fundamental economic assumptions in psychological experiments, what is the point in deriving elaborate and carefully proved deductions from these assumptions? Why do economists believe that they learn something useful from analyzing models that are based on wrong assumptions?” Their answer to these question is based on an identification of economic models with *case-based reasoning* rather than *rule-based reasoning*. Rule-based reasoning requires the formulation of general rules or theories. In contrast, case-based reasoning requires one to draw inferences based on similar past cases. The purpose of economic models, in this view, is to add to the bucket list of cases and analogies that can be used to draw inferences now, or at some point in the future.

These views give a fair bit of insight into contemporary thinking in economics about how we should go about practicing our craft. I also believe that acceptance of these views is widespread in the economics profession and many economists challenged on these views are outrightly dismissive. Initial intuition about economic models, whether motivated by existing empirical evidence, or a desire to make novel predictions, must begin from somewhere. Here, the role of initial conjectures as parables, useful stories, or fables to inform one’s intuition about better and more complete models is surely important. But this cannot be the justification for continued reliance on a set of models that have faced persistent refutation, or to wish to shield them from refutation by seeking a special status for them.

Indeed, many of the contemporary methodological views in economics are retrogressive and a license to engage in defensive methodology to protect the status quo. Friedman’s approach has been much misused in economics. Consider the following entirely reasonable description of Friedman’s approach to *model building* (as distinct from evaluating theories) in Gintis (2015, p. 223) that this book concurs with: “The goal of model-building [is] to create a tractable analytical structure, analyze the behavior of this structure, and test the fruitfulness of the results by comparing them with empirical data.”⁷

Consider “normal” practice in physics; I encourage the reader not to judge natural sciences by a few well-publicized outliers. In a letter to the *London Times*, dated November 28, 1919, Albert Einstein described his *theory of relativity* in comparison to *Newtonian physics*, to a lay audience. Einstein mentioned two predictions of his theory that had been confirmed (both in domains where his theory was most likely to fail, hence, these are “stringent tests”): (1) Revolution of the ellipses of the planetary orbits round the sun, which was confirmed for Mercury. (2) The curving of light rays by the action of gravitational fields. He mentioned one prediction that had not yet been confirmed (displacement of the spectral lines toward the red end of the spectrum in the case of light transmitted to us from stars of considerable magnitude); indeed at the time Einstein published the theory of relativity, it was not even clear how to test this prediction. Einstein then

⁷ Readers interested in pursuing this approach further can consult Godfrey-Smith (2006, 2009) and Wimsatt (2007).

wrote (p. 4): “The chief attraction of the theory lies in its logical completeness. If a single one of the conclusions drawn from it proves wrong, it must be given up.”

I invite the reader to pause for a moment to compare Einstein’s approach with the “mainstream” views in economics that I have outlined above. Indeed, as Bardsley et al. (2010, p. 8) note: “But it is surprisingly common for economists to claim that the core theories of their subject are useful despite being disconfirmed by the evidence.”

The tendency to discount experimental evidence in economics, despite its growing importance and prominence, is probably an important factor in ignoring “uncomfortable” evidence in economics; see Section 3 below. Another important factor is that Friedman’s instrumental position has been used as a license by some to make ad-hoc auxiliary assumptions, and others to genuinely believe that their assumptions are literally true in an “as if” sense. Any empirical rejection of the “as if” assumptions is often rejected on the grounds that the evidence is flawed, untrustworthy, based on dubious experimental methods, or lacks external validity. This is a form of defensive methodology that is inimical to the progress of economics, and I urge the reader to resist it.

Behavioral economics offers an easier resolution of the “as if” approach. There is now compelling evidence (see Sections 5, 6) which shows that some of the central tenets of neoclassical economics are neither true in an “as if” sense, nor are their predictions always satisfactory when subject to stringent tests.

Schotter (2015) offers the following critique of Friedman’s position. Suppose that assumptions x , y , and z lead to some theory T . Suppose also that one or more assumptions are violated by the empirical evidence, yet T makes a successful prediction. Then there are three possibilities. (1) The violated assumptions are superfluous for the theory, at least in the context where the theory was tested. (2) The violated assumptions counteract each other perfectly, so they do not affect the prediction. (3) The successful prediction is a fluke. Conversely, if the assumptions are correct and the model is complete then we expect T to make successful predictions anyway. Thus, it is difficult to justify a theory based on patently false assumptions. Schotter (2015, p. 63) observes, correctly: “after all, the assumptions are the theory.”

My colleague, Ali al-Nowaihi, likes to give the following example that applies to birds who cannot swim (e.g., gannets can swim, so they are excluded). Birds fly, so one may theorize that they behave “as if” they understand the laws of aerodynamics. This is an admissible hypothesis, but then one must test the “as if” assumption. Given that air is basically a fluid, so birds might also be assumed to know the laws of hydrodynamics. If the “as if” presumption were true in this case, then birds released under water should try to swim, but they actually try to fly, and drown. Thus, the original “as if” supposition is false. If the “as if” assumptions are not tested properly, then we can never have any degree of confidence in the models based on these assumptions.

A common view in economics (shared unfortunately by some behavioral/experimental economists, I must add) appears to be that there is something rather difficult and unique about testing economic theories, relative to the natural sciences. So, at least implicitly, the argument goes, one needs to accord a “special status” to economic theories. Consider the following representative quote from Richard Lipsey’s wonderful introduction to economics [Lipsey, 1979, p. 8] cited in Bardsley et al. (2010, p. 6–7) that, I suspect, many economists would agree with: “Experimental sciences, such as chemistry and some branches of psychology, have an advantage because it is possible to produce relevant evidence through controlled laboratory experiments. Other sciences, such as astronomy and economics, cannot do this.”⁸ This mainstream view is contestable

⁸ And here is a similar view from a celebrated text in economics (Samuelson and Nordhaus, 1985, p. 8): “Economists (unfortunately) . . . cannot perform the controlled experiments of chemists or biologists because they

and there appears to be a misunderstanding about the relative difficulty of testing theories in the natural sciences and in economics.

The view that testing of theories is somehow easy or easier in the natural sciences, as compared to economics, must surely be deeply offensive and insulting to experimenters in the natural sciences. The Higgs boson or Higgs particle was proposed by British physicist Peter Higgs in the early 1960s, and it took 50 years of incredibly hard efforts to confirm the particle in 2013. Particle physicists did not seek a *special status* for this theory that could insulate it from rejection. The enormously high energies required to test for the Higgs particle required the construction of a very expensive and complex experimental facility, CERN's Large Hadron Collider, that eventually confirmed the theory. Note also that Peter Higgs was made to wait 50 odd years and given the Nobel Prize in physics only after his theory was confirmed. He was not given the Nobel on any of the following criteria: elegant and beautiful theory, useful model that helped the intuition of particle physicists, or a fable or useful story that aids in the understanding of how the universe began.

Astronomers who dealt with the question of the distance of earth from distant objects, or the chemical composition of stars that are millions of light years away, did not also seek a special status for their subject. They got on with the difficult job of seeking the relevant measurements, often using indirect evidence and clever implications of theory. They were eventually successful after several decades of work. Are economists seriously arguing that their measurement problems are more difficult than the problems in the natural sciences? Cosmic microwave background radiation was first proposed in 1948, but experimentally confirmed due to an accidental discovery in 1964. DNA was first isolated in 1869, but it took the most part of a century to find the double-helix structure of DNA, and confirm it by experimental evidence in 1953. The germ theory of disease was proposed in the mid sixteenth century, yet confirmation of the theory occurred in the seventeenth century. The pool of such examples is very large. The process of discovery, measurement, and of testing the theory, can be a long and arduous one; seeking a special status for the subject is defeatist and put bluntly, lazy.

Economists opposed to lab/field data are likely to argue that the behavior of humans is too noisy, heterogenous, and fickle, which is not a problem in the natural sciences (e.g., atoms are, after all, not subject to mood swings). This overstates the degree of difficulty in testing economic theories, relative to those in the natural sciences on at least two grounds.

1. Experimental economics has discovered systematic human behavior in many of the most important domains in economics. A small sample includes reference dependence, loss aversion, non-linear probability weighting, conditional cooperation, intention-based reciprocity, present-biased preferences, and importance of emotions such as regret, guilt, and disappointment. These behaviors are also underpinned by neuroeconomic evidence. Replication of experimental results is routine, and if similar subject pools and protocols are used, experiments produce replicable data. Examples are results from double auction experiments, and a range of games that demonstrate human prosociality, such as the ultimatum game, the gift exchange game, the trust game, and the public goods game; these examples can be multiplied many fold, as the results in this book attest.
2. If indeed human behavior is inherently too noisy and heterogenous, then economic theory needs to focus more efforts in this direction. When Brownian motion was discovered in 1827 by Robert Brown, in the behavior of pollen grains, physicists did not throw up their arms in despair. Important work in the late part of the nineteenth century, and by

cannot easily control other important factors. Like astronomers or meteorologists, they generally must be content largely to observe."

Einstein in the early 20th century, paved the way for describing not only the mathematics of Brownian motion, but also predicting the probability distribution of particles in Brownian motion. Perhaps, in an analogous manner, economic theories need to predict the probability distribution of economic behavior, which can then be tested in experiments.

Experimental economics in the lab, and in the field, has made enormous progress in developing new econometric techniques for small samples, and in novel experimental methods. It has also deeply enhanced our understanding of human behavior and allowed for stringent testing of economic theory. This progress is inconsistent with the view that we should grant a special status to economic theories that exempts them from careful and stringent testing. The differences in experiments in economics and the natural sciences are much smaller relative to the differences in attitudes and institutions in the two fields of study. Progress in economics will be substantially enhanced if we learn from best practice, and give up our implicit demand for special status.

3 The experimental method in economics

Work on experiments in behavioral economics gained particular momentum following the seminal work of Daniel Kahneman and Amos Tversky in the 1970s. However, a number of important experiments in economics were also conducted in the late 1940s, the 1950s, and the 1960s. These include Richard Chamberlain's testing of general competitive equilibrium (Chamberlain, 1948); Maurice Allais's work on demonstrating violations of the independence axiom in expected utility theory (Allais, 1953); Vernon Smith's work on induced value elicitation and double auction experiments in competitive settings (Smith, 1962); and Sidney Siegel's experiments on bargaining (Siegel and Fouraker, 1960). Other prominent figures who were either involved in experimental economics, or expressed an interest in it during the 1950s and 1960s included Ward Edwards, Reinhard Selten, Martin Schubik, Herbert Simon, Charles Plott, Donald Davidson, and Pat Suppes; for a brief historical sketch see Guala (2008) and Bardsley et al. (2010).

Experimental economics is now mainstream by most yardsticks, particularly in terms of its presence in peer reviewed journals in economics. In his early surveys on experimental economics, Roth (1987, 1988) hoped that experimental economics would perform three kinds of functions. Speaking to theorists (testing economic theory), searching for facts (generating novel empirical regularities that could be modeled by subsequent theory), and whispering in the ears of princes (offering reliable policy advice). Roth (2015) takes stock of experimental economics on these criteria and finds that it is thriving. One of his case studies, on bargaining behavior, is outlined in detail in Part IV of the book.

At one level, there has been a complete denial of the usefulness of experiments in economics. Friedman (1953, p. 10) views the domain of empirical testing in economics to be naturally occurring field data: "Unfortunately, we can seldom test particular predictions in the social sciences by experiments explicitly designed to eliminate what are judged to be the most important disturbing influences. Generally, we must rely on evidence cast up by the 'experiments' that happen to occur." A modern critique of the experimental method in economics is offered by Levitt and List (2007). They list several objections to experimental results that I address in subsequent sections.

(1) Participants in experiments are subjected to unprecedented experimental scrutiny. Since subjects may perceive that they are being watched over by the experimenter, they may give responses that the experimenter really desires (*experimenter demand effects*; see Zizzo, 2010) or

they may not reveal their true underlying preferences. For instance, they worry that participants may engage in more prosocial behavior than they really intend to.

Whilst I reserve my detailed responses to later sections, I find it somewhat curious that if subjects are accused of being influenced by experimenter demand effects, say out of guilt or shame, then they appear to exhibit social preferences (or beliefs may directly enter their utility functions, as in psychological game theory), which is precisely what is being disputed by the critics.

(2) In actual practice, human decisions are context-dependent and influenced by cues, social norms, and past experiences. It is not clear that experiments can capture these factors. For instance, participants in experiments may import an inappropriate “outside context” into their responses in experiments.

(3) Actual human behavior is strongly affected by stake sizes in experiments. Experiments are typically conducted with small stakes, so they might not capture the richness of human behavior that arises from varying stakes.

(4) There could be self-selection biases caused by student volunteers who might be particularly prosocial, younger, more educated, and have a higher need for approval, as compared to the average human population. In contrast, people who self-select themselves into real market situations, might be particularly suitable to do well in real markets.

(5) Choice sets in experiments might be particularly restrictive relative to the real world. For instance, there could be more prosocial options in experiments relative to the real world.

(6) The results of lab experiments may generalize poorly to real-world behavior for all of the reasons mentioned in (1) through (5), above. This issue of *external validity* of lab experiments is the main concern raised by the authors who write (p. 170): “Perhaps the most fundamental question in experimental economics is whether findings from the lab are likely to provide reliable inferences outside of the laboratory.”

This discussion briefly encapsulates the modern case against experimental economics. Let us now briefly examine these claims.

3.1 Experiments and internal validity

Experiments allow for unprecedented control over the economic environment, hence, they have high *internal validity*, which is critical for stringent tests of economic theories. Internal validity is reduced when there are, for instance, selection issues, confounds in treatments, and unclear experimental instructions, all of which are carefully addressed in modern experimental work. Thus, in well-conducted experiments, the complicated identification strategies of field studies can be replaced by clever and much simpler experimental design.

For instance, suppose that a researcher is interested in testing if higher wages elicit higher effort in a firm; this is known as a *gift exchange game*. A field experiment is likely to be influenced by strategic behavior and reputational concerns of the workers and firms; field experiments in general, are likely to have lower internal validity. However, in a lab experiment, these factors are easily controlled, allowing one to cleanly separate the relation between a fair-wage and effort. The high degree of experimental control in lab experiments allows for replication of lab results. For the converse reason, the results of field experiments are more difficult, and sometimes impossible to replicate when one is given access to a unique field environment.

Experiments can also test the predictions of theory in a parameter space that might be difficult to observe in the field. This is similar to extreme stress tests of aircraft frames under conditions that are not normally encountered in the actual operation of the aircraft, or the exposure of bridge designs to extreme environmental conditions. In a nutshell, all this allows for more stringent tests

of economic theory. Experiments are sometimes criticized on the grounds that the sample sizes are small. Falk and Heckman (2009) term this issue as a “red herring” on the grounds that there have been important developments in small sample econometrics, and many experiments do use large subject pools.

Camerer (2015) argues that there is no evidence of experimenter demand effects, despite the suspicion that there might be such effects; see also his discussion of the alternative interpretations of experimenter demand effects in Hoffman et al. (1998). There are several reasons why experimenter demand effects may be weak or non-existent. Such demand effects require two conditions. First, subjects must know the experimenter’s preferred hypothesis. Second, they should be willing to sacrifice their own experimental earnings in order to favor the experimenter’s preferred hypothesis.

On a-priori grounds, arguably, it is often quite hard for subjects to know the experimenter’s preferred hypothesis. This arises particularly when (i) experimental instructions are carefully worded to prevent any such inference, and (ii) the experimenter might not be sure which of the competing hypotheses actually hold. However, if subjects can somehow guess the preferred hypothesis, then stakes can be raised to levels where they are too difficult to sacrifice for the sake of pleasing the experimenter. However, in most cases, the results with high stakes are not dramatically different from those with modest stakes (Camerer and Hogarth, 1999).⁹ In three preference reversal experiments, Lambdin and Shaffer (2009) find that the percentage of subjects who were successfully able to guess the preferred hypothesis of the experimenter was 7%, 32%, and 3%.

The degree of anonymity in lab experiments can be varied, so it is an ideal environment to test for the effects of variation in the degree of anonymity (Bolton et al., 1998). One’s actions are often observed by others in real-world situations, and in many field situations, where controlling for such scrutiny, and varying its level is typically even more difficult. The criticism of lab experiments on grounds of scrutiny (by the experimenter and other participants), also applies to field experiments, insofar as field subjects realize that they are in an experiment. Such experimenter demand effects may arguably, in many cases, be even stronger in field experiments, which are typically run in collaboration with governmental and semi-governmental bodies, and NGOs.

It is indeed the case that when subjects are observed in dictator game experiments in the lab, they give higher amounts (Dana et al., 2007; Haley and Fessler, 2005). In many real-world giving situations, one is also observed by others; for instance, church collections that take the form of passing along a collection plate/basket, or having to declare one’s charitable contributions for tax purposes. However, the effect of being observed disappears if one introduces a minimal element of strategic interaction as, say, in an ultimatum game (Barmettler et al., 2012). A more important determinant of giving in dictator games is whether income is earned or not. Giving in dictator game experiments falls to about 4.3% of an endowment of \$10, when income is earned, relative to about 15% of the endowment in the case of unearned income (Cherry et al., 2002); the figure of 4.3% is closer to the corresponding field benchmark of charitable giving in the US, which stands at about 1% of income (Camerer, 2015).

A commonly heard critique of behavioral models of social preferences is that if experimentally observed social preferences are so important, then, putting it rather starkly, why do we not observe people giving envelopes stuffed with money to others (Bardsley et al., 2010, p. 53)? When

⁹ Andersen et al. (2011) consider extremely high stakes ultimatum game experiments; the stakes vary from the equivalent of 1.6 hours of work to 1600 hours of work. The median offer by the proposer is to give 20% of the share to the respondent, but the rejection rate falls with the increase in the stake. In real life, we rarely make decisions involving 1600 hours of work, yet social preferences were not eliminated in the experiment.

dictators in experiments give out of earned income, then the extent of giving is not too far off from the rate of charitable giving (4.3% versus 1% for the case of US; see above). In the real world, subjects give charity out of after-tax income, which is not the case in the lab. So imagine that in dictator games in the lab with earned income, the dictator was told: “Here is your endowment of \$10, which you have earned. We are taking 30% off as taxes, which we will partly use for redistributive purposes to the recipient in the experiment. How much of the rest will you offer to the recipient?” I will be surprised if the 4.3% giving in lab dictator games does not get closer to 1% figure for charitable giving in the field. Similar observations apply to proposer offers and responder rejections in lab experiments that do not include a tax redistributive component. If this is the case, then giving in experiments may be tapping into the innate human desire to redistribute to others, that is, at least partly, codified institutionally in the social welfare state.

3.2 Subject pools used in lab experiments

It is not unusual in many quarters to dismiss experiments conducted on students, the typical lab subject pool, as having limited or no relevance to testing economic theories. There are several objections to this claim that we now outline.

Economic theory does not specify the subject pool on which its predictions are to be tested. Gilboa et al. (2014, p. F516) write “the economic theorist is typically not required to clearly specify where his model might be applicable and how.” Clearly, you cannot have it both ways by not specifying a subject pool and then objecting to a particular subject pool. This view has been popularized in Vernon Smith’s *blame the theory argument*. Writing in the context of incentives in experiments, Smith (2001) writes in his abstract: “The rhetoric of hypothesis testing implies that game theory is not testable if a negative result is blamed on any auxiliary hypothesis such as ‘rewards are inadequate.’ This is because either the theory is not falsifiable (since a larger payoff can be imagined, one can always conclude that payoffs were inadequate) or it has no predictive content (the appropriate payoff cannot be prespecified).”

One concern with the student subject population is that students might not have the necessary and relevant experience to conform to the predictions of the theory. However, one can allow lab subjects to gain experience in the lab by repeatedly making decisions; indeed, many lab experiments examine such learning effects and the effects of experience. We postpone a fuller discussion of these issues to Section 3.4, where we consider the external validity of lab experiments.

Students possess higher than average education and intelligence, which should be rather favorable to tests of neoclassical economic theory that requires economic agents to possess high levels of cognitive ability. It often comes as a surprise to the critics, but student subjects are much less prosocial relative to non-student subject pools (Falk et al., 2013; Carpenter and Seki, 2011; Anderson et al., 2013).¹⁰ In a review of 13 studies that satisfy stringent tests of comparability, Fréchette (2015) finds that either there was very little difference between the behavior of students and professionals, or students were actually closer to the predictions of neoclassical theory. CEOs are often more trusting as compared to the student population (Fehr and List, 2004). More prosocial students do not self-select themselves as subjects in experiment (Cleave et al., 2012). Students who self-select themselves into experiments are motivated by monetary rewards (Abeler and Nosenzo, 2015), or interest in experimental lab tasks (Slonim et al., 2013). This evidence stands in contrast to their characterization in Levitt and List (2007) (based on two studies conducted in the 1960s) as scientific do-gooders who cooperate with experimenters to seek social approval.

¹⁰ However, student subjects might be more prosocial when it comes to volunteering time (Slonim et al., 2013).

3.3 Stake sizes in experiments

Economic theory does not specify the size of the stakes for which its predictions hold. Experimental economics is typically criticized for its low stakes. The evidence on stake size effects is mixed. However, many experimental results continue to hold, at least qualitatively, even with higher stakes (Slonim and Roth, 1998; Cameron, 1999). The most prominent effect of stakes arises when one moves from hypothetical payoffs to some strictly positive incentives. However, there is much less difference between moderate and high stakes; in particular, the main effect is a reduction in the variance of responses (Camerer and Hogarth, 1999).

There are two issues with high stakes, which are understated in many critiques of experimental economics. First, the vast majority of decisions that we make in real life are low stake decisions. How many times do we buy a car, a house, or a consumer durable such as a TV/laptop? Second, the main evidence for stake effects comes from experiments themselves. Third, as Thaler (2015) notes, the insistence on high stakes arises presumably because we are supposed to pay greater attention to economic decisions involving high stakes. But our success and expertise in making economic decisions is as much a matter of practice and learning. Since high stakes decisions are rare, we get limited opportunities to learn and make optimal decisions; the converse is true of low stakes decisions. Hence, there is no supposition that high stakes decisions should be closer to the predicted outcomes in neoclassical economics. So, he argues, correctly, that economists need to make up their minds whether they wish to insist on high stakes or low stakes as the appropriate test of their theories. Either way, experiments still offer the most natural environment to test the effect of stakes, which is an argument for more, not fewer, experiments.

3.4 The issue of the external validity of lab findings

Camerer (2015) distinguishes between the *policy view* and the *scientific view*. In the policy view, generalizability of lab findings to the field, or *external validity*, is essential. In the scientific view, all properly gathered evidence, including lab and field evidence, serves to enhance our understanding of human behavior. In this view, there is no hierarchical relation between lab and field evidence, and it is a mistake to pose the issue as if one had to make a choice between the two kinds of evidence. Camerer (2015, p. 251) explains cogently: “In this view, since the goal is to understand general principles, whether the ‘lab generalizes to the field’ (sometimes called ‘external validity’ of an experiment) is distracting, difficult to know (since there is no single ‘external’ target setting), and is no more useful than asking whether ‘the field generalizes to the lab.’”

To understand Camerer’s argument more fully, consider the following simple formalization in Falk and Heckman (2009). Suppose that we are interested in some variable Y that can be explained fully by the variables $X_1, X_2, \dots, X_n$ and the “true” functional relation between them is given by $Y = g(X_1, X_2, \dots, X_n)$, which is sometimes known as an *all causes model*. A researcher may be interested in examining the causal effect of X_1 on Y , holding fixed all other variables $\hat{X} = (X_2, X_3, \dots, X_n)$. For instance, in gift exchange experiments, Y is the level of effort of a worker and X_1 is the level of wage paid by the firm. The all causes model will typically include many factors in the vector $\hat{X}$, such as the number of firms and workers, choice sets, payoff functions, incentives, demographic characteristics, regulatory environment, and moral and social characteristics of the parties involved.

When the relevant hypothesis is tested in the lab, the researcher estimates a model of the form $Y = f(X_1, X^L)$, rather than the all causes model $Y = f(X_1, \hat{X})$, where $X^L \neq \hat{X}$; X^L includes

variables such as incentives given in the experiment, the endowments of subjects, the subject pool, context, and the structure of payoffs. One may also conduct field experiments in which one estimates a model of the form $Y = f(X_1, X^{F_1})$, where $X^{F_1} \neq \widehat{X}$, and typically $X^{F_1} \neq X^L$. Field experiments are conducted with a particular subject pool, such as sports card traders in List (2006).

The typical claim by critics of the experimental method is that $f(X_1, X^L)$ does not satisfy external validity, but $f(X_1, X^{F_1})$ does satisfy it. Now suppose that we are interested in examining the gift exchange relation in yet another population of subjects in the field, say, part time employees at General Motors. This gives rise to yet another estimated relation $Y = f(X_1, X^{F_2})$, where X^{F_2} reflects the set of variables and their characteristics in this field experiment. Is there any particular reason why the results based on the model $Y = f(X_1, X^{F_1})$ are more relevant, as compared to $Y = f(X_1, X^L)$, for predicting the causal effects of X_1 in the relation $Y = f(X_1, X^{F_2})$? Camerer (2015, p. 256) offers his assessment (expressed in our notation): “If the litmus test of ‘external validity’ is accurate extrapolation to X^{F_2} , is the lab X^L necessarily less externally valid than the field setting X^{F_1} ? How should this even be judged?” Falk and Heckman (2009, p. 536) go slightly further: “The general quest for running experiments in the field to obtain more realistic data is therefore misguided. In fact, the key issue is what is the best way to isolate the effect of X_1 while holding constant $\widehat{X}$.”

Since the criterion for external validity is unclear, it is best to treat lab and field evidence as complementary. Lab evidence allows for much tighter control of the variables in $\widehat{X}$. Field experiments allow for a larger variation in some aspects of $\widehat{X}$ (e.g., different subject pools with different demographic and social characteristics) while lab experiments allow for larger variation in other aspects of $\widehat{X}$ (e.g., exploration of the parameter space for values that can be hard or rare to find in the field). Lab experiments allow for greater replication because they are less costly and the economic environment in the lab can be more tightly controlled, while any specific field environment could be fairly unique.

We review the evidence for the generalizability of lab evidence to the field in many parts of this book. We end this section with the following bold claim from Camerer (2015, p. 277) made from studies where the lab and field evidence can be well matched: “There is no replicated evidence that experimental economics lab data fail to generalize to central empirical features of field data (when the lab features are deliberately closely matched to field features).” Readers interested in these issues can further consult Camerer (2015, pp. 281–5) for a list of studies that show a good association of lab behavior with field behavior in studies where it is more difficult to match the lab and the field evidence due to differences in the design or subject population.

Not everyone within the experimental economics community is willing to dismiss the importance of external validity. For instance, Kessler and Vesterlund (2015) believe that aiming for external validity is important, that qualitative experimental results have high degree of external validity, and that the concerns about external validity pertain only to quantitative experimental results. However, using a relatively large number of studies, Herbst and Mas (2015) show that peer effects on worker output in lab experiments and field studies from naturally occurring environments are quantitatively very similar. Hence, experiments appear to have external validity, even for quantitative estimates.

In light of the discussion above, the reader may perhaps appreciate better the view taken in this book that all sources of evidence, lab experiments, field experiments, and field data are equally valid and complementary in nature.

3.5 The role of incentives in economics

The norm in experimental economics is that all decisions made by subjects, or any elicitation of their underlying preferences in experiments, should be incentive compatible. It would currently be near impossible to get an experimental paper published in an economics journal if it did not respect incentive compatibility, preferably by using monetary rewards. In contrast, psychology does not require similar adherence to incentives. Camerer and Hogarth (1999) perform a meta-study of 74 studies, over the period 1990–98, in the leading economics journals: *American Economic Review*, *Econometrica*, *Journal of Political Economy*, and *Quarterly Journal of Economics*. Not a single experimental study in this sample was published without the use of incentives. In a meta-study over the period 1988–1997, Hertwig and Ortmann (2001) found that only 26% of the articles published in a leading journal in psychology, the *Journal of Behavioral Decision Making*, used monetary task related incentives.

Economists have traditionally viewed effort as aversive and requiring scarce cognitive resources. Hence, they believe that people will exert effort only if the marginal disutility of effort is exceeded by the marginal utility of monetary incentives (*extrinsic motivation*). By contrast, psychologists have stressed the *intrinsic motivation* of subjects in experiments that does not require task related monetary incentives. Psychologists do pay subjects a show-up fee and/or course credit because it may be unethical to pay less than a minimum wage for student's time and also because the fixed fee may elicit a reciprocal response by priming intrinsic motivation. Another difference between experiments in economics and psychology is that in the former, experimental practice is much more standardized and regulated (Hertwig and Ortmann, 2001). For instance, deception is taboo in economics experiments, but its practice is variable in psychology experiments.

Smith (1976) made an early recommendation for the use of incentives, based on his experience in double auction experiments. Employing monetary incentives in experiments is also a common recommendation by some of the pioneers in the field (Davis and Holt, 1993; Roth, 1995; Smith, 1991; Smith and Walker, 1993). The effectiveness of incentives, particularly in settings outside double auction experiments, is an empirical question. Several meta studies of the effect of incentives are now available, which show that the effect of monetary incentives on effort and performance in experiments is quite subtle. Neither of the following two extreme positions is supported by the evidence: Incentives make no difference and incentives remove all behavioral anomalies.

However, there is a caveat to these meta studies. Since the underlying preferences are unobserved, in experiments, task performance and incentive compatibility are measured by compliance with the predictions of the underlying theory. However, the underlying theory might itself be refuted by the evidence. One might then fall back on measuring the effect of incentives on the effort undertaken by the subjects (an input), rather than the performance in the task (the output); for a brief survey of these studies, see Bardsley et al. (2010, Ch. 6). However, in many tasks, effort might not improve performance. For instance, consider subjects in experiments engaged in non-trivial strategic interaction. In most games, it would be hard for the players to hit upon the Nash equilibrium purely by deduction, even if incentives were high.

The results reported in the careful meta-study by Camerer and Hogarth (1999) are particularly instructive. In many economic environments, such as choosing among risky gambles, market trading, and bargaining, the weight of the evidence (and the modal result) shows that increasing incentives does not change the average behavior. However, incentives often reduce the variance of results. Incentives may even harm the performance of subjects in experiments, for instance, subjects may choke under pressure (Ariely et al., 2009); their intrinsic motivation may be crowded out, and they may engage in too much payoff-reducing experimentation to avoid

negative incentives such as punishments (Hogarth et al., 1991);¹¹ they may place too great a reliance on personal judgment, when relying on public information would have improved pay-offs (Arkes et al., 1986); they may experience motivational crowding-out, e.g., in the presence of incentives, individuals do not feel responsible for their own behavior (Gneezy and Rustichini, 2000b); individuals might be insulted by the low level of incentives (Gneezy and Rustichini, 2000a). However, for another set of problems, mostly in judgment tasks, financial incentives affect average performance. Typically, the nature of these tasks is such that increased effort improves performance. Examples include memory retrieval tasks, tasks where recalling the play in previous rounds improves predictions, or mundane tasks such as piece rate clerical work where one might be easily bored, e.g., counting the number of occurrences of a particular alphabet on a page of English text.

The meta-study by Bonner et al. (2000) supports the results of the Camerer–Hogarth study. They find that incentives have little effect in problem solving tasks, highest positive effects in tasks where effort improves performance (e.g., clerical tasks, pain endurance, or detecting typos on a page), and weak positive effects in judgment tasks. Hertwig and Ortmann (2001) find moderately positive effect of incentives. However, their sample size is small. Only 48 out of 186 studies (26%) that they considered used financial incentives. Of these, only 10 explored the effect of monetary payments. One of their suggestions is to use different treatments that do use and do not use monetary incentives to build a better picture of the effect of incentives.

Clearly, a reduction in the variance of responses in the presence of incentives, in some experiments, is desirable on the grounds that it improves the statistical power of tests. However, economics advocates a cost–benefit approach, and research funding for experiments is a scarce resource. Hence, one has to trade off the reduced variance against performing more experiments with bigger subject pools in the absence of monetary incentives. Greater variance of responses produces more outliers, but as Camerer and Hogarth (1999, p. 31) point out, there are alternative solutions: “Of course, other methods might work the same magic more cheaply. Trimmed means and robust statistical methods also reduce the influence of outliers.”

So what explains the absolute adherence of experimental economists to incentives as a matter of norm? Bardsley et al. (2010, pp. 249–50) conjecture a cynical explanation: “A reason often given to psychologist’s opposition to them is that incentives present an obstacle to less-experienced researchers who find it harder to secure research funding. It may be that economists are less concerned by the presence of such obstacles, or even that some actively promote the use of incentives as a barrier to entry.” Many psychologists remain unconvinced of the arguments that experimenters in economics propose in favor of using incentives. For instance, Read (2005, p. 265) concludes his paper as follows: “. . . there is no basis for requiring the use of real incentives to do experimental economics.”

The decision to employ incentives is a part of the experimental design. But as a practical matter, experimenters interested in publishing their work in economics journals have little choice in this matter. A second issue in the design of experiments is the level of incentives and the context. Gneezy and Rustichini (2000a) use four different levels of incentives, zero, low, medium, and high, to check for the level effect and conduct two different experiments to check for context effects. In one experiment, the IQ experiment, they find that incentives improve performance when one compares the levels low and high; but there was no difference between medium and high incentives. In the second experiment, the donations experiment, donations were highest

¹¹ For a formal model that shows the trade-off between intrinsic and extrinsic motivation, see Bénabou and Tirole (2003).

when no incentives were given; in this case, intrinsic motivation is superior to extrinsic motivation. Using the same dataset, Rydval and Ortmann (2004) find that ability levels of individuals explained greater variation in performance relative to incentives. The assessment offered by Bardsley et al. (2010, p. 253) on the level of incentives is this: “The message seems to be that in terms of the impact on cognitive effort allocation the presence of task related incentives matters more than their level.”

How can an experimental economist who has decided to run experiments using incentives ensure that the responses by the subjects are incentive compatible? Typically, incentive compatibility in experiments has been defined with respect to expected utility theory; however, the behavior of the majority of people is not consistent with expected utility theory. Hence, judging incentive compatibility is a vexed task. For instance, the leading method of ensuring incentive compatible choices, the *random lottery incentive scheme* (RLI), that we outline in Appendix A, depends crucially on the *independence axiom* of expected utility theory. However, if the decision maker does not follow expected utility but follows, say, rank dependent utility or prospect theory, then the independence axiom does not hold. Thus, choices are not guaranteed to be incentive compatible.

The evidence against expected utility is now overwhelming (see Part I) and the independence axiom has been rejected in many empirical tests (e.g., *Allais paradox*). Bardsley et al. (2010) argue, however, that despite the rejection of the independence axiom, empirical evidence shows that the random lottery mechanism does not bias results. They write (p. 270): “it may simply be a happy coincidence that the RLI works, by engaging a particular mental heuristic that promotes unbiased task responses.” Some readers will wish to seek further clarification of the underlying mechanism that makes RLI an empirically attractive option.

Another popular incentive compatible mechanism in experiments that rely on an expected utility formulation is the *Becker–DeGroot–Marshak mechanism* (BDM) (Becker et al., 1964). This is typically employed in eliciting the willingness to pay and the willingness to accept for an object. Subjects in their roles as sellers are asked to state their valuation, v , for an object. A random price, p , is then drawn. If $p \leq v$, then a seller is not allowed to sell but if $p > v$, then the seller must sell. Subjects who follow expected utility should report their true valuations. To see this, suppose that the seller’s valuation for an object is believed to be distributed randomly over the interval $[\underline{v}, \bar{v}]$, the true valuation is v , and the price is drawn randomly from the interval $[\underline{v}, \bar{v}]$, using the distribution F . Suppose that the seller chooses to declare a low valuation, $v_L < v$. In this case, the seller chooses to forgo a profitable opportunity to sell if the price turns out to be in the interval $(v_L, v]$ and the expected forgone profits are $\int_{v_L}^v (v - p) dF(p) > 0$. Thus, the seller never declares a low valuation. Now suppose that the seller declares a high valuation, $v_H > v$. In this case, the seller risks the potential expected loss $\int_v^{v_H} (v - p) dF(p) < 0$, so he never overstates his valuation. In short, the BDM is incentive compatible under expected utility. A direct implementation of the BDM has been found not to be incentive compatible (Bohm et al., 1997). However, suppose decision makers are presented with ascending or descending prices and must choose to trade or not at these prices. In the end, one of the trades is picked at random and paid-off. There is some evidence that this version of the BDM may be incentive compatible (Braga and Starmer, 2005).

Similar calculations as employed in the BDM show that the *second price auction* (bidders bid simultaneously, and the winner pays the second highest bid) is also incentive compatible under expected utility. Bardsley et al. (2010, Chapter 6) review evidence which shows that the responses in BDM and the second price auction are sensitive to experience and, in later rounds, valuations are likely to converge to the true values. But the evidence on the incentive compatibility of variants of the *Vickery auctions* (which is a more general form of second price auctions) is mixed. They

write (p. 274): “the efficacy of particular elicitation mechanisms can vary across different types of tasks. Indeed, whether an elicitation mechanism works may depend on the kind of research question that motivates its use.”

To summarize, the discussion on incentive compatible elicitation mechanisms in experiments raises uncomfortable questions for the practice of experimental economics. First, predicted incentive compatibility relies on theories such as expected utility theory that have inadequate empirical support. Second, there are no universally agreed on mechanisms that will guarantee incentive compatibility in experiments and it very much comes down to a matter of judgment of the experimenter. Third, despite the firm insistence of experimental economists to use task related incentives, the theoretical response to these problems (e.g., the development of incentive compatibility under behavioral theories) appears to be inadequate.

3.6 Is survey data of any use?

Section 3.5 shows that there is often merit in the use of incentives, e.g., usefulness in double auction experiments, judgment tasks, and effort related responses as in clerical tasks, and reduction in variation of responses. However, it is fair to say that the empirical case for the use of incentives in experiments is not as strong as is typically made out within the experimental economics community. Hence, one needs to consider other sources of data too. Most economists mistrust *survey data* and data based on *hypothetical questions* on the grounds that subjects are not incentivized to reveal their true preferences. This is due to the traditional belief in economics that people mainly have *extrinsic* rather than *intrinsic* motivation.

In the history of behavioral economics, the reliance on hypothetical questions and survey data often gave rise to deep insights that allowed for major advances. Two prominent examples include the work of Daniel Kahneman and Amos Tversky on prospect theory (Kahneman and Tversky, 1979) based on hypothetical questions in lab experiments, and the work of Daniel Kahneman, Richard Thaler, and Jack Knetsch on fairness motivations in humans (Kahneman et al., 1986) based on hypothetical questions posed in telephone surveys. Kahneman and Tversky (1979) is the second most cited paper in all of economics, the foundation for the Nobel prize to Kahneman, and the source of prospect theory, which is currently the most satisfactory decision theory under risk, uncertainty, and ambiguity.¹² Yet the paper is based on hypothetical, non-incentivized, lab experiments. Any guesses if it would have been published in an economics journal today?

Another merit of studies based on hypothetical, non-incentivized, questions is that they can complement lab experiments in certain areas. We give an example from fairness concerns in Section 6.1. Here, we consider two examples from the work of Daniel Kahneman and Amos Tversky.

Example I.1 *Consider the identification of loss aversion (losses bite more than equivalent gains) in Kahneman and Tversky (1979). In lab experiments, it is not always easy to induce loss aversion over “large losses” because it is considered unethical in experimental economics to leave subjects out-of-pocket. The main alternative in lab experiments is to give subjects money upfront to ensure they are never out-of-pocket. However, this may contaminate their responses by the “house money effect” (gamblers who win money in a casino are more likely to gamble with it; as they say, easy come easy go).*

Example I.2 *Consider the following hypothetical, non-incentivized, lab questions from experiments conducted in a seminal paper by Kahneman and Tversky (1984). The percentage of subjects choosing each response is given in brackets.*

¹² The source for the “second most cited paper” claim is Table 2 in Kim et al. (2006).

Imagine that the US is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows:

Positive Framing: If program A is adopted, 200 people will be saved. If program B is adopted, there is a one-third probability that 600 people will be saved and a two-thirds probability that no people will be saved. Which of the two programs would you favor? (72% chose A, 28% chose B).

Negative Framing: If program C is adopted, 400 people will die. If program D is adopted, there is a one-third probability that nobody will die and a two-thirds probability that 600 people will die. Which of the two programs would you favor? (22% chose C, 78% chose D).

Relative to the status-quo, positive framing presents data in terms of “lives saved”, while negative framing presents the same data in terms of “lives lost”. Options A and C are identical, and options B and D are also identical. Under frame-invariance, the typical assumption in neoclassical economics, if A is chosen over B then C must be chosen over D (and vice versa). However, in the domain of gains (lives saved relative to the status quo), the majority (72%) chose the safe option (A) over the risky option (B). In the domain of losses (lives lost relative to the status quo) the majority (78%) chose the risky option (D) over the safe option (C). This establishes several results. First, behavior is not frame-invariant. Second, the results are not consistent with expected utility. Third, subjects are risk averse in gains and risk seeking in losses; this is one of the important insights of prospect theory.

These results have survived several challenges (e.g., choices over incentivized lotteries) and even professional physicians and World Bank staff behave in the same manner as lab subjects; see Part I for more details. Clearly this experiment cannot be run in the lab because it would be unethical to take a life.

Thaler (2015, p. 47) writes on the aversion to survey data: “This disdain [for survey data] is simply unscientific. Polling data, which just comes from asking people whether they are planning to vote and for whom, when carefully used by skilled statisticians . . . yield remarkably accurate predictions of elections. The most amusing aspect of this anti-survey attitude is that many important macroeconomic variables are produced by surveys! . . . “jobs” data . . . come from surveys conducted by the Census Bureau. The unemployment rate, . . . is also determined from a survey that asks people whether they are looking for work. Yet using published survey data is not considered a faux pas in macroeconomics. Apparently economists don’t mind survey data as long as someone other than the researcher collected it.”

Happiness economics, which we consider in Part 6 is based almost entirely on survey data. Policymakers appear to take serious interest in this area and survey measures of well-being and happiness are highly correlated, which is consistent with the standard expected of incentivized responses. Furthermore, firms often use information gleaned from market surveys to introduce new products about which consumers have no prior experience, and alter the characteristics of existing products.

Consider an example where incentives make the problem even worse. Using hypothetical, non-incentivized questions, Lichtenstein and Slovic (1973) asked people to choose between a less risky P -bet (high probability of winning a low prize) and a more risky $\$$ -bet (a low probability of winning a high prize). When asked to choose directly between the two bets, most people chose the low risk option, the P -bet. Yet, most people assigned a higher certainty equivalent to the $\$$ -bet as compared to the P -bet. Thus, we have a case of *preference reversals* that cannot be explained by any model of consistent preferences. Grether and Plott (1979) suspected that the

problem was caused by hypothetical choices. However, when they re-ran the experiments using incentivized subjects, the incidence of preference reversals worsened (see Part I of the book for details).

In sum, there are no grounds to outrightly reject survey data. Survey data has pitfalls, for instance, poll data sometimes gives misleading results, however the benefits of survey data outstrip the potential pitfalls. But other data sources also have limitations; hence, it is important to use all sources of data, surveys, experiments, and field evidence, to build a more complete picture of human behavior.

4 Approach and organization of the book

This book is an attempt to take stock of behavioral economics and is aimed to serve several purposes: A course text for advanced students in economics and other social sciences, a research handbook for behavioral economists, and an invitation to economists and other social scientists of all persuasions to explore this exciting new field. In its teaching role, the book would ideally be taught in a year long course in behavioral economics, supplemented by readings that reflect the interests of the instructor. A good example is a typical two semester North American style course that has 13 weeks of teaching in each semester; students meet each week for a two hour lecture and a one hour problem solving class. This book tries to standardize the material in behavioral economics so that students can see behavioral economics as a coherent and widely applicable body of theory, much as they might see any other established area in economics. It is also intended to ‘nudge’ many idiosyncratic course outlines, that pass for behavioral economics at the moment, to adopt a more balanced approach.

Given the scope of the book, one can also construct a large number of coherent and interesting single semester courses not just in behavioral economics but also in behavioral decision theory, behavioral macroeconomics, behavioral industrial organization, behavioral contract theory, topics in behavioral economics, and the like. I believe that the book should also be essential background reading for any advanced course in microeconomics in order to address Herbert Gintis’s (2009, p. xvi) concern raised above, till the time that microeconomics books respond to his suggestion.

Despite the size of the book, it is, I believe, the “minimum” amount of material that any academic who declares behavioral economics or experimental economics as their research interest, must have deep, rather than passing, familiarity with.

As the reader may have guessed from Sections 2 and 3, my main criterion for including models in this book is their consistency with the evidence, unless there is significant merit in using a model for pedagogical or other reasons. Most of the models that I use can typically explain at least as much as neoclassical models, but can also better explain data from other domains of human behavior. It is entirely possible, of course, to append auxiliary assumptions to standard neoclassical models to explain almost any set of stylized facts. However, I try to stay away from models in which these auxiliary assumptions are ad-hoc fixes.

I encourage the reader to keep an open mind about behavioral economics models and judge models by the relevant empirical evidence. Models in behavioral economics are as theoretically rigorous as models in neoclassical economics, if rigor is an important criteria for the reader. Gilboa (2009, Section 7.1) may well be right that all models are ultimately wrong. Even if this statement were true, we must consistently strive to improve our models in the light of rejections, in order to make better predictions and improve our understanding of human behavior. The

sequence of these improved models may or may not ever converge to the “true model.” Yet, each time we reject a model and put in its place a new one that makes better predictions, we make progress.

The level of honesty that we need as a profession is captured in the words of the Nobel Prize winning physicist Richard Feynman (1965, p. 158) that economists would do well to embrace: “But experimenters search most diligently, and with the greatest effort, in exactly those places where it seems most likely that we can prove our theories wrong. In other words, we are trying to prove ourselves wrong as quickly as possible, because only in that way can we find progress.” The “all models are wrong critique” is unhelpful at best and, if used inappropriately, it is likely to hinder progress in the subject.

For many of the topics, the book is organized in the following general style. A neoclassical theory is outlined, followed by a review of the evidence for it. If the weight of the evidence is inconsistent with the theory, then the behavioral alternatives and the evidence for or against them is described. There is a range of applications in the book for most of the behavioral theories that the reader/instructor can choose from. In order to enable readers who might not be familiar with the full range of topics in behavioral economics, I err on the side of longer introductions to most of the chapters. As a first pass, readers may wish to skim through the introductions to the various parts/chapters, in order to get a birds-eye view of the material. This should enable readers to draw up a priority list of topics that they would wish to read or teach in their courses. I deliberately shy away from offering such a list of topics to the reader, and like a good restaurant, prefer to fall back on offering a wide menu of tempting choices, hoping that you, the reader, will keep coming back, and spread the word.

Some readers may feel that it is presumptuous to include the word “foundations” in the title of this book. However, in justification, I believe that there is now a sufficiently rich body of lab and field evidence that is well described by models of behavioral economic theory. The richness of human behavior that we can account for with these models is unprecedented in modern economics. A few of the areas in behavioral economics, notably *neuroeconomics*, are very young and I have some hesitation in bringing this material under the rubric of a book that claims to describe the foundations of a subject. Yet, omitting this material is not an attractive option, and I have erred on the side of inclusion.

I also try to flag up material that is speculative, yet promising, or where the evidence base is not sufficiently established at this point in time. Behavioral economics continues to be work in progress and despite the huge increase in understanding attained over the last several decades, a lot needs to be done. Writing this introduction upon completion of the book, and reflecting on the material, I must confess that I have never personally felt so enriched, yet so ignorant, in my life. I hope that the serious, and just the plain curious, among you, come away with a similar sentiment after reading this book.

The book is organized into nine parts and many parts have multiple chapters; there are problems at the end of each part, except for Part 9 on neuroeconomics. The order of chapters reflects, at least partially, the historical development of the subject, and also bears some similarity with the organization of a typical course in microeconomics. Each part begins with an introduction to the material in all the constituent chapters in the part, and each chapter has a separate, but more detailed, introduction.

The table of contents give a bird’s eye view of what is inside the book and I will not bore you with a verbatim description; anyone bold enough to read so far, probably has enough on their plate to think about. I should make a few broad comments though, about my organization of the material. I focus mainly on topics that readers may find to be in an unexpected location within the book.

The book opens with Part 1 on *behavioral decision theory* because playing a game against nature under risk, uncertainty, and ambiguity, in a static setting, is one of the simplest economic problems; this abstracts from temporal and strategic concerns. However, there is an important overlap between risk preferences and time preferences, which is split between Part 1 and Part 3 on time preferences. Many readers might have wished to see a more thorough treatment of ambiguity. However, most of the modern developments in ambiguity are within the confines of the neoclassical models, although an important behavioral literature on source-dependent preferences, that we describe in Part 1, is now beginning to develop that highlights the role of prospect theory.

Microeconomics texts typically begin with a discussion of the properties of human preferences, which are regarded as the primitives of the model. Most evidence indicates that there is a mixture of individuals with purely *self-regarding preferences*, as in neoclassical economics, and those with *other-regarding preferences*. Part 2 considers the evidence on *human sociality* and behavioral models that take account of this evidence. There is also a discussion of the implications of social preferences for competitive general equilibrium in this part, which typically comes much later in a microeconomics course. Two important topics that lie at the heart of human motivation in economics, *social identity* and *human virtues*, appear in a stand-alone chapter in Part 2. The inclusion of these major topics should be less surprising than the fact that most microeconomics texts have somehow contrived to exclude them.

An explicit time dimension is not fundamental to the material in Parts 1 and 2. In Part 3, we consider the evidence on time preferences and *behavioral models of time discounting*. The treatment of time preferences is excessively narrow in neoclassical microeconomics texts. This is because the entire psychology of time preferences is captured by a single parameter, the discount rate, in the exponential discounted utility model—the main model of time preferences in neoclassical economics. This is not just unsatisfactory in light of the richness in observed time preferences, but exponential discounting is strongly rejected by the evidence, perhaps even more so than expected utility theory.

There will be much that is new and unfamiliar to the reader in Part 3 that will require investment in learning new machinery such as *subadditivity*, *attribute-based models* of time preferences, and *models of reference time preferences*. I advise the serious reader to persevere with Part 3. It is worth telling an anecdote that the reader may wish to keep in mind when an unfamiliar topic (and there will be many) is encountered in Part 3. When George Loewenstein, one of the leaders in time discounting, first saw Part 3, he particularly commended me for including *subadditivity* (i.e., discounting depends on how one partitions a given interval of time). However, when another group of non-specialists read Part 3 as part of a group reading seminar in a leading university, they were put off by the “unusual and new” concept of subadditivity. The reason subadditivity is there in the book is because it is currently supported by the evidence; its current popularity/aesthetic appeal (or lack of it) are criteria that, for this book, are not relevant.

Part 4, the longest part of the book, is on *behavioral game theory*. Classical game theory revolutionized economics by forcing it to specify explicitly the economic environment, set of players, sequence of moves, and the mapping of histories to payoffs. Anyone comparing industrial organization theory before the advent of game theory with modern industrial organization theory will immediately notice the much greater clarity achieved by modern game theoretic models (think, e.g., of models of entry deterrence). Gintis (2009) advocates keeping game theory as the common toolkit for all social sciences despite pointing out serious shortcomings in classical game theory, for instance, the assumption of common priors and the justification for a Nash equilibrium. The approach in this book concurs with this sentiment. Despite widespread belief to the contrary, the evidence has not been too kind to classical models of game theory. This is particularly the

case for observed behavior in the early rounds of a game, and for equilibrium refinements that require high cognitive requirements. In particular, it stretches the imagination to believe that players could arrive at a Nash equilibrium purely by a deductive process. For instance, the cooperation rate in many static prisoner's dilemma game experiments is about 60%; cooperation here is a dominated strategy, which in classical game theory, ought never to have been played.¹³

Part 4 considers at length the evidence for classical models of game theory and the leading behavioral alternatives. Many of the main behavioral alternatives, such as *level-k models*, the *cognitive hierarchy model*, and the *quantal response equilibrium*, relax the “equilibrium in beliefs” assumption, or the assumption that players play a best response. There are several other behavioral models of strategic interaction that are in the fray. While the empirical evidence is not always consistent with these models, the most heartening aspect of the discourse in behavioral game theory is that we are actually using empirical evidence to choose among the models. You cannot convince the leaders in this field, say, Vincent Crawford or Colin Camerer, about your proposal for a new solution concept in game theory by simply appealing to the aesthetic appeal of your proposed model, or its ability to tell a useful story, or a fable, if it is not supported by the evidence. I have decided to include the topic of *psychological game theory* in Part 4 but some instructors may have wished to place this topic in Part 6 on emotions, because it also deals with *anger* and *guilt*, and others in Part 2, because it also deals with intentions-based reciprocity. My relatively detailed treatment of this material reflects my preference to see this material closer to the centre stage of modern behavioral economics.

Part 5 outlines the evidence on human learning and traditional as well as *behavioral models of learning*. I also give chapter length treatments of *evolutionary game theory* and *stochastic social dynamics*. An appendix briefly introduces the reader to the necessary technical machinery on deterministic and stochastic dynamical systems. Some instructors would have wished to see the chapter on evolutionary game theory in Part 4. However, one way of looking at evolutionary game theory is that it deals with learning at the level of the gene. One aim of this chapter is to provide evolutionary foundations for human sociality that we covered in Part 2, and also introduce some newer and interesting topics, such as *gene-culture coevolution*. Part 5 is particularly instructive in evaluating the common claim that a Nash equilibrium arises if sufficient learning opportunities are provided to the subjects. A separate chapter-length treatment on stochastic social dynamics in Part 5 might be unexpected, even for many behavioral economists who do not typically include this topic in their courses. However, I believe that this material needs to be taken more seriously, and all students of behavioral economics should have at least a basic familiarity with it.

Neoclassical economics typically focuses on cold and emotionless deliberation. Part 6 considers the role of emotions in explaining economic phenomena. Instructors interested in teaching a course that focuses on emotions will need to combine the material here with that from other parts of the book, such as psychological game theory (Part 4), models of regret and disappointment aversion (Part 1), and issues of self-control and present-biased preferences (Part 3). Part 6 also considers the Gul–Pesendorfer model of temptation that some readers might have preferred to situate in Part 3, where issues of present-biased preferences are discussed. I have also placed the important policy relevant topic of *happiness economics* in Part 6, insofar as happiness may be considered as a type of emotion. The evidence base in happiness economics is typically constructed using survey data that not all economists agree with (see Section 3.6 above), which has hampered its acceptability within academic economists. I urge economists to take this material

¹³ Social preferences alone cannot explain behavior of the players in the prisoner's dilemma game; see al-Nowaihi and Dhami (2015).

more seriously, particularly given the correlations between various measures of well-being and happiness. Some readers may find *models of dual selves* to be somewhat out of place in this part; I partly share this concern but it was not immediately clear where better to situate this important topic.

Part 7 considers *models of bounded rationality*, and it is split into three chapters. The chapter on *judgment heuristics*, a topic that owes its origin and importance to Tversky and Kahneman (1974), brings into sharp focus the debate about the relative suitability of the *optimization approach*, and the *heuristic-based approach* to economic models.¹⁴ The chapter on *mental accounting* considers work that Richard Thaler almost single-handedly pioneered. It deserves to see much more development and interest among theorists working in behavioral economics. The reader may not have expected a chapter on *behavioral finance* in Part 7 but there is no better place to illustrate bounded rationality and the inefficiency of markets in the very market that is typically held up as a model of efficiency in neoclassical economics.

Part 8 considers issues in *behavioral welfare economics*, which is an area that is likely to experience more development in the future. If people have behavioral biases relative to the neoclassical model, then should we respect those biases or not? The debate is not just about *libertarianism* and *paternalism*, but also about: what exactly do these terms mean in the presence of behavioral biases?

Part 9 gives a very brief tour of *neuroeconomics* that highlights the neuroeconomic foundations of just a few selected aspects of human behavior from the first four parts of the book. In particular, I make no attempt at completeness, or balance, in the treatment of Part 9.

There are some notable omissions from the book that are partly dictated by technological considerations such as the physical size of the book, the current importance of these topics in behavioral economics, and last but not least, by my own lack of competence in these areas. *Complexity theory* and its implications for public policy do not figure in this book; readers can pursue Mitchell (2009) for an introduction, followed by Colander and Kupers (2014) for policy issues. I have also omitted a discussion of *agent-based computational models*; readers can begin with Tesfatsion and Judd (2006) and Farmer and Foley (2009). I also give insufficient attention to the *epistemic foundations of equilibrium concepts* in classical and behavioral game theory; see Dekel and Siniscalchi (2015) for a recent review and Gintis (2009) for a critique of the epistemic foundations of classical game theory. My discussion of neuroeconomics is highly incomplete and selective; for an authoritative survey, see Glimcher et al. (2009) and Glimcher and Fehr (2014).

5 Five theoretical approaches in behavioral economics

I would now like to pick up the story of behavioral economics where I left off at the end of Section 1. I shall consider five case studies of theoretical models in behavioral economics, chosen from five different parts of this book.¹⁵ In each case, I also briefly comment on some/all of the following: the historical context of these models, reactions to these models in the profession, evidence on the models, applications, and lessons learnt. The reader should not construe that I

¹⁴ For a discussion of alternative views on the optimization versus heuristics debate, the reader may consult the June 2013 issue of the *Journal of Economic Literature*, Volume 51, No. 2.

¹⁵ An idea whose development I wished to trace out in this section, since Gary Bolton pointed it out to me a few years ago, was the development of models of relative concerns for others that arose surprisingly from experiments on bargaining games conducted by Alvin Roth with others (see Part 4 of the book). However, Roth (2015) has recently described this story very lucidly, so I have omitted it (see Part 4 of the book for details).

consider the behavioral models not mentioned in this section as hierarchically inferior. An identical observation holds for the material in Section 6 below, where I consider five examples of behavioral evidence.

5.1 A case study of prospect theory

Two papers in the 1970s by Daniel Kahneman and Amos Tversky may be identified with the beginning of modern behavioral economics. In the first, Tversky and Kahneman (1974) proposed a radical, non-optimization approach to decision making in economics in which people use heuristics or simple rules of thumb; we discuss this work in Part 8 and Section 5.5 below. In the second, Kahneman and Tversky (1979) outlined a new theory of decision making under risk, uncertainty, and ambiguity that distilled in a simple, conservative, extension of expected utility theory, the main experimental insights from their own work and that of others. They called it *prospect theory* (PT), which is not a particularly illuminating name for the theory.¹⁶

Consider a lottery that describes outcome, probability pairs in n states of the world,

$$L = (x_1, p_1; x_2, p_2; \dots; x_n, p_n), \quad (1)$$

where $x_1 < x_2 < \dots < x_n$ are monetary prizes (positive or negative) and $p_i \geq 0$ is the objective probability with which the prize x_i occurs, such that $\sum_{i=1}^n p_i = 1$. In the dominant decision theory in neoclassical economics, *expected utility theory* (EU), lottery L is evaluated by the function $EU(L) = \sum_{i=1}^n p_i u(x_i)$, where u is a utility function for monetary prizes. When well-defined subjective probabilities replace objective probabilities (uncertainty) then EU is referred to as *subjective expected utility*. Two features stand out in this formulation. First, utility derives from the final outcomes in each of the n states. Second, the objective is linear in probabilities. EU is underpinned by a compelling and plausible set of preference axioms (von Neumann and Morgenstern, 1944). However, a very large body of evidence rejects the hypothesis that all humans follow expected utility theory (Kahneman and Tversky, 2000; Starmer, 2000; Wakker, 2010). Axiomatic plausibility on its own cannot be used as a criterion for the continued acceptability of a theory if the evidence rejects it. Indeed, the independence axiom, which appears eminently sensible is often rejected by the evidence; the Allais paradox is one of the earliest violations.

PT is a *descriptive* theory of choice that strives to explain actual human behavior not just for risk, but also for uncertainty and ambiguity. Like many behavioral theories, it also has rigorous axiomatic foundations. PT not only accounted for the known violations of EU, it helped to successfully predict and explain a range of new phenomena. In this sense, it satisfies the criteria of theoretical and empirical progress in the Lakatosian framework. A key idea in PT is *reference dependence*. This implies that humans derive utility from changes in outcomes relative to some reference outcome; if the outcome is equal to or greater than the reference outcome it is in the *domain of gains*, otherwise it is in the *domain of losses*.

One could renormalize all outcomes relative to a reference outcome under EU but this is hardly enlightening, because EU has little to offer towards an understanding of human behavior in the domain of gains and losses. In contrast, PT gives a rich account of the difference in human behavior in the domain of losses. A key idea that drives many results in behavioral economics is *loss aversion*, i.e., losses bite more than equivalent gains. Kahneman and Tversky (1979) report a

¹⁶ If you are mystified by the name, prospect theory, as I was once, then you can seek a resolution to this mystery by reading p. 25 in Thaler (2015).




Figure 1 The power form of the utility function under prospect theory.

median figure of loss aversion of 2.25. So, for instance, assuming linear utility, a monetary gain of 100 feels like a utility gain of 100, while a monetary loss of 100 feels like a utility loss of 225 under PT; under EU, a loss of 100 would just feel like a utility loss of 100. Loss aversion is empirically very robust, and it may help some of us to understand our own past behavior.¹⁷

Furthermore, evidence shows that the utility function, v , under PT is concave in the domain of gains and convex in losses. This is shown in Figure 1 for the case of the following power form of utility

$$v(y) = \begin{cases} y^{\gamma^+} & \text{if } y \geq 0 \\ -\lambda (-y)^{\gamma^-} & \text{if } y < 0 \end{cases}, \gamma^+ > 0, \gamma^- > 0, \lambda > 1, \quad (2)$$

plotted for the case of the empirically observed values of $\gamma^+ = \gamma^- = 0.88$ and $\lambda = 2.25$ in Kahneman and Tversky (1979). Notice the kink at the reference point of zero and a sudden increase in slope as we enter into the domain of losses. This captures loss aversion, as reflected in the parameter λ in (2). Given the different shapes of the utility function in gains and losses (concave in gains, so risk averse, and convex in losses, so risk loving),¹⁸ the reader should now readily make sense of the results in Example I.2.

A third component of prospect theory is *non-linear probability weighting*; recall that under EU, probabilities are weighted linearly. In particular, at least for probabilities bounded away from the endpoints of the interval $[0, 1]$, evidence suggests that low probabilities are subjectively over-weighted, while high probabilities are subjectively underweighted. This is captured by a continuous, strictly increasing, *probability weighting function*, $w(p) : [0, 1] \rightarrow [0, 1]$ such that $w(0) = 0$ and $w(1) = 1$. Figure 2 shows a plot of the Prelec probability weighting function due to Prelec (1998); this is given by $w(p) = e^{-\beta(-\ln p)^\alpha}$, where $\alpha > 0$, $\beta > 0$. Arguably, Figures 1 and 2 are the two most important diagrams in Part I of the book.

¹⁷ I undoubtedly received many gifts as a child that I only have a dim memory of, but I can never forget the 5 rupee note that I lost at age 5, or the toy car given to me by an uncle from the US that was stolen on the very first day in school. Losing the toy car was certainly a far more painful and long-lasting experience as compared to the joy that I felt at receiving it. Loss aversion may be one explanation behind these feelings.

¹⁸ We have not, so far, said anything about non-linear probability weights, which, with the shape of the utility function, jointly determine attitudes to risk under PT (see below).




Figure 2 A plot of the Prelec function for $\beta = 1$ and $\alpha = 0.5$.

Suppose that an individual faces the following lottery in *incremental form*

$$\tilde{L} = (y_{-m}, p_{-m}; \dots; y_0, p_0; \dots; y_n, p_n), \quad (3)$$

where $y_{-m} \leq y_{-m+1} \leq \dots \leq y_{-1} \leq y_0 = 0 \leq y_1 \leq y_2 \leq \dots \leq y_n$ are monetary outcomes, each expressed as an increment relative to a reference point r ; so if the monetary outcome is x_i , then $y_i = x_i - r$. Here, we have m outcomes in the domain of losses and n outcomes in the domain of gains. Kahneman and Tversky (1979) proposed that a decision maker who follows PT, evaluates lottery $\tilde{L}$ as: $PT(\tilde{L}) = \sum_{i=-m}^n w(p_i) v(y_i)$.¹⁹ However, it was soon realized that this could lead to stochastically dominated options being chosen.

Quiggin (1982) suggested a way out of this problem with his proposal of *rank dependent utility* (RDU) that alters EU in a minimal, but important way. A decision maker who follows rank dependent utility evaluates the lottery L in (1) as: $RDU(L) = \sum_{i=1}^n \pi_i u(x_i)$; π_i , known as the *decision weight* corresponding to outcome x_i , is simply a cumulative transformation of the probability weighting function given in Figure 2.²⁰ Crucially, the utility function, u , stays as it is under EU. RDU can explain everything that EU can (because for $\alpha = \beta = 1$, the Prelec function is $w(p) = p$), disallows stochastically dominated choices, and can explain the celebrated Allais paradoxes.

Tversky and Kahneman (1992) incorporated insights from RDU into their 1979 version of prospect theory and came up with *cumulative prospect theory*.²¹ Following convention, we call the 1979 version *original prospect theory* and the 1992 version, simply, *prospect theory* (PT), which we focus on below. Under PT, a decision maker evaluates the lottery in (3) as $PT(\tilde{L}) = \sum_{i=-m}^n \pi_i v(y_i)$, where the decision weights, π_i , are constructed as under RDU but computed separately for gains and losses relative to the reference point. Decision makers who follow PT do not choose stochastically dominated options.

¹⁹ In the main, Kahneman and Tversky (1979) were interested in two-outcome lotteries but their model is generalizable to more than two outcomes.

²⁰ For readers who are itching for more information at this stage, before they get to Part I, $\pi_i = w\left(\sum_{j=i}^n p_j\right) - w\left(\sum_{j=i+1}^n p_j\right)$ for $i = 1, 2, \dots, n$.

²¹ The work of Starmer and Sugden (1989) in introducing non-linear probability weights in a prospect theoretic framework ought also to be mentioned here.

PT has been incredibly successful in understanding human behavior under risk, uncertainty and ambiguity; for a book length treatment, see Wakker (2010). Arguably, no other decision theory in economics has been so successful in explaining such a wide range of phenomena. Part 1 gives the details, but a small sample of these applications is as follows. Why does the equity premium puzzle (too high a return on equities relative to bonds) exist (Benartzi and Thaler, 1995)? Why is it so hard to find a taxi on a rainy day in New York (Camerer et al. 1997)? Why do people pay taxes (Dhami and al-Nowaihi, 2007, 2010a)? Why is there a disparity between willingness to pay and willingness to accept (Thaler, 1980)? How does ex-ante competition serve as a reference point in the formation of ex-post contracts (Fehr et al. 2011). Why are stabilizations delayed (Weyland, 1996)? Why are price elasticities asymmetric (Camerer, 2000)? Why do people hold on to loss making stock for too long and sell winners too early (Odean, 1998)? Why are people influenced by goals (Locke and Latham, 1990)? Why do people simultaneously gamble and insure (Kahneman and Tversky, 1979)? Why are incentives low-powered (de Meza and Webb, 2007)? Why do markets price skewness in asset returns (Barberis, 2013)? Furthermore, most of what economists think of as *risk aversion* might well be just *loss aversion* (Novemsky and Kahneman, 2005). It is hard to think of an area in economics that has not been touched directly or indirectly by prospect theory.

EU is a special case of RDU, which in turn, is a special case of PT. Thus, PT can explain everything that the other theories can explain, but the converse is false. For the purists, PT has axiomatic foundations (Chateauneuf and Wakker, 1999) and the individual components of prospect theory have axiomatic foundations, such as the utility function in Figure 1 (al-Nowaihi et al., 2008) and the Prelec probability function in Figure 2 (Prelec, 1998; al-Nowaihi and Dhami, 2006a). There is strong empirical support for the individual components of PT in choice data (Wakker, 2010), and in neural data (De Martino et al., 2006; Roiser et al., 2009; Tom et al., 2007). When a two-horse race is run between EU and PT, the predictions of EU are qualitatively and quantitatively unsatisfactory, while those of PT fit well the evidence (Dhami and al-Nowaihi, 2007). Mixture models indicate that there is possibly a 80%–20% split of people who follow PT and EU (Bruhin et al., 2010). Field evidence also supports PT, even among experienced subjects (Pope and Schweitzer, 2011; Fryer et al., 2012; Camerer, 2015).²² Close primate relatives, such as capuchin monkeys, also exhibit loss aversion, suggesting that loss aversion already existed prior to the separation of the human line from other primate relatives, and is hardwired in us (Chen et al., 2006).

A non-economist reading this introduction would think that surely PT must be the main decision theory taught in microeconomics courses. Wrong! Most of the standard texts in microeconomics either omit any mention of PT, or refer to it only in passing (Varian, 1992; Mas-Collel et al., 1995; Jehle and Reny, 2011). Such a situation is unlikely to have arisen in the natural sciences, at least in modern times. Incredibly, it is still possible to get a degree in economics in many universities without having undertaken a study of prospect theory, or even a course in behavioral economics. So why this relative neglect of PT, despite its growing usefulness and importance, and the continued reliance on EU? One can conjecture several reasons.

²² The List (2003) findings are commonly cited as showing that lab evidence on loss aversion does not generalize to the market; indeed, the title of the paper is: Does market experience eliminate market anomalies? However, the data in the paper shows that experience, whether in the lab or in the market, reduces loss aversion and the magnitudes of reduction are comparable. Significantly loss aversion still remains despite market experience, as is established by many other empirical studies.

1. The institutional structure in economics has a strong status-quo bias, partly because of homespun and self-serving methodological positions that do not recognize the primary role of evidence and of refutability (see discussion in Section 2 above). So, for instance, EU is often justified as telling a useful story of human behavior under risk and uncertainty or justified on the grounds that people behave “as if” they follow EU. An EU analysis also underpins many important results such as the existence of a Nash equilibrium and problems in mechanism design, thus, refutations of EU could be alarming for many economists. These defences of EU are untenable on scientific grounds.
2. Many economists appear to believe that EU is successful in making qualitative predictions, which justifies its use, while PT impinges mainly on quantitative results. However, a-priori, it is often not clear which qualitative predictions made under an EU analysis will be empirically supported. Within the same model, EU may make predictions that are qualitatively and quantitatively incorrect, while PT makes the correct predictions in each case (Dhami and al-Nowaihi, 2007, 2010a).
3. Economists critical of PT often argue that an exogenous reference point in PT is an unacceptable auxiliary assumption. In contrast, Kahneman and Tversky (1979) identified an endogenous reference point, the *status quo*, which is eminently sensible in most applications. As they also recognize, one’s entitlement or bargaining power, historical norms, or the expected value of a lottery, may also prove to be compelling reference points. Every falsifiable theory makes auxiliary assumptions, which can and should be tested. Evidence shows that people do indeed have reference points. It is well known that rationality, by itself, has very little predictive power in the neoclassical model in the absence of auxiliary assumptions; see a nice discussion in Thaler (2015, Ch. 17). The current state of the art in neoclassical economics is not even remotely close to the position that “all auxiliary assumptions must be endogenized.” Indeed it might not be sensible to aspire to do so.

One response to this criticism has been the development of a theory of *endogenous reference points*; roughly this is similar to imposing rational expectations on the formation of reference points (Kőszegi and Rabin, 2006). This solution has served to satisfy many neoclassical theorists although it sits uneasily with the evidence on cognitive limitations that we review at length in this book. Whether endogenous reference points in the sense of Kőszegi–Rabin is a step in the right direction is entirely an empirical question that its authors are aware of. Initial empirical evidence (see Part I) is consistent with the model, but we must await more stringent tests of the theory before we can be more confident.

4. It is not uncommon to hear the following view from some quarters: “PT has more parameters than the other decision theories, so it is not surprising that it explains better?” Incidentally, it is conveniently glossed over that exactly the same a-priori argument applies to the use of EU over expected value theory that it replaced.

There are two main statistical methods of addressing this criticism. First, estimate the model for a subset of the sample and use the estimated parameters to test the predictions on the remaining data. Second, use statistical criteria that allow one to test the trade-off between a greater number of parameters and better explanatory power, even among non-nested models, such as the Akaike information criteria (AIC), or the Bayesian information criteria (BIC). One may then run a horse race between the alternative theories.

I am unaware of completed studies in the first category, but there are some in the second. In a horse race between eight theories including EU and RDU (but not PT), Hey and Orme (1994) conclude (p. 1321): “Expected utility (and its special case, risk neutrality) emerges from this analysis fairly intact . . . However, we should emphasize once more

that this interpretation of our results needs to be taken with caution—since our sample of 80 subjects was in no sense representative.” The overriding feeling from these results (see also the survey by Hey, 2014), is that victories for one decision theory over another are rather messy and require one to buy into functional restrictions, assumptions on the error processes, and data limitations that one is never entirely comfortable with.

It is not commonly realized that PT also applies to situations of ambiguity; this is shown explicitly in Tversky and Kahneman (1992). In a comprehensive and careful study, Kothiyal et al. (2014) run a horse race between 11 competing theories of ambiguity that include subjective expected utility, PT, and several theories that are natural extensions of expected utility to ambiguity; PT emerges as the clear winner in this competition. Erev et al. (2010) run an unusual choice prediction competition pitting several decision theories against each other, many of which are more popular in psychology. In choice tasks that involved predictions from actual decisions, a stochastic version of PT outperformed the other theories. In contrast, in choice tasks involving predictions from experience, models that relied on small samples (see Section 5.5 below) performed best.

Under EU, attitudes to risk are captured entirely by the shape of the utility function. However, if we make any departure from linearity in probabilities in an EU type functional, as e.g., in RDU or PT, then attitudes to risk are *jointly* determined by the shapes of the probability weighting function and the utility function. In particular, the *four fold classification of risk* is one of the central insights of PT (Tversky and Kahneman, 1992), and despite subsequent empirical criticisms, it remains a robust finding (Fehr-Duda and Epper, 2012). Thus, under PT, one needs to elicit the probability weighting function and the utility functions separately. This can be accomplished by the trade-off method (Wakker and Deneffe, 1996; see Part 1).

Using a cleverly designed elicitation domain, the trade-off method is valid for EU and RDU. Employing this procedure, Abdellaoui (2000) checked if people transform probabilities; in particular, this test is not subject to the criticism that RDU does better because it has a higher number of parameters. The test showed that people transformed probabilities, which is inconsistent with EU but consistent with RDU. Abdellaoui then conducted a stringent test of the difference between RDU and PT and found that the probability weighting function was different for gains and losses, which is consistent with PT but not RDU. He then checked for simple lotteries, a condition (the duality of probability weighting functions) that would have established that RDU and PT coincide, but this condition was rejected.²³

A greater challenge to PT is posed by the following developments that appear to be lost in the rather narrow and confused mainstream critique of PT.

1. Michael Birnbaum has done a series of experiments whose results are not consistent with PT (Birnbaum, 2008). However, these results are not consistent with any of the other mainstream decision theories. His own proposals (e.g., *configural weights models*) are likely to appear ad hoc to many economists in the manner in which they allocate decision weights to outcomes. Further, there are no preference foundations of his proposed theory (see Part 1). For simple lotteries, the predictions of configural weights models are identical to PT.

²³ In personal communication, Mohammed Abdellaoui shared some findings of ongoing research that compares PT with alternative theories based on a comparison that takes account of the extra parameters under PT. These results are consistent with those reported above for Abdellaoui (2000).

2. There could be potential interaction between probability weights and the size of outcomes (Fehr-Duda et al., 2010). These results hit at the very foundation of models where the relevant objective functions are separable in probabilities and outcomes, such as EU, RDU, PT, and most other known decision models.
3. As Kahneman and Tversky (1979) so presciently realized, human behavior near the end-points of the probability interval is variable. Some people treat very low probabilities as very salient (notice the steep slope of the Prelec function in Figure 2 as $p \rightarrow 0$). However, others simply ignore very low probabilities (but this cannot be captured in Figure 2). Since the probability weighting functions, such as the Prelec function, underlie many behavioral theories, such as RDU, PT, and configural weights models, this evidence is troublesome for all of them. The neglect of low probability events, at least for a sizeable fraction of the population, impinges on a wide range of phenomena, such as buying inadequate insurance for low probability natural hazards, running red traffic lights, deterrence effects of capital punishment, and driving and talking on mobile phones (see Part I for details). One way of accounting for these phenomena is the proposal by Dhami and al-Nowaihi (2010b) to modify the Prelec function to a richer, axiomatically founded, probability weighting function that allows both overweighting and underweighting of very tiny probabilities. Whether this proposal will work is an open empirical question.

The rather detailed case study of PT in this section is not only useful for those seeking insights into behavioral theories of decision making, but also to illustrate the approach that I take in this book. I shall ask uncomfortable questions of both neoclassical and behavioral models, highlight the progress that has been made, and, whenever possible, use empirical evidence to judge between alternative theories.

5.2 Human sociality and inequity averse preferences

The exclusive pursuit of self-interest and the lack of any consideration of character virtues in neoclassical economics, is at variance with the evidence. Gintis (2009, p. 49) captures the heart of the issue: “Self-regarding agents are in common parlance called *sociopaths* . . . We conclude from behavioral game theory that one must treat individuals’ objectives as a matter of *fact*, not *logic*.” Evidence indicates that humans: exhibit *altruism* and *envy*; have an inherent tendency to cooperate yet display *conditional reciprocity* (respond to kindness with kindness and unkindness with unkindness); judge the kindness or otherwise of the *intentions* of others; and value human virtues such as “promises” for their own sake. Despite a great amount of evidence for *human sociality*, the acceptance of these forms of other-regarding behaviors is relatively slow among mainstream economics.

Many economists will probably dispute this criticism on the grounds that neoclassical economics allows people to have *other-regarding preferences*, and cite the rich tradition for this work (Veblen, 1899; Duesenberry, 1949). While this is certainly true, however, neoclassical economics has not generated models that can explain the richness of modern experimental findings on social preferences and human virtues, which is a minimum standard for aspiring contenders (Fehr and Schmidt, 2006). By contrast, behavioral economic theory has provided a rich range of models that are consistent with the evidence. Behavioral economics does not argue that human virtues, such as honesty and keeping one’s promises, are absolutes. Human adherence to virtues also depends on the cost of being virtuous and on the behavior (virtuous or otherwise) of others, e.g., tax evaders feel less guilty, if there is greater incidence of tax evasion in society.

Some of the earliest experiments on fairness concerns were conducted in the mid 1980s by Daniel Kahneman, Jack Knetsch, and Richard Thaler, based on telephone surveys in Canada.²⁴ Consider the following survey questions from Kahneman et al. (1991):

Question 1a. “A shortage has developed for a popular model of automobile, and customers must now wait two months for delivery. A dealer has been selling these cars at list price. Now the dealer prices this model at \$200 above list price.” Out of 130 respondents, 29% found the behavior of dealers acceptable, while 71% found it unacceptable.

Question 1b. “A shortage has developed for a popular model of automobile, and customers must now wait two months for delivery. A dealer has been selling these cars at a discount of \$200 below list price. Now the dealer sells this model only at list price.” Out of 123 respondents, 58% found the behavior of dealers acceptable while 42% found it unacceptable.

This example illustrates a number of interesting points.

1. Under self-regarding preferences, the responses to the two questions should be similar, but they are different.
2. The two different economic environments in Q1a and Q1b may generate ‘entitlements’ or status-quo reference points, such that receiving less than the entitlement is considered unfair. The entitlement is not necessarily mutually agreed upon. In other contexts, this can give rise to a *self-serving bias* in negotiations, which can potentially account for strikes, wars, and bargaining failures (Babcock and Loewenstein, 1997).
3. This example highlights some issues that play no role in neoclassical economics: For instance, forgoing a discount is less aversive relative to a surcharge. This leads to a better understanding of many business practices. Consider, for instance, charging for credit card usage in the early days of credit cards (Thaler, 2015). Initially, there was a 3% *surcharge* on the use of credit card payments in retail stores, which a consumer could avoid by paying cash. The credit card companies successfully lobbied the retail price of a \$1 item to be raised to \$1.03. So cash buyers could pay a *discounted price* of \$1, while card users could pay the *full price* of \$1.03. Card users thus do not pay a surcharge for their purchase. As Thaler puts it (p. 18): “Paying a surcharge is out-of-pocket, whereas not receiving a discount is a ‘mere’ opportunity cost.” Clearly direct out-of-pocket costs are more aversive relative to the more indirect concept of opportunity costs that is less well understood by many.

There is now overwhelming evidence of social preferences from lab and field data that has been widely replicated around the world with small and large stakes. Experimental games that have been used to build up the evidence base for social preferences include the *dictator game*, the *ultimatum game*, the *gift exchange game*, the *trust game*, and the *public goods game with and without punishment*.

Of the experimental games on human sociality, the *ultimatum game* (Güth et al., 1982) is the most replicated. In the ultimatum game, a proposer is given a fixed sum of money to split with the responder. The responder can then accept or reject. The backward induction outcome using classical game theory is that the proposer will offer the minimum divisible unit of money (e.g., 1 cent out of a dollar) to the responder; after all, something is better than nothing for the responder, so he accepts. By contrast, the average offer made by the proposer is 30–40 percent of the pie. Low offers are rejected by the responders who often invoke the “unfairness of low offers” as their justification for rejection; in neoclassical economics, people are amoral, so these considerations

²⁴ See the interesting background to these surveys in Thaler (2015, Ch. 14) and a discussion of the other experiments that they conducted.

do not arise. The results survive with sensibly high stakes, but rejection rates fall with extremely high stakes, although proposers still offer 20% of the pie (see Part 2 for details). These results have been replicated across the globe (Henrich et al., 2001) and reveal large cross-cultural differences in fairness norms.

The evidence from these experimental games shows that self-regarding preferences are not insignificant. However, a majority of the subjects also exhibit other-regarding preferences. A critical insight of this literature is that introducing a minority of individuals with other-regarding preferences in a pool of self-regarding individuals can potentially produce large changes (Fehr and Schmidt, 1999; Fehr and Gächter, 2000b). Empirical success in many areas that traditionally focused exclusively on self-regarding preferences, e.g., contract theory and political economy, has been considerably enhanced by incorporating other-regarding preferences. Over many years, Herbert Gintis and Samuel Bowles have made a compelling case for incorporating reciprocity into economic models, based on its evolutionary foundations (Gintis, 2009; Bowles and Gintis, 2011).

How should one proceed in light of this empirical evidence? Economists have traditionally been justifiably worried about preference-manipulation to explain economic phenomena. However, these concerns do not, and should not, apply to the modern literature on other-regarding preferences. The main models of other-regarding preferences have been subjected to extensive and stringent empirical testing, and fit well the evidence from a range of experimental games. Yet, it is not unusual to hear the comments that the incorporation of other-regarding preferences will “compromise the hard-won discipline encapsulated in the neoclassical framework with self-regarding preferences,” or that “other-regarding preferences are ad-hoc,” or that “they take the easy way out.” There is no empirical basis for these arguments; and we seek to demonstrate this in Part 2. Anyone who has been following the publications in the leading journals in economics in the last few years would have noticed that this state of affairs is fortunately changing.

Inequity aversion is an important human motivation. People chose to live in a welfare state with democratically elected governments where taxes finance redistribution to the poor and socially disadvantaged, and provide social insurance. Most of us will find it morally repugnant to live in a system without these features, even if we disagree on the exact levels of social redistribution. People also give their money and time to charitable causes; in the US, charitable donations alone account for 1% of GDP over and above redistributive taxes paid by people. Clearly, we are also driven by altruistic concerns. At the same time, we are envious of others who might be better-off than us. This is particularly the case when we believe that the higher income of others is not deserved; there is often a public outcry over high salaries and bonuses to top management in the financial sector especially at a time when they might be perceived as being responsible for a financial crises.

Fehr and Schmidt (1999) propose a simple, tractable, yet powerful model to capture these concerns. Suppose that n individuals are ordered in such a way that they have incomes $y_1 < y_2 < \dots < y_n$. Then, the Fehr–Schmidt utility function of any individual i is given by

$$U_i = y_i - \frac{\alpha_i}{n-1} \sum_{j>i} (y_j - y_i) - \frac{\beta_i}{n-1} \sum_{j<i} (y_i - y_j), \alpha_i \geq 0, 0 \leq \beta_i < 1. \quad (4)$$

An individual is said to have *inequity averse preferences*, a form of other-regarding preferences, if at least one of β_i and α_i is non-zero. In contrast, the neoclassical model with self-regarding preferences requires $\beta_i = \alpha_i = 0$ for all i . The first term in (4) captures self-regarding preferences and the next two terms capture, respectively, disadvantageous and advantageous inequity. Most experiments find that $\beta_i < \alpha_i$, so individuals suffer relatively more from disadvantageous

inequity; Eckel and Gintis (2010) review empirical estimates of the parameters from various studies. Empirically, $\beta_i < 1$ because we do not observe individuals throwing away their own incomes to reduce disutility from advantageous inequity. As with most other leading models in behavioral economics, this model has axiomatic foundations (Neilson, 2006). The Fehr–Schmidt model can be extended to a non-linear version, and this avenue has been fruitfully explored in several applications of the model.

Another popular model of inequity aversion is a model of *equity, reciprocity, and competition* (ERC), due to Bolton and Ockenfels (2000). Suppose that there are n players, $i = 1, 2, \dots, n$. The monetary payoff of player i is $y_i > 0$, and the sum of the total monetary payoffs is $S = \sum_{i=1}^n y_i$. The relative monetary payoff of player i is given by $r_i = y_i/S$. Then, the utility function of player i is given by

$$u^i(y_i, r_i). \quad (5)$$

u^i is twice continuously differentiable in y_i, r_i , and increasing and concave in y_i . For any given monetary payoff, u^i is maximized at $r_i = 1/n$ and is strictly decreasing and strictly concave in r_i around this point. Thus, individuals prefer to increase their monetary payoffs, yet for any monetary payoff, the individual prefers equal division; different individuals may view this trade-off differently, hence, the index i on the utility function. In the ERC model, individuals compare their payoff with the total payoff of others, but they do not care about interpersonal comparisons of their own payoffs with others (i.e., there is no notion of advantageous or disadvantageous inequity). In this sense, ERC has a less nuanced notion of inequality relative to the Fehr–Schmidt model.

The Fehr–Schmidt model is a bold one. It sticks its neck out and proposes a particular functional form that is eminently testable; this is quite distinct from the neoclassical approach, which argues that its underlying framework can potentially account for other-regarding preferences, yet does not specify how. The Fehr–Schmidt model has two eminently attractive features. First, it is able to explain the results from the key experimental games on human sociality such as the dictator game, ultimatum game, trust game, and public goods game with punishment; Part 2 gives the details. Thus, it meets the minimum standard for any aspiring theory of other-regarding preferences. Second, it predicts cooperation and sharing in bilateral/small group encounters and self-regarding outcomes under market competition. For instance, in ultimatum games, the proposer offers the responder a share of the pie because advantageous inequity is aversive and the responder rejects low offers because disadvantageous inequity is aversive. Ex-post punishment of non-contributors by contributors in public goods games can also be explained by disadvantageous inequity suffered by contributors. By contrast, under self-regarding preferences, we should never observe such punishments because bygones are bygones; see Section 6.2 below for more details. However, in an ultimatum game with responder competition, low offers are not rejected by any responder, because another responder may accept the offer, removing the opportunity to punish the proposer for an unfair offer. Anticipating this, the proposer makes a low offer. In this case, the outcome is as if individuals had self-regarding preferences although they really have other-regarding preferences.

The ERC model is also able to explain several important experimental results such as those from ultimatum games (including those under proposer competition), differences in the relative amounts given by proposers in ultimatum and dictator games, and the positive association between effort and rents in the gift exchange game. The ERC model also allows for asymmetric information, which gives it a wider domain of applicability. It can also explain better the data

in three player ultimatum games where a passive third party appears less cared-for relative to the predictions of the Fehr–Schmidt model (Bolton and Ockenfels, 1998). In other domains, it is less successful. Mean preserving spreads of payoffs leave a player’s utility unaffected but this is not supported by the experimental results (Charness and Rabin, 2002). In public goods games with punishments, players with ERC utility will not punish non-contributors because they are not concerned with a comparison of individual payoffs. Anticipating this, non-contributors will not contribute. Hence, the ERC model predicts that the outcome of the public goods games with and without punishments should be identical, a finding that is rejected by the evidence (Fehr and Gächter, 2000a,b). The ERC model may perhaps give even more realistic predictions in cognitively challenging situations where players might not know the precise income distribution of others but still have a rough idea of their relative income.

Clearly, the Fehr–Schmidt model and the ERC model cannot explain all aspects of social preferences because they are not designed to do so. In many situations, inferring intentions is important. Most people would punish others for intentional unkindness, but not for unintentional unkindness. Even the law makes a distinction between murder and manslaughter. However, inferring intentions of others requires a specification of belief hierarchies and a rigorous analysis requires the use of *psychological game theory*, which, in contrast to classical game theory, allows beliefs to directly enter the utility function (Geanakoplos et al., 1989; Rabin, 1993; Dufwenberg and Kirchsteiger, 2004; Battigalli and Dufwenberg, 2009); see Part 4 for the details. For instance, we may offer a tip to a taxi driver if we believe that he expects a tip. Indeed, the insights offered by Fehr–Schmidt preferences have been successfully combined with endogenous beliefs in the framework of psychological game theory (Falk and Fischbacher, 2006). In other domains, such as equal splits in ultimatum games, it appears that people follow certain *fairness norms* that are not directly captured either by the Fehr–Schmidt model or the ERC model (Rotemberg, 2008; Andreoni and Bernheim, 2009). The interested reader can pursue an empirical critique of both models in Cooper and Kagel (2015).

The Fehr–Schmidt model has led to several important theoretical advances in the literature on human sociality. For instance, it has played an important role in contract theory in the understanding of long-term contracts (Brown et al., 2004, 2012), and the choice of contracts in the presence of contractual incompleteness, such as the choice of bonus contracts over incentive contracts (Fehr et al., 2007). In each of these cases, the experimental evidence is consistent with the predictions of the underlying model.

Redistributive concerns are often center stage in national and local elections. In experimental games, people often prefer a smaller more equitably distributed cake to a larger less equitably distributed cake; these results fit well the Fehr–Schmidt model (Tyran and Sausgruber, 2006; Ackert et al., 2007). These empirical insights have led to new theoretical results in *behavioral political economy*, such as the existence of a Condorcet winner with inequity averse preferences, and the equilibrium redistributive tax when there is a mixture of voters who have self-regarding and Fehr–Schmidt preferences (Dhami and al-Nowaihi, 2010c,d). There has been progress in establishing key results on the effects of other-regarding preferences in a Walrasian competitive equilibrium (Dufwenberg et al., 2011); Fehr–Schmidt preferences provide an important example of a key separability assumption in this framework (Dhami and al-Nowaihi, 2010d). The standard machinery of stochastic dominance for self-regarding preferences (e.g., first and second order stochastic dominance and Lorenz dominance) does not apply when people have other-regarding preferences. However, the analogues of all these concepts for inequity averse preferences have been developed rigorously (Dhami and al-Nowaihi, 2013). Part 2 describes the details of these developments.

5.3 The quasi-hyperbolic model and self-control problems

Temporal analyses of economic models, for descriptive and normative purposes, almost exclusively rely on the *exponential discounted utility* (EDU) model. The EDU model arose purely in response to technical considerations. Samuelson (1937) proposed the EDU model and it turned out to be the unique method of discounting among delay-discounting models that leads to *time consistent choices* (Strotz, 1956). Let $\mathbf{c}_0 = (c_0, c_1, \dots, c_T)$ denote the real-valued *temporal consumption profile* of an individual at time $t = 0$. Under EDU, this consumption profile is evaluated by the following utility function,

$$U(\mathbf{c}_0) = \sum_{t=0}^T \delta^t u(c_t); \delta = \frac{1}{1+\theta}, \theta > 0. \quad (6)$$

The EDU model captures the key idea of impatience in a simple and parsimonious manner. However, a single parameter, the constant *discount rate*, θ , captures all relevant psychological considerations. In contrast to many who followed him, Samuelson was very careful in introducing his proposal. On the descriptive validity of his model he writes (p. 159) “It is completely arbitrary to assume that the individual behaves so as to maximize an integral of the form envisaged in (2) [i.e., the EDU model in (6)].” On the normative implications of EDU he writes (p. 161): “. . . any connection between utility as discussed here and any welfare concept is disavowed.”

The evidence does not support the EDU model. Arguably, the EDU model is rejected even more strongly than expected utility theory. The evidence can be gleaned from many sources (Loewenstein and Prelec, 1992; Frederick et al., 2002). There are several anomalies of the EDU model that make it a poor candidate of choice on empirical grounds. For instance, empirical evidence shows that losses are discounted less than gains (*the sign effect*), greater magnitudes are discounted relatively less (*the magnitude effect*), discounting over an interval depends on how that interval is subdivided (*subadditive discounting*), temporal choices can be cyclical (*apparent intransitivity of preferences*), and choices may depend on the *shape of the consumption profile*, despite no changes in discounted utility. Furthermore, the discount rate depends on age (younger people are more impatient), intelligence (higher IQ is correlated with lower impatience), and addictions (gamblers, smokers, drug users, and heavy alcohol drinkers are more impatient).

In this section, we concentrate on an important anomaly of the EDU, namely, the rejection of the constant discount rate, sometimes known as the *common difference effect*. Consider the following example from Thaler (1981): “Most people prefer one apple today to two apples tomorrow but they prefer two apples in 51 days to one apple in 50 days.” This *temporal preference reversal* violates constant discounting. A range of other examples highlight an increase in impatience when a future reward is shifted in date towards the present. For instance, short-lived New Year resolutions on smoking or drinking alcohol made with perfect sincerity at an earlier date, procrastination in undertaking various tasks or activities, splurging too much on consumption now, and saving inadequately for the future.

The sort of preferences exhibited in Thaler’s apples problem are known as *present-biased preferences*. Consider outcome–time pairs of the form (z, t) , where z is an outcome (possibly but not necessarily monetary), and t is time. Consider two possible outcomes z, z' and three time periods t, τ, t' , all strictly positive. The decision maker considers alternative outcome–time pairs at time 0. Let $\succsim$ be a preference relation between outcome–time pairs (e.g., if $(8, 3) \succsim (10, 4)$, the 8 units




Figure 3 A plot of the discount rate for the one-parameter hyperbolic discounting function with varying time delay.

received 3 periods from now are “at least as good as” 10 units received 4 periods from now). A great deal of evidence indicates *preference reversals* of the following form.

$$(z, t) \succsim (z', t') \text{ but } (z', t' + \tau) \succsim (z, t + \tau); \tau > 0, t' > t, z' > z. \quad (7)$$

One may fix z, t, t' in an experiment and ask subjects to state a value of z' such that they are indifferent between (z, t) and (z', t') . A straightforward calculation can then be used to infer the implied discount rate. Early experiments by Thaler (1981) showed that the implied discount rates for a one month, one year, and ten year horizon, were, respectively, 3.45, 1.2, and 0.19. Thus, individuals are more impatient, the closer the reward is to the present.

Fundamental work was done by George Ainslie in fitting discount functions to human and non-human temporal choice data that is summarized in Ainslie (1992). Figure 3 plots the implied per-period discount rate when the future reward is received with various time delays, for one of his functional forms that fitted his data well; see Part 3 for the details. As in Thaler’s experiments, the discount rate is higher when the reward is more proximate to the present and the discount rate is hyperbolically declining, hence, the term *hyperbolic discounting*.

One behavioral response to the empirical data has been to suggest a two-parameter *general hyperbolic discount function* that nests EDU as a special case (Loewenstein and Prelec, 1992).²⁵ However, a much simpler version, *quasi hyperbolic discounting* or the (β, δ) form, based on initial work by Phelps and Pollak (1968), but popularized by Laibson (1994, 1997), has paved the way for major applications of present-biased preferences. Under quasi-hyperbolic discounting, the consumption profile $c_0 = (c_0, c_1, \dots, c_T)$ is evaluated in the following manner.

$$U^0(c_0) = u(c_0) + \beta \sum_{t=1}^T \delta^t u(c_t), \quad 0 < \beta < 1, \delta = \frac{1}{1 + \theta}, \theta > 0. \quad (8)$$

²⁵ See also al-Nowaihi and Dhami (2006b) for a simple axiomatic derivation of the general hyperbolic discount function that corrects some errors in the seminal work of Loewenstein and Prelec (1992).

Here, δ plays the same role that it plays in the EDU model. The real bite comes from the term β that shrinks the utility from future consumption levels, relative to the present. This creates an additional *present bias for current consumption* over and above impatience that is captured in the discount factor δ . This feature explains preference reversals such as in Thaler's apples example, which the reader can readily verify. Unlike the general form of hyperbolic discounting, under quasi-hyperbolic discounting, the discount rate jumps up discontinuously just as one approaches the current date; the reader can do a few quick plots to check this.

Empirical evidence typically shows declining discount rates as in Figure 3, which implies *present-biased preferences*. However, there is somewhat less agreement about whether discount rates just decline, or if they decline in a hyperbolic manner. A great deal of empirical evidence from the lab and the field is consistent with quasi-hyperbolic discounting (Viscusi et al., 2008; Fang and Silverman, 2009; Tanaka et al., 2010).

Quasi-hyperbolic discounting has given rise to a wide range of applications of behavioral time discounting and a real improvement in our understanding of temporal decisions; an entire chapter is devoted to these applications in Part 3. Do individuals with present-biased preferences have adequate awareness that they might also exhibit a present bias in the future? This turns out to be a rather critical question whose consequences have been well explored (O'Donoghue and Rabin, 2001). In this context, it is best to think of individuals as a collection of *multiple selves*, each self coinciding with a distinct time period (e.g., a morning-self who presses the snooze button, overriding the decision to wake up early made by a night-self). Empirical evidence suggests that individuals lie somewhere between the two extremes of full awareness of future self-control problems, and no awareness of such problems; thus, they may be termed as *partial-naifs*.

Partial-naivety about future self-control problems, in conjunction with quasi-hyperbolic preferences, makes interesting predictions in a range of problems (see Part 3). There are important applications to macroeconomics, where one can derive the behavioral analogue of the Euler equation and explain several critical puzzles, for instance (Laibson, 1997, 1998): Why does consumption track income so closely? Why do individuals undersave for retirement? Why is there a sharp drop in consumption at retirement? Why do individuals hold illiquid assets and credit card debt simultaneously? Why is high consumption variability not associated with high consumption growth? What is the appropriate relation between intertemporal substitution and risk?

But there is also an equally important range of microeconomic questions that can be addressed with this framework: Why do many people procrastinate (O'Donoghue and Rabin, 2001)? Why do many people retire early (Diamond and Köszegi, 2003)? What are appropriate sin taxes (taxes on harmful addictions) in the presence of present bias (O'Donoghue and Rabin, 2006)? Why do smokers sometimes welcome an increase in the retail tax on cigarettes (Gruber and Mullainathan, 2005)? Why might some people buy expensive annual gym memberships, when a pay-per-visit contract would have saved them money (DellaVigna and Malmendier, 2004)? Why might farmers use fertilizers suboptimally (Duflo et al., 2011)? Hopefully these questions may whet your appetite to read Part 3 in more detail.

Given the evidence, the rigorous axiomatic developments, and wide-ranging applications of present-biased preferences, the dominance of EDU might well be an interesting topic for a student of economic thought writing a dissertation in 2030. But I would be lying if I said that the current generation of behavioral economists do not find this state of affairs to be frustrating. There have been several criticisms of hyperbolic discounting. But whether time preferences are exactly

hyperbolic or not will be eventually less important than whether they are present-biased.²⁶ Ultimately, a more complete picture of choice over temporal sequences might also need to relax the high degree of cognitive requirement that is embodied in all delay-discounting models. *Attribute-based models* that invoke choice heuristics may well outperform existing delay-discounting models (Ericson et al., 2015).

5.4 Level-k and CH models: disequilibrium in beliefs in strategic interaction

Non-cooperative game theory represents one of the most significant advances in economic theory. It has enabled economists to make very precise and testable predictions in strategic situations. In the economics profession there is widespread acceptance and use of the main solution concept, *Nash equilibrium*. In a Nash equilibrium and its extensions, players play a best response to their beliefs, and the beliefs and actions of the players are mutually consistent. Furthermore, Bayes' rule is used to update beliefs, whenever possible.

How do the actions and beliefs of players come to be mutually consistent? The most common justifications are an appeal to an unmodeled, underlying, process of learning, or that a Nash equilibrium is simply a social norm. None of these arguments is satisfactory. First, as we show in Part 5, the outcomes of experiments on learning do not necessarily converge to a Nash equilibrium. Second, many norms are not Nash equilibria, particularly in the domain of cooperative human outcomes. Furthermore, realistic epistemic conditions that justify Nash equilibrium play are not easily satisfied (e.g., common knowledge of rationality only implies the weaker notion of rationalizability); see Gintis (2009, p. 83) for related commentary.

Whether people do play a Nash equilibrium is an empirical matter. Let me cull out, very briefly, some interesting aspects of the evidence that is reviewed over several hundred pages in Parts 4 and 5; I suppress the very large number of references.²⁷ In many games, the experimental evidence, particularly in the early rounds of a game, is not consistent with a Nash equilibrium. In many others, the outcomes do not converge to a Nash equilibrium, even when the experiments are repeated over a large number of rounds. Even in relatively simple games that involve more than two to three steps of iterated deletion of dominated strategies, the outcome is often not a Nash equilibrium. However, it is also the case that in more complicated games, experimental evidence is sometimes surprisingly consistent with a Nash equilibrium, such as in signaling game experiments. In other games where complexity and cognitive requirements are quite high, such as bargaining under one-sided asymmetric information, the evidence is inconsistent with a Nash equilibrium. Yet, rather miraculously, when the degree of complexity is raised even further, as in bargaining under two-sided asymmetric information, the evidence fits well with a Nash equilibrium. In some games, where the Nash equilibrium is incredibly difficult to compute, such as in traffic network games, one sometimes observes relatively quick convergence to the Nash equilibrium; see, for instance, Mak et al. (2015). In this game, there are 3.3 billion pure-strategy equilibria, yet, incredibly, all players coordinate to converge on one of these equilibria! In other fairly complex games, the outcome is not even close to a Nash equilibrium, e.g., players might be subject to a winner's curse in common value auctions.

²⁶ In this book we also consider several other theories of present-biased preferences. These include the *planner-doer model* (Shefrin and Thaler, 1981), *models of dual-selves* (Fudenberg and Levine, 2006) and the Gul-Pesendorfer model of *temptation preferences* (Gul and Pesendorfer, 2001).

²⁷ For an excellent, but slightly dated review of the evidence, see Camerer (2003a).

In other cases, such as *market entry games*,²⁸ the number of entrants exactly matches that predicted by a Nash equilibrium. Kahneman (1988), who ran these experiments, exclaimed that “it looks like magic.” However, it is now well known that behavioral theories, such as the *cognitive hierarchy model*, discussed in Part 4, explain the evidence on market entry games equally well.

In sum, the case for focusing exclusive attention on a Nash equilibrium in economics is difficult to justify. In far too many cases, the behavioral alternatives provide a much superior explanation of the relevant evidence. One approach to relaxing the cognitive requirements embodied in a Nash equilibrium is to allow players to play noisy best replies, rather than strict best replies. This is the approach taken in *quantal response equilibrium* (QRE). A second approach, that we focus on briefly in this section, is *level- k models* (Nagel, 1995; Stahl and Wilson, 1995) and the closely related *cognitive hierarchy model* (Camerer et al., 2004).

The classic game that is used to illustrate the level- k model is the *p -beauty contest*. Suppose that in a winner-take-all contest, each of n participants chooses a number, x_i , from the interval $[0, 100]$. The winning guess is the number closest to $p\bar{x}$, where $p > 0$, and $\bar{x}$ is the mean guess. Let us assume that $p = \frac{2}{3}$. In a symmetric Nash equilibrium, all participants should choose the guess 0. This is easily checked. Conditional on the equilibrium guess of the other players, any player who chooses $0 + \varepsilon$, where $\varepsilon \in \{0, 1, \dots, 100\}$, moves the target ($\frac{2}{3}\bar{x}$) upwards from 0 to $\frac{2}{3n}\varepsilon$. But for $n \geq 2$, the new target is still closer to 0 than to ε , so no profitable deviations are possible. Likewise, the reader can check that there is no symmetric pure strategy Nash equilibrium in guesses that is strictly positive, because it would be destroyed if one individual chooses a lower guess.²⁹

In what follows, we shall, on empirical grounds, ignore Breitmoser’s effect that was raised in the last footnote. The typical findings in p -beauty game experiments are shown in Figure 4 for $p = \frac{2}{3}$ and for a diverse subject pool, comprising high ability Caltech students, economics Ph.D. students, portfolio managers, and high school students. The results are in stark contrast to the Nash equilibrium prediction. The original experimental findings are due to Nagel (1995), the first replication was by Ho et al. (1998), and since then, the experiment has been replicated many times. In particular, a peculiar feature of the data is that there are spikes around guesses of 33, 22, 14, . . . This is inconsistent with a Nash equilibrium.

Level- k models offer a parsimonious way of explaining these findings. In a level- k model, players have types Lk , $k = 0, 1, \dots, K$. Each type, except type $L0$, plays a best response, assuming that all other players are of type $Lk - 1$. The specification of type $L0$ depends on the problem and in some applications, this type is assumed to be simply randomizing equally among the available actions. In other domains, uniform randomization by type $L0$ players is behaviorally implausible, and one may have to modify it. For instance, in Crawford and Iriberri (2007b), type $L0$ players prefer the label-salient action in games where labels convey salience. The theory does not require type $L0$ players to physically exist; this type may simply represent a model of other players that type $L1$ players have in mind. The frequencies of the various types of players can be estimated

²⁸ In a market entry game, a set of firms with unit supply must simultaneously enter into a market with a fixed demand of $d > 0$ integer units. If just d firms enter, all make profits. If more than d firms enter, all firms earn zero, and if less than d firms enter, then the entrants earn a low but positive level of profits.

²⁹ Breitmoser (2012) questions the suitability of a level- k approach for the p -beauty contest. His main point is that by ignoring the effect of one’s own guess on the average guess, the predictions of level- k models are misspecified. However, it appears to be the case, in most experiments, that subjects ignore the effects of their own actions on the average and the guesses of subjects in experiments are in conformity with the predictions of level- k models; see Part 4 for the details.




Figure 4 Distribution of choices by different subject pools in a p-beauty contest game.

Source: Reprinted from *Trends in Cognitive Sciences*, 7(5): 225–31. Colin F. Camerer, “Behavioural studies of strategic thinking in games.” Copyright © 2003, with permission from Elsevier.

from the data. Empirical studies show that the largest percentage of players is of types L_0 , L_1 , and L_2 . Players of type L_3 or beyond, are rare.

Level- k models directly speak to the cognitive limitations of players and truncate the “infinite regress in beliefs” feature of classical game theory at some low, cognitively plausible, level. In particular, the solution concept that is proposed here is non-circular in the following sense that is well described by Selten (1998): “Basic concepts in game theory are often circular in the sense that they are based on definitions by implicit properties . . . Boundedly rational strategic reasoning seems to avoid circular concepts. It directly results in a procedure by which a problem solution is found.”

Let us apply level- k models to the p-beauty contest game. Type L_0 players simply randomize equally among the various guesses so their average guess is 50. Type L_1 players play a best response to a guess of 50, which is $\frac{2}{3}50 = 33.3$. Type L_2 players play a best response to a guess of 33.3, which is $\frac{2}{3}33.3 = 22.2$. Similarly, type L_3 players choose $\frac{2}{3}22.2 = 14.8$. The level- k model predicts spikes at 33.3, 22.2, and 14.8, which matches well with the data in Figure 4. If we accept that type L_0 players randomize equally among their choices, then this is an incredibly precise prediction that is confirmed by the data. This model is a “disequilibrium in beliefs” model so that, ex-post, many players will be surprised at the outcome, because their initial beliefs about other players are not confirmed.

A perseverant classical game theorist may not accept these results, arguing that a Nash equilibrium is about learning through repeated play; so, if you played the p-beauty game several times, the Nash equilibrium prediction will hold. This argument does not hold for several reasons. First, even if subjects use simple rules of thumb in a repeated p-beauty contest, and play a best response to the observed mean in the last period, then under reasonable assumptions the outcomes can be shown to converge to a guess of zero. Second, unlike the repeated p-beauty contest, real-world environments are not stationary, so a repeated p-beauty contest is not a stringent




Figure 5 Choices of game theorist subjects in a p-beauty contest game.

Source: Bosch-Domènech et al. (2002) with permission from the American Economic Association.

test of the theory. Third, a Nash equilibrium does not explain the temporal pattern of guesses. In a 10 round repeated p-beauty contest where groups of 3 subjects are matched by cognitive ability, a level- k model describes well the temporal pattern of guesses (Gill and Prowse, 2015). In the final two rounds, the high cognitive ability group reaches the Nash equilibrium prediction 29% of the time; the corresponding figure for the mixed and low cognitive ability groups is, respectively, only 15% and 13%.

A particularly simple version of the p-beauty contest occurs when there are two players, $n = 2$. For $p = \frac{2}{3}$, the winning guess is simply the lower of the two guesses. Thus, the cognitive requirements in guessing a Nash equilibrium are substantially reduced. Yet, before the p-beauty contest was part of the common folklore, when the experiments were run with professionals at conferences, the Nash outcome was not achieved for the vast majority of subjects. Four of these five conferences were economics conferences that included game theorists (Grosskopf and Nagel, 2008). 36.2% of professionals and 9.85% of students chose the Nash outcome. However, any choice greater than zero is a weakly dominated strategy, which is chosen by 62.8% of professionals and 91.15% of students. There is no statistical difference between the cumulative distribution of choices in the cases $n = 2$ and $n > 2$ for both pools of player. The results from game theorist subjects in Figure 5 show that their behavior is similar to other humans, although 22% of the choices in this group conformed to the Nash outcome, the highest among the six groups considered (Bosch-Domènech et al., 2002).

The *cognitive hierarchy* (CH) model is similar to the level- k model. However, in the CH model, each type, Lk , $k = 1, 2, \dots$, assumes that there is a distribution, typically Poisson, of players of all lower types, $Lk - 1$, $Lk - 2$, $\dots$, $L0$. The predictions of the level- k model and the CH model are similar for a large class of games and it is very difficult to put one ahead of the other on empirical grounds (Crawford et al., 2013). When the early-round results of a large class of 55 games are pitted against the predictions of the CH model and the Nash equilibrium, then there are two findings (Camerer et al., 2004). Whenever the Nash equilibrium predicts well, CH predicts well too, and when the Nash equilibrium predicts poorly, CH predictions are more accurate. Field data supports the superior predictions of the CH model relative to a Nash equilibrium (Goldfarb and Yang, 2009; Goldfarb and Xiao, 2011). Level- k models are supported by a great deal of other evidence (Costa-Gomes et al., 2001; Costa-Gomes and Crawford, 2006).

The applications of level- k and CH models have grown enormously. These range from testing the significance of focal points when they come into conflict with label salience in symmetric and asymmetric games (Crawford et al., 2008); the role of preplay communication in enhancing coordination (Crawford, 2007); the explanation of the “magical” results in the market entry game

that were originally thought of as a remarkable confirmation of Nash equilibrium (Camerer et al., 2004); and the explanation of the winner's curse in common-value auctions (Crawford et al., 2009; Crawford and Iriberri, 2007a).

Level- k models do not have a theory of level-0 players and the behavior of these players needs to be inferred from the data. So, a critic might argue, is it legitimate to use the same data to infer parameters of the behavior of level-0 players, and then use these estimated parameters to estimate the remaining parameters of interest? Crawford and Iriberri (2007b) solve this problem by first using a part of the sample to estimate the behavior of level-0 players. They then successfully make predictions for out-of-sample data (which in this case is a set of closely related games) that are confirmed.

Hargreaves Heap et al. (2014) question the issue of portability of level- k behavior across a wider range of games than are considered in Crawford and Iriberri (2007b). They do assume, however, that the specification of the level 0 players is identical across these games. In some games, particularly coordination games, *team reasoning* (see Part 4) rather than level- k models appear to better describe the evidence. This should alert us to the importance of different contexts that may trigger different behavioral responses—in some kinds of games, level- k reasoning may be invoked, while in others, team reasoning may be invoked. Perhaps a hybrid of level- k and team reasoning models, provided team reasoning can be formally defined in a satisfactory manner, may explain the data even better; similar conclusions are reached by Crawford et al. (2008, p. 1448).

5.5 *The heuristics and biases program: radical behavioral economics*

The neoclassical approach views economic agents as Bayesian, subjective expected utility maximizing decision makers who have unlimited cognitive abilities and follow the rules of classical statistics. These economic agents also routinely engage in mathematical optimization, using the latest developments in mathematics, statistics, engineering, and computer science. This is one possible hypothesis of human behavior to start with. However, as always, its descriptive validity hinges on the empirical evidence. Following Milton Friedman's approach (see Section 2 above), and some may argue, its misuse, it is not an uncommon response in economics to handwave away the evidence and argue that real humans behave "as if" they were textbook economic agents in neoclassical theory.

The crucial significance of the *heuristics and biases program* of Daniel Kahneman and Amos Tversky was that they showed systematically that this "as if" view is inconsistent with the evidence (Tversky and Kahneman, 1974). *Heuristics* are any rules of thumb that individuals actually use, as distinct from normative recommendations. Heuristics are *fast*, in terms of the computation time required, and *frugal*, in terms of information requirements. The term *biases* in the program refers to the difference between actual human behavior and the predictions of neoclassical theory.

The heuristics and biases program, one of the most important advances in all of social science, created great disagreement among psychologists. Some of the attacks on Tversky and Kahneman were vicious, and even personal (see Part 7 and Stanovich and West, 2000, p. 649).³⁰ Most economists ignored this work, and continue to do so, despite its ubiquitous presence around us in the media, public debates, and in behavioral finance. Luckily for behavioral economics, a young

³⁰ Here are some quotes from published journal articles (references in Part 7). 1. "it is Kahneman & Tversky, not their subjects, who have failed to grasp the logic of the problem." 2. "If a 'fallacy' is involved, it is probably more attributable to the researchers than to the subjects." 3. "In the examples of alleged base rate fallacy considered by Kahneman and Tversky, they, and not their experimental subjects, commit the fallacies."

economist, reading this work in 1976 did not ignore it; he wrote about it recently (Thaler, 2015, p. 22): “As I read, my heart started pounding the way it might during the final minutes of a close game. The paper took me thirty minutes to read from start to finish, but my life had changed forever.”

Tversky and Kahneman identified a range of judgment heuristics, which are underpinned by a great deal of evidence. Part 7 gives the details. Here, we consider some aspects of this research program. Let us call a textbook neoclassical economic agent who follows the rules of classical statistics a *Bayesian*.

People are often asked to answer questions of the form: How likely is that event A was generated by category/class/process B? In answering this question, people often use the *representativeness heuristic*, i.e., they give an answer based on how much event A looks like B.³¹ Relative to a Bayesian, people often assign too high a probability that small samples share the properties of large samples. Such individuals behave as if they obeyed a *law of small numbers*, rather than the statistically correct, *law of large numbers*. This has important implications. For instance, individuals find it difficult to produce hypothetical sequences of random numbers; these sequences exhibit too much *negative autocorrelation* (Tversky and Kahneman, 1974; Bar-Hillel and Wagenaar, 1991; Rapoport and Budescu, 1997). This behavior produces the *gambler's fallacy*, i.e., avoiding bets on recent winning numbers in gambling or recent winners in the stock market (Clotfelter and Cook, 1993; Terrell, 1994; Odean, 1998).

People may often observe many consecutive, identical, outcomes. In this case, they may use the representativeness heuristic to infer that the sample is drawn from a population that contains a disproportionate share of such outcomes. Thus, they may expect more of these outcomes in the future; this *positive autocorrelation* in outcomes is known as the *hot hands fallacy*. Clearly, a Bayesian would also update in the same direction, however, under representativeness, individuals react more strongly. The hot hands phenomenon has been documented in sports, such as basketball or baseball, where players who score heavily in a game may be believed to be on a hot streak. Lotto stores that sell winning lotto tickets experience unusually high sales for up to 40 weeks following the lotto win, although the probability of purchasing a winning ticket from that store is unchanged (Guryan and Kearney, 2008). For an “animate” stochastic process that relies on human skills (e.g., basketball and baseball), one is more likely to observe a hot hands fallacy relative to an “inanimate” process (e.g., casino betting), where a gambler's fallacy is more likely to be observed (Ayton and Fischer, 2004). The law of small numbers can also explain the existence of overconfidence (Rabin, 2002).

Individuals use the *availability heuristic* when they judge the probability of an event by the ease with which similar events are accessible or “available” to the mind. Accessibility is easier for salient and vivid events, which highlights the importance of media coverage of events. There is a close link between the perceived frequency of deaths by various causes and proxies for the availability of such events in memory (Lichtenstein et al., 1978).

The *affect heuristic* identifies the role of emotions in human decisions (Slovic et al., 2002; Loewenstein et al., 2001). People often classify and tag events by their *affect*. For instance, a holiday in New York may leave a positive affect, if it was enjoyable but, *ceteris paribus*, a negative affect, if one was mugged during the holiday. The affect heuristic argues that the pool of these positive and

³¹ For instance, suppose that three successive balls are drawn with replacement from either urn A that has 1 red and 2 black balls or urn B, which has 2 red and 1 black balls; it is common knowledge that urn A is chosen with probability 0.6. The three sampled balls turn out to be 1 red and 2 black. What is the probability that the balls came from urn A? The statistically correct answer is 0.75, but an individual subject to the representativeness heuristics gives an answer greater than 0.75 and, in the extreme case, an answer of 1.

negative affects determines our choices, rather than objective statistical probabilities of events, e.g., do I go to New York for my next holiday? The affect heuristic may economize on cognitive efforts in making a decision and quicken decision making, because it is easier to think “what do I feel about it?” rather than “what do I think about it?” (Kahneman, 2011). Even experts rate the risk of hazardous substances by the feeling of dread that these substances invoke (Fischhoff et al., 1978; Slovic et al., 1999). Loewenstein et al. (2001) make a powerful case for understanding choices among risky options in terms of their emotional impact.

Anchoring, the influence of initial suggestions as anchors on subsequent judgments, is one of the most robust and important heuristics. When anchors are self-generated, people adjust insufficiently towards the correct response, but when they are provided by others (e.g., the experimenter), people are often biased towards finding confirmatory evidence for them. In a famous experiment, the response of subjects to the question “how many African countries are there in the United Nations?” is influenced by the prior observed outcome of a random wheel of fortune that is rigged to come up with the numbers 10 or 65 (Tversky and Kahneman, 1974); clearly irrelevant information influences choice in this case. In field studies, experienced judges and estate agents are influenced by anchors (Englich and Mussweiler, 2001; Northcraft and Neale, 1987). Implausible and irrelevant anchors can have a significant effect on subjects’ choices (Strack and Mussweiler, 1997). Regulation, such as caps on damages awarded to litigants, serves an anchoring role and influences the rate of pretrial settlements and the level of damages (Babcock and Pogarsky, 1999; Pogarsky and Babcock, 2001).

Anchoring has been documented in a wide range of contexts, such as poetry reading sessions (Ariely et al., 2006); violation of a basic implication of efficient markets (Mussweiler and Schneller, 2003), and first offers in price negotiation (Galinsky and Mussweiler, 2001). Anchoring has also been used to explain a range of other behavioral phenomena, such as hindsight bias, non-linear probability weighting, and preference reversals.

Bayes’ rule fundamentally underpins modern economic theory. However, most people do not employ Bayes’ rule when the relevant information is presented in a probability format. In particular, individuals routinely exhibit *base rate neglect*, i.e., give insufficient attention to the prevalence of a particular trait or feature in the overall population (Kahneman and Tversky, 1972; Bar-Hillel, 1980; Tversky and Kahneman, 1980). Consider the following problem, adapted from Casscells et al. (1978) that was given to students at Harvard Medical School. Suppose that a police Breathalyzer test discovers false drunkenness in 5% of the cases when the driver is sober. However, the Breathalyzers are always able to detect a truly drunk person with complete certainty. In the general population, 1/1000 of drivers engage in driving while drunk. Suppose that a driver is checked at random, and takes a Breathalyzer test, which shows that the driver is drunk. How high is the probability he or she really is drunk? The modal response was 95%, the mean response was 56%, and only a sixth of the students gave the correct answer, which is approximately 2%. Most people ignore the low base rate of 1/1000 of drivers who drink and drive.

Hindsight bias occurs when individuals make a prediction about some future event, but after they observe the actual outcome, their postdictive judgment (i.e., recall of their initial prediction) is too close to the actual outcome (Fischhoff, 1975). Thus, people give the impression as if “they knew it all along.” This bias is extremely important in legal cases where the judge and jurors must infer if the defendant took enough preventive care in the past, having already observed the outcome of the defendant’s decision. A meta-study finds that experts also exhibit the hindsight bias (Guilbault et al., 2004). In finance, hindsight bias reduces estimates of volatility, and people engage in less corrective activity of incorrect models (Biais and Weber, 2009).

Confirmation bias is an important heuristic in which people interpret the evidence in a self-serving manner so as to confirm their initially held views (Lord et al. 1979). Higher cognitive ability does not reduce this bias (Stanovich and West 2007). As one may expect, since confirmation bias vindicates initially held views, it can also explain the prevalence of overconfidence and underconfidence (Rabin and Schrag, 1999).

As with many other great works, some of the initial charges against the heuristics and biases approach could be easily refuted. These included the charge that subject errors and misperceptions caused the biases, or that Kahneman and Tversky did not understand probability theory well enough. Other charges took longer to refute, some were already actually answered by Tversky and Kahneman in their published work, others helped to modify the theory without serious changes in the message. A prime example is the belief that if individuals are presented with information in *frequency format* rather than in a *probability format*, then human biases disappear (Cosmides and Tooby, 1996; Gigerenzer and Hoffrage, 1995). Empirical evidence, examined in some detail in Part 7, supports two kinds of conclusions. (1) Natural frequency format reduces the extent of biases only for a few heuristics. For instance, in the best case scenario for the frequency format, for base rate neglect, less than half the subjects conform to Bayes' rule. (2) For the vast majority of heuristics, results for the probability format and the frequency format are comparable. Most real-world data of interest to economists are presented in a probability format, e.g., interest rate increases announced by central banks, interest rates quoted by mortgage providers, percentage changes in unemployment, crime, and GDP changes announced by various government agencies. Thus, the debate between the proponents of the natural frequency format and the probability format might be of more academic than substantive interest.

In his seminal work on bounded rationality, Herbert Simon argued that cognitive limitations preclude individuals from engaging in neoclassical optimization exercises. As Simon (1978, p. 347) famously put it: "But there are no direct observations that individuals or firms do actually equate marginal costs and revenues." Reinhard Selten has often argued that many economic problems are hard in the sense that they are *NP-Complete*. This means that the number of steps needed to solve a problem in any algorithm of the solution, grows exponentially with the size of the problem, as in the well-known travelling salesman problem. Moderately difficult problems in this class require very high computation time. It is a leap of faith to assume that humans can routinely solve these problems or act "as if" they could.

Herbert Simon's solution was to replace *optimization behavior* with *satisficing behavior*, i.e., individuals have an *aspiration level* and adjust gradually towards it (Simon, 1978, 2000; Selten, 2001); for supportive empirical evidence see Caplin et al. (2011). A second line in the literature is associated with the work of Gerd Gigerenzer and colleagues (Gigerenzer et al., 1999; Gigerenzer and Selten, 2001). The main thrust of this approach is that heuristics are not mistakes, but a rational response to bounded rationality, and heuristics often outperform selected optimization methods. However, the superiority of fast and frugal heuristics is often demonstrated in domains where the relevant optimization benchmark is unclear. In contrast, Tversky–Kahneman were interested in domains where the relevant benchmark is completely clear. Furthermore, in the fast and frugal heuristics approach, the heuristics are typically experimenter provided; by contrast, in more complicated situations, it is not always clear what the relevant heuristics should be. Tversky–Kahneman are able, in contrast, to typically specify the environments in which their heuristics apply.

Gigerenzer's work, in particular, has led to a heated debate with the Tversky–Kahneman approach with most commentators treating the positions as adversarial. This has come to be known as the *great rationality debate* in psychology that is outlined in Part 7 (Stanovich and

West, 2000; Stanovich, 2012). My own reading is that this debate is extremely muddled; the two approaches are complementary and deal with quite different domains. Both approaches need to be encouraged by social scientists, but for entirely different reasons. Kahneman and Tversky's (1996, p. 584) frustration is well captured in this quote: "The position described by Gigerenzer is indeed easy to refute but it bears little resemblance to ours. It is useful to remember that the refutation of a caricature can be no more than a caricature of a refutation."

How would you feel about economics, and particularly about economists, if you were Daniel Kahneman? (1) Prospect theory, the most empirically satisfactory decision theory for risk, uncertainty, and ambiguity, has not yet made it into the mainstream economics curriculum, 36 years after it was published. (2) The heuristics and biases approach was initially viciously attacked, continues to be ignored by most economists despite having the greatest relevance for their work, and then pitted unfairly and unnecessarily against Gerd Gigerenzer's approach. One can sense the underlying disappointment if one reads Kahneman (2011) between the lines. However, the literature on behavioral finance, and a new generation of economists, most notably led by Matthew Rabin, have rekindled interest in this approach. Part 7 reports this work in detail.

6 Five examples of behavioral evidence

In this section, I consider five examples of behavioral evidence. This evidence often challenges some of the assumptions or implications of neoclassical models that are taken to be true on a-priori grounds. A more detailed account can be found later in the book.

6.1 *Does competitive market equilibrium survive fairness considerations?*

The development of the theory of competitive market equilibrium (CE) is one of the great intellectual achievements of modern economics. Empirical tests of a CE require a specification of the hypothetical Walrasian auctioneer that is a central part of the theory. Chamberlin (1948) ran some of the first CE experiments with Harvard students. He allowed students to interact freely in their roles as buyers and sellers in a market, and strike voluntary and mutually agreeable deals. This experimental protocol accords well with the intuition of many people about how competitive markets function. He found that prices did not converge to a CE. Modern experimentalists will disagree with many features of Chamberlin's experiments, despite their intuitive appeal. For instance, subjects were not incentivized and did not play the game anonymously without face to face interaction. I suspect strongly that Chamberlin's results will go through even if the subjects were incentivized; the main drivers of his result are the other elements of his experimental design. However, Chamberlin's results have not had a fundamental impact on modern economics.

Vernon Smith (1962) used a novel method of inducing demands and supplies among subjects in experiments. This is known as *induced value methodology*, and is now a standard technique to induce preferences in experiments. Each buyer is given a card with a valuation on it that is privately known to the buyer only. The buyers cannot pay a price above their valuation. Similarly, sellers are given a card with a cost on it, and they are not permitted to sell below it. In a *double auction setting*, in each round, buyers place 'bids' and sellers post 'asks.' Buyer-seller pairs voluntarily trade within a given period; successful matches drop out for the remainder of the period. Subjects keep playing within the period until they strike no more voluntary deals, or the

period ends. Within this setting, there is relatively quick convergence to a CE, a result that has been replicated many times (Davis and Holt, 1993; Kagel and Roth, 1995). This result also holds with zero intelligence robots (Gode and Sunder, 1993), which suggests that it is the underlying induced value methodology and the institution of double auctions that lead to a CE.

Vernon Smith's findings were readily accepted into economics, unlike Chamberlin's and Kahneman and Tversky's findings. However, Smith's results leave open the question whether we get the CE outcome in the absence of (i) the double auction setting, which is not a particularly pervasive institution, (ii) the induced value methodology, and (iii) in richer competitive market settings. Regarding (ii), Thaler (2015, p. 40) writes: "But I had some worries about this methodology [induced value methodology]. When you go to the store and decide whether to buy a jacket for \$49, no one is telling you how much you are willing to pay for it. You have to decide that for yourself, and that value might depend on all sorts of issues." Indeed, in Thaler's own experiments on exchange asymmetries, when subjects had to assign a value to an owned object and an unowned but identical object, owners assigned higher values to owned items. In this section, we focus on (iii).

Fehr et al. (1998) tested for a CE by enriching the market environment to include issues of endogenous quality. In this case, the outcomes did not converge to a CE. It is worth exploring these findings and the subsequent literature. Suppose that risk-neutral sellers have some fixed cost, f , of producing a single unit of a good and they must choose the quality of the good from some interval $[q_0, q_1]$. The variable cost of producing alternative quality levels is $c(q)$ such that $c' > 0$, so higher quality is more costly to produce. Risk-neutral buyers have unit demands, know the seller's cost, and make price offers to the sellers. Hence, this is a one-sided oral auction with many buyers. Buyers prefer higher quality, and their valuation for one unit of the good is $v(q)$, such that $v' > 0$. Sellers observe the auction price, p , and then decide on non-enforceable quality that they will supply to the buyer. There are more sellers than buyers.

Since higher quality is costly to produce, if all parties have self-regarding preferences, then we should expect sellers to act opportunistically and supply the lowest quality, q_0 . Anticipating this, buyers will offer the lowest price $p = v(q_0)$. Thus, in a competitive equilibrium, sellers make a profit equal to $v(q_0) - f - c(q_0)$, the utility of the buyer who wins the auction is $v(q_0) - p = 0$, and the utility of non-winning buyers is also zero. The parameters of the experiment are chosen so that $v(q_0) > f + c(q_0)$, where $f + c(q_0)$ is the minimum cost to the seller of supplying a unit of the good at the lowest price, hence, it is the *reservation price* in the experiments.

The outcome will be very different if market participants have other-regarding preferences. In particular, if sellers are known to exhibit reciprocity, then buyers will offer prices in excess of the seller's reservation price, and the seller will reciprocate with higher quality. It is likely that in equilibrium, buyers and sellers derive, respectively, higher utility and higher profits relative to the predictions of a CE (with self-regarding preferences). This may be termed as the reciprocity treatment (RT). The parameters in the experiments were chosen so that reciprocal quality increases were associated with Pareto improvements. In order to check if reciprocity is the driving motive, a control treatment (CT) was also implemented in which the quality was publicly known to be exogenously fixed and the seller had no control over it.

There was intra-group homogeneity among the sellers and the buyers. The matched parties did not know each other's identities; price messages were transmitted by telephone to hide the identities of players. While prices were publicly announced, the information on quality was revealed only to the two contracting parties. These features ensured that there were no reputational effects across different rounds of the game. For the RT, Figure 6 shows the response of quality (vertical axis, $[q_0, q_1] = [0, 0.7]$) to various levels of price (horizontal axis, in blocks of 10s, and ranging




Figure 6 Response of quality to price in a competitive equilibrium.

Source: Reprinted from *European Economic Review*, 42(1): 1–34, Ernst Fehr, Georg Kirchsteiger, and Arno Riedl, "Gift exchange and reciprocity in competitive experimental markets." Copyright © 1998, with permission from Elsevier.

from 30 to 110). Numbers over the bars indicate the number of observations in each category. The main finding is that higher prices offered by buyers elicit higher quality by the sellers. Furthermore, in the RT, the average price offered by the buyers was twice the level of the reservation price. However, in the CT, where the quality was fixed exogenously, the prices go down to the sellers' reservation price, hence, sellers' reciprocal behavior is the main driver of departures from the CE outcome. When sellers were offered low prices, they reacted by offering the minimum quality, which is also consistent with the reciprocity hypothesis.

There is experimental evidence of correlation between competition, market institutions, and trust/reciprocity. In their seminal cross-cultural study of the ultimatum game, Henrich et al. (2001) found that the rejection rates of lower offers are higher in societies where there is predominance of buying/selling, working for a wage, and production is carried out on a cooperative basis (as in modern firms). However, causality can run both ways. In societies with greater trust and reciprocity, markets may tend to work well because it is easier to engage in information contracts. Conversely, markets may have an impact on sociality and, therefore, affect the willingness to trust and reciprocate. Given the current evidence, we cannot establish the direction of causality. Evidence shows that in the trust game, competition improves human sociality (Huck et al., 2012); this line of research needs to be developed further.

List (2006) seeks to address the issue of sellers' quality differences between lab and field experiments, based on data gathered from sports card dealers. In the typical marketplace for sports cards, dealers display their sports cards and buyers mill around, requesting cards of a particular quality and offering a price for that quality. In the field data, quality can often be hard to verify and sellers do not know if buyers will later engage in costly quality verification from a third party. Buyers and sellers are recruited by monitors from a sports card market to participate in an hour long experiment. Buyers can offer either a low price, \$20, or a high price, \$65. In contrast to the field study, sellers in the lab know that the quality will be subsequently measured, a factor that may influence their quality supplied in the lab relative to the field.

The lab experiments were conducted in close proximity to the sports card market, hence, these experiments offer an unusually clean comparison of lab and field results. There were two kinds of subjects under study, local dealers (70% of the sample), and non-local dealers (30% of the sample).




Figure 7 Response of quality to price for local and non-local dealers in the lab and in the field.

Source: Camerer, Colin F. (2015). "The promise and success of lab-field generalizability in experimental economics: a critical reply to Levitt and List." In G. R. Fréchette and A. Schotter (eds.), *Handbook of Experimental Economic Methodology*, pp. 249–96. Adapted from List, J. A. (2006). "The behavioralist meets the market: measuring social preferences and reputation effects in actual transactions." *Journal of Political Economy* 114(1) 1–37.

List (2006, p. 5–6) used the results of this data to conclude that “even though the data collected from one-shot laboratory experiments suggest that social preferences are quite important among these agents, parallel treatments in the field suggest that such effects have minimal influence in naturally occurring transactions. In this sense, dealer behavior in the marketplace approaches what is predicted by self-interest theory.”

Following Camerer (2015), let us re-examine the data arising from a within-subjects analysis, i.e., from instances where the same sports card seller or buyer generated data in the lab and in the field; clearly, this is the most meaningful comparison between the lab and the field. This also ties in nicely with the discussion on lab-field generalizability in Section 3, above. The results are shown in Figure 7; some of this data is based on analysis that was not reported in List (2006). When data are pooled for local and non-local dealers, the quality level supplied (on a 1 to 5 scale) is lower in the field setting, relative to the lab, for each price level: 2.3 versus 3.1 when the price is \$20, and 3.1 versus 4.1 when the price is \$65. However, the response of quality to changes in price is similar in the lab and the field.

Interpreting the results of List (2006), Levitt and List (2008) argue that sports card dealers in the field behave more selfishly, and (p. 10) “increase quality little in response to a generous offer from the buyer.” In their abstract, they say: “Economic models can benefit from incorporating insights from psychology, but behavior in the lab might be a poor guide to real-world behavior.” Is this a reasonable conclusion for the more relevant case of within-subjects design? Using Figure 8, for field data, non-local dealers increase their quality in response to higher prices by just 0.13, but the corresponding figure for local dealers (who constitute 70% of the sample) is 1.40, which is significant. Levitt and List draw their strong conclusion because they only report the data for non-local dealers in their 2008 paper. Local dealers do offer higher quality in response to higher prices, hence, their behavior could be the outcome of social preferences.

Camerer (2015) conducts tests of statistical significance of the difference in quality between the lab and the field, for the within-subjects design, but the results are inconclusive. He notes: “To be crystal clear, I am not sure whether dealers behave differently in the lab and the field.” List

(2006) also gives data on the difference between actual quality supplied by sellers and the seller's claimed quality, a measure of the sellers' mendacity. Mendacious behavior is more common in the field, yet the statistical tests lack power to differentiate lab and field data.

6.2 Why do we not let bygones be bygones?

An important insight of neoclassical economics is that if self-regarding preferences are combined with classical game theory, then we should not observe costly ex-post punishments. This is because, such punishments only reduce the size of an individual's payoff, so why punish? In other words, *let bygones be bygones*.

Consider a typical public goods game experiment. A group of n risk-neutral individuals are given an initial endowment of money, $y > 0$. They simultaneously decide on their respective contributions, $g_i \geq 0$, $i = 1, \dots, n$, to a public good, giving rise to $G = \sum_{i=1}^n g_i$ units of the public good (the production technology for the public good is linear). The public good gives a return $r > 0$ per unit, to each individual. Thus, the utility of the i th individual is given by $y - g_i + rG$; typically $\frac{1}{n} < r < 1$, so that it is privately optimal to free-ride (i.e., $g_i = 0$), although the efficient solution is that everyone should contribute their entire endowment (i.e., $g_i = y$). Now introduce a costly ex-post punishment stage in which, at the end of the game, the individuals can observe the contributions of others, and engage in costly punishment of non-contributors. An individual who desires to punish another, gives up one monetary unit and the experimenter takes away three monetary units from the intended target of the punishment.

Classical game theory predicts that the punishment stage is irrelevant, because bygones are bygones. Looking ahead, rational players will realize that the punishment stage has no bite, hence, the public goods game with and without punishment is predicted to yield identical results. Fehr and Gächter (2000a) put this to the test. They ran 10 rounds of a public goods game with punishment, followed by 10 rounds without punishment (rounds 11–20), with $y = 20$ and $r = 0.4$. There were three different treatments. In the partner treatment, the same subjects were matched over all rounds. In the stranger treatment, they were randomly rematched in groups at the start of each round. In the perfect stranger treatment, they were randomly rematched in groups at the start of each round. The perfect stranger treatment is similar to the stranger treatment, but when subjects are




Figure 8 Public good contributions in the presence and absence of punishments.

Source: Fehr and Gächter (1999).

rematched at the start of each round, the experimenter ensures that they have not been matched in any of the past rounds. The results are shown in Figure 8.³²

The prediction of classical game theory, based on self-regarding preferences, is rejected by this evidence. Contributions are high in the presence of punishment for each of the three treatments (rounds 1–10), but decline rapidly once the punishment is removed, reaching near perfect free-riding levels (rounds 11–20). Furthermore, in the presence of punishments, there is a marked difference in the contributions in the three treatments. Contributions are near first-best levels in the partner treatment, but even in the perfect stranger treatment, subjects contribute about half their endowments. Subjects seem to anticipate that they will be punished if they contribute less, particularly when they are repeatedly matched with the same partners. This induces high levels of cooperation.

Punishments, in contrast to the predictions of classical game theory, perform a prosocial role. When the punishment option is removed, subjects begin with high levels of cooperation. However, once contributors observe low contributions in their groups by opportunistic players, they respond in the next round by contributing less. This leads to successive rounds of negative reciprocity and reduced contributions. The efficiency enhancing role of punishments has also been documented in other games, such as in the gift exchange game with punishment (Fehr et al., 1997); we consider this and other evidence in Part 2.

Consider a few other related points about this analysis.

1. The fact that in the absence of punishment, we observe widespread free-riding does not save the theory with self-regarding preferences, because it cannot explain the pattern of decay in contributions. An explanation of this pattern requires a fraction of the players to exhibit reciprocity (Fischbacher et al., 2001; Fehr and Fischbacher, 2005; Fischbacher and Gächter, 2010).
2. When the experiment is repeated outside Western countries, one also observes antisocial punishment, i.e., punishment of contributors by non-contributors who react to past punishments. Antisocial punishments are more likely to be found in societies where the indices for the *norms of civic cooperation* and *rule of law* are low (Herrmann et al., 2008). Experiments possibly overstate the role of punishments because in actual practice, punishments are institutionally meted. Democratic societies do not sanction antisocial punishments. The desire for prosocial punishments, and the efficiency enhancing role of punishments (see below) in most Western societies, possibly underpins public sanctions for antisocial behavior.

In an important advance, Fehr and Williams (2015) allow subjects in experiments to voluntarily choose one of the following four institutions. (1) No punishments. (2) Fehr–Gächter type peer punishment. (3) As in (2), but with the additional opportunity for each individual to announce how much everybody should contribute. The average contributions are revealed at the end of each round. This captures the formation of social norms of contributions. (4) Centralized punishment by a democratically determined leader, plus social norms of contributions, as in (3). The main result is that virtually all subjects choose institutions (3) and (4). Furthermore, institution (4) wipes out antisocial punishment. We also review other results in Part 2, which demonstrate that antisocial punishments are eliminated if, ex-ante, individuals choose the punishment institution by a referendum.

³² One of the most versatile and sagacious social scientists of our times, Herbert Gintis, once mentioned to me that Figure 8 was one of the most important graphs that he had seen in his life.

3. The efficiency enhancing role of costly punishments has also been debated using evolutionary arguments. In experiments, the net effect of costly punishments may be to reduce overall payoffs (Page et al., 2005). However, this is not the case, if sufficiently large numbers of rounds of punishments are allowed. Gächter et al. (2008) have shown that if the credibility of the punishment threat is established, and if the punishment opportunity arises over sufficiently many rounds (10 and 50 round cooperation games are considered), then punishments are efficiency enhancing; in fact, only 5–8 rounds of punishments may suffice to increase overall payoffs. Social punishments and reciprocity are also likely to have been essential to the formation of social norms of cooperative behavior (Gintis, 2009).
4. Subjects do not let bygones be bygones, which is a puzzle under self-regarding preferences. Punishments in the public goods experiments can be explained by Fehr–Schmidt preferences in which people have a concern for inequity aversion. Neural evidence is consistent with the view that individuals derive hedonic satisfaction from the act of punishing others who they believe have wronged them. A region of the brain, the striatum, that is associated with the brain's reward circuits, shows greater activity when social punishments are carried out (De Quervain et al., 2004). This also illustrates the powerful role of neural evidence in informing us about human motivations.

6.3 Are financial markets efficient?

Let us call the neoclassically rational agents in financial market *arbitrageurs*. Most neoclassical economic analyses allow for the possibility that some economic agents could be irrational, also known as *noise traders*, relative to arbitrageurs. Yet, it is argued, we can ignore them due to two a-priori arguments.

1. The individual errors of *noise traders* will cancel in the aggregate.
2. Noise traders, if any, will lose money due to the arbitrage activity of arbitrageurs, hence, they will eventually exit the market.

Anyone who has done an introductory course in economics would have come across these plausible sounding arguments. Most, including those teaching these arguments, are likely to have been convinced. Empirical evidence shows how flawed this a-priori reasoning can be. Formulating plausible a-priori hypotheses is essential in all science, as is testing them stringently, but the stringent testing bit appears to have received less attention, possibly motivated by bogus “as if” arguments.

The first argument has been debunked by the heuristics and biases approach of Kahneman and Tversky (Section 5.5, above). Individual errors of noise traders are often systematic and not random, so they can accumulate in one direction. For instance, individuals subject to the hindsight bias consistently underestimate volatility of stock prices, and fail to update their incorrect models fully (Biais and Weber, 2009). In the formation of *bubbles*, for instance, house price bubbles, individual beliefs can systematically depart from fundamental values for a sufficient length of time. In financial markets, the *disposition effect* caused by loss aversion may prevent people from selling stock below its purchase price precisely at a time when such arbitrage activity is required to equilibrate the market. Insofar as loss aversion is hardwired in individuals, they systematically engage in the disposition effect.

The second argument also turns out to be problematic. Consider the example of house price bubbles, and let prices exceed fundamental values. Yet, suppose that investor sentiment, as

captured in the beliefs of noise traders, is that prices will continue to increase. If there are enough noise traders in the market, then arbitrageurs might be forced to bet their money in the direction of a price increase, which exacerbates the departure of price from fundamental values. The seminal work of De Long et al. (1990) shows that noise traders may make more profits than arbitrageurs, particularly if they are bullish, and they may continue to make money for a time, while arbitrageurs might not be able to hold on to loss making positions for long enough. But, there are many other reasons why financial markets might not be efficient that we discuss next.

Traditional finance has been built on the foundations of the *efficient markets hypothesis* (EMH). It requires that at any instant in time, the price of a security must equal its fundamental value, which we may take to be the present discounted value of all future income flows from the security. Several theoretical arguments can be advanced for why the EMH need not hold. The demonstration of De Long et al. (1990) for market inefficiency based on noise trader risk provided an early theoretical challenge to the EMH. The full range of substitute securities that are required to guarantee EMH might not exist. Transaction costs of arbitrage (cost of finding mispricing information and brokerage fees for taking a short position) might also impede market efficiency.

The difficulty of beating the stock market is often cited as proof of its efficiency. EMH implies that there are no arbitrage opportunities; however, the absence of arbitrage opportunities does not logically imply that EMH holds. This is simply a confusion between necessary and sufficient conditions (Barberis and Thaler, 2003). EMH relies on instantaneous flows of information. Yet, evidence shows that some kinds of information may flow only gradually through the market, leading to a relationship between momentum and trading volumes of stocks, and aggressive trading, even when price equals fundamental values (Hong and Stein, 1999, 2007).

Professional arbitrageurs, such as mutual funds managers, invest on behalf of small and dispersed investors who have little financial expertise. However, the greater financial nous of professional arbitrageurs cannot be relied on to induce market efficiency (Shleifer and Vishny, 1997). Small investors only observe indirect signals of the ability of professional arbitrageurs, such as the end of year financial performance of professionals, and their performance relative to other professionals. This forces professional arbitrageurs to (i) liquidate their positions in loss making assets too quickly before the year is out, which makes it difficult to counter the activity of noise traders, and (ii) to invest in relatively similar portfolios to avoid looking worse than the others, which is a form of *herding behavior* and can lead errors to accumulate in one direction. These distortions may even reduce the performance of professional arbitrageurs below that of passive investment strategies (Ippolito, 1989; Lakonishok et al., 1992).

A great deal of evidence from financial markets is inconsistent with the EMH. While substitute securities are essential for EMH, their prices can differ systematically, which is inconsistent with EMH. Particularly clean evidence of this effect comes from *Siamese-twin companies* that have merged in the past on an $x : y$ equity basis, yet may trade separately in different markets. If EMH holds, then their traded prices should be in the $x : y$ ratio, but they differ systematically from this ratio (Froot and Dabora, 1999). Security prices often react to no-news and many stock market crashes occur suddenly without any apparent changes in fundamental values (Cutler et al., 1991). The price of freshly inducted stocks into the S&P 500 index immediately jumps upwards, but no new information about the company has been revealed (Shleifer, 1986). The volatility in stock prices is excessive, relative to the predictions of EMH (Shiller, 1981).

Stale information can predict future prices. For instance, stocks that have done well over a long horizon give low subsequent returns and past losers give higher subsequent returns (De Bondt and Thaler, 1985). Over shorter horizons, a *momentum effect* has been observed; recent stock market performance of stocks can carry over to the future (Jegadeesh and Titman,

1993). All public information should immediately have been incorporated into stock prices, yet there is evidence of a *post-earnings-news drift* (Michaely et al., 1995). The decision to buy stocks depends on whether news about the performance of stocks is framed as *high, long-term, returns*, or a *more finely presented series of fluctuating short-term stock returns* (Benartzi and Thaler, 1995). The evidence from closed-end funds shows that there is a systematic difference between “net asset value” and the market price of shares, which is not consistent with the EMH (Lee et al., 1991).

The behavior of investors who follow prospect theory creates another reason for the inconsistency of evidence with the EMH, which assumes that investors follow expected utility theory. Recall the probability weighting function in Figure 2. This is central to the emerging topic of *psychology of tail events* (Barberis, 2013). Since individuals overweight small probabilities, the market will price skewness in the returns to assets (fat tails on the right of the returns distribution); although, the probability of these returns is low, the probability is overweighted. If there is a small probability of returns on an asset collapsing, then this can potentially explain the equity premium puzzle (De Giorgi and Legg, 2012); since the low probability downside risk is overweighted, a premium must be paid to induce investors to hold it. Other aspects of prospect theory also mitigate against the EMH. Loss aversion contributes to two phenomena already discussed above, the disposition effect (Odean, 1998), and the equity premium puzzle (Benartzi and Thaler, 1995).

6.4 Is expert behavior consistent with neoclassical economics?

A common response to the anomalies of the neoclassical model is that its predictions apply only to experienced subjects who operate in domains familiar to them. A particularly good test of this response is to examine the behavior of experts in their own field. The available empirical evidence shows that experts are not immune to biases and exhibit some biases even more than lay people (Gilovich et al., 2002; Kahneman, 2003; Kahneman, 2011; and Tetlock, 2006). Tetlock (1999, p. 351) evaluates the evidence on the predictions of experts on political events as follows: “Across all seven domains, experts were only slightly more accurate than one would expect from chance . . . Most respondents thought they knew more than they did. Across all predictions that were elicited, experts who assigned confidence estimates of 80% or higher were correct only 45% of the time . . . Expertise thus may not translate into predictive accuracy, but it does translate into the ability to generate explanations for predictions that experts themselves find so compelling that the result is massive over-confidence.”

Experts exhibit many of the biases that we outlined in Section 5.5. They exhibit greater *over-confidence* relative to lay people (Heath and Tversky, 1991; Frascara, 1999; Glaser et al., 2007a,b), and more experienced experts are even more overconfident (Kirchler and Maciejovsky, 2002). Managerial overconfidence may explain excessive merger activity, despite the fact that relatively few mergers are successful (Malmendier and Tate, 2005, 2008). Another manifestation of overconfidence is the predicted distributions of stock market returns by senior finance professionals; the predicted intervals are too narrow and realized returns are within their 80% confidence intervals, only 36% of the time (Ben-David et al., 2013). Mathematical psychologists exhibit the *law of small numbers* (Tversky and Kahneman, 1973). Experts’ perceived riskiness of various hazardous substances is determined by their *affective response* (Slovic et al., 1999). The affective response of experts also determines the likelihood of reoffending that they assign to discharged mental patients (Slovic et al., 2000). Clinical psychologists *underweight the base rate* relative to a Bayesian decision maker (Meehl and Rosen, 1955). Finance professionals exhibit the *hindsight bias* (Guilbault et al., 2004; Biais and Weber, 2009).

Estate agents *anchor* on the list price when evaluating the value of properties (Northcraft and Neale, 1987). The anchoring heuristic has also been found to influence expert legal judgment (Chapman and Bornstein, 1996; Englich and Mussweiler, 2001; Englich et al., 2006). Judges exhibit the false consensus effect (Solan et al. 2008) and there have been various legal and institutional developments to minimize the effect of such biases (Rachlinski, 1998). In imputing the risk preferences of others, finance professionals are disproportionately influenced by their own risk preferences—this is the *false consensus effect* (Roth and Voskort, 2014). When presented with identical options, the responses of physicians are influenced by *framing effects*; their responses depend on whether outcomes are framed as gains or losses relative to a reference point (McNeil et al., 1982; Kahneman and Tversky, 1984). Physicians display limited attention and cognitive overload with the addition of extra choices, even when the total number of choices is relatively small (Redelmeier and Shafir, 1995).

6.5 Do people play a mixed strategy Nash equilibrium?

The ability to play a *mixed strategy Nash equilibrium* (MSE) is critical to the theoretical foundations of neoclassical economics; many critical results, such as existence results in game theory, depend on it (Nash, 1950). We have already noted that the epistemic conditions for playing a Nash equilibrium are extremely stringent, although these results are inadequately highlighted in most game theory courses. Harsanyi (1973) interprets an MSE as the limiting behavior in a set of games that have pure strategy equilibria, but the payoffs are perturbed; this is the *purification agenda*. Yet, these results do not hold for an important range of games that economists are interested in, such as repeated games and principal–agent games, which limits the practical usefulness of this argument (Gintis, 2009; Chapter 6).

An extra layer of cognitive sophistication is required for an MSE relative to a pure strategy Nash equilibrium. Players must randomize among their pure strategies in such a manner that leaves their opponents exactly indifferent among the pure strategies in the support of their mixed strategies. It takes most undergraduates a while to internalize this breathtaking degree of sophistication, but most buy into the idea within a few weeks. However, the evidence on MSE is rarely presented in game theory courses. The history of science is replete with ideas that were ex-ante counterintuitive, yet surprisingly turned out to be true when confronted with the empirical evidence. Based on the available evidence, MSE cannot be classified into this category. The empirical argument is twofold: (i) Many people play a mixed strategy, but not an MSE. (ii) The best evidence supporting the MSE comes from professional sports, a domain that arguably is not the most suitable for reasons given below.

Consider the rather complicated choice facing a football penalty kicker. It requires choosing the height, direction, and swerve of the ball, as well as body feints and gestures. For simplicity, let us assume that the choice is either to kick to the left, or to the right of the goalkeeper. Clearly, over successive penalty shots, kickers would need to randomize somehow, otherwise, given the widespread use of video analysis of past games in modern sports, they will be ineffective. Similar comments apply to tennis serves, which are complicated by the spin, the angle, and the slice on the ball, but let us assume that the choice is to serve either to the left or to the right. In these domains, support is found for the MSE in professional soccer games (Palacios-Huerta, 2003) and in professional tennis matches (Walker and Wooders, 2001). However, in each case, there is serial dependence in the actions (penalty kicks and tennis serves) that is not detected by the opponents, thus, they are not able to exploit it. Serial dependence is not surprising, given the difficulty that humans have in generating random sequences of numbers

(see Section 5.5, above). However, leaving aside the modelling of other complexities in penalty kicks and tennis serves that are mentioned above, serial dependence is strictly not consistent with an MSE.

The results on MSE in a sports context are clearly important because they give us a concrete domain in which MSE is reasonably successful. However, the main skills employed by professional sportsmen in playing mixed strategies are motor skills and hand–eye–foot coordination. The development of these skills has an evolutionary basis in hunting, and in intergroup conflict, where they were probably critical for physical survival. Indeed, with practice, humans can achieve a high level of these skills, as evidenced by professional sportsmen. There is less basis to believe that motor skills come into play in important economic decisions such as choosing between alternative pension plans, portfolios of assets, savings decisions, choice of a University course or degree, tax evasion decisions, number of hours to work, or in dealing with self-control issues. Data from professional sports may, thus, have low external validity in establishing conclusive evidence for an MSE.

The ability to play a MSE in a sports context does not translate into an ability to play MSE in other contexts, i.e., the ability to play mixed strategies is not portable. The initial evidence was supportive of the portability of MSE (Palacios-Huerta and Volij, 2008), however, this was overturned by the results in Wooders (2010). In a carefully conducted study, Levitt et al. (2010) appear to have largely settled the debate by providing a clear negative answer to the question of the portability of MSE.

The serial dependence in football kicks and tennis serves suggests that subjects may be using some mixing heuristics that were not explicitly tested in the empirical model. This is an important consideration. In an influential experiment to test an MSE, O'Neill (1987) employed a two-player zero-sum game in which each player had four strategies (Ace, 2, 3, and Joker) and there was a unique MSE (details in Part 4). The experimental results, on the face of it, provided stunning confirmation of an MSE. The marginal frequencies of play by the row and column players appeared close to the MSE, as did the win percentages of each player. However, in an exemplary piece of empirical work, Brown and Rosenthal (1990) showed that O'Neill's findings constituted an unequivocal rejection of MSE. Indeed, several plausible heuristic-based models could account for the observed empirical results. Despite the proximity of the empirical frequencies of play with the MSE predictions, a chi-square test rejected MSE at any level of significance. Furthermore, they showed that the O'Neill design had insufficient power to distinguish between MSE and the alternative explanations.

There is some evidence that an MSE is more likely to be supported by the evidence when it predicts uniform randomization over the available pure strategies (Slonim et al., 2003). However, moving away from this special case and from sports contexts, the evidence often does not support an MSE. Using an extensive empirical analysis, Mookherjee and Sopher (1994, 1997) conclude in their 1994 paper (p. 109) that: "The rejection of the minimax theory at the aggregate level, and not just at the level of individual decision making, implies that it cannot even be treated as an acceptable as if theory."

Suppose that a game has a unique MSE and it is played repeatedly over a finite number of rounds. One expects players to draw repeatedly in each round from a multinomial distribution that reflects the MSE of the game. However, a common finding is that there is serial correlation in such play. One defence of the MSE is that although players can play an MSE in a static game, they are unable to draw randomly from the relevant multinomial distribution in repeated play. In order to test for this possibility, subjects were given a randomizing device, yet they exhibited serial correlation in play (Shachat, 2002). A reasonable interpretation is that subjects do play

mixed strategies, but not in the proportions predicted by MSE. We consider several empirical studies of MSE in Part 4.

In the *equilibrium in beliefs* interpretation of MSE, each player plays a pure strategy (but never a mixed strategy) and the proportion of players who play each pure strategy is identical to the predicted MSE. Thus, each player acts as if he faces an opponent who plays a mixed strategy (Binmore, 1991). However, the evidence that people do play mixed strategies (although typically not the MSE) also rejects the equilibrium in beliefs interpretation of MSE.

Appendix A: The random lottery incentive mechanism

Consider the *random lottery incentive mechanism* (RLI) that is widely employed to pay subjects in experiments. Suppose that expected utility maximizing subjects in an experiment perform n tasks from a set of tasks $T = \{t_1, \dots, t_n\}$. Let task t_i involve choosing from a set of lotteries, $\mathcal{L}_i$ and define $\mathcal{L} = \cup_i \mathcal{L}_i$ as the set of all lotteries considered in the n tasks. Consider a subject in the experiment whose preferences over pairs of elements in $\mathcal{L}$ are represented by $\succeq$. At the beginning of the experiment, subjects are told that one of the tasks will be picked at random once all tasks are completed, and their choice in that task will be paid off for real. Let us abstract from issues of a possible show-up fee for the experiment that we have already commented on above.

Suppose that in any task, t_i , the most preferred lottery of the subject is L_i^* . Conditional on the incentive structure described above, should the subject choose L_i^* ? An elicitation method is incentive compatible if the answer is yes. The probability that task i is chosen to be rewarded at the end of the experiment is $\frac{1}{n}$. With the complementary probability $1 - \frac{1}{n}$, any of the other $n - 1$ tasks might be chosen to be rewarded. Let the choice made by the subject on any task t_j , $j \neq i$ be L_j . Let us denote by $(L_1, \frac{1}{n}; \dots; L_i^*, \frac{1}{n}; \dots; L_n, \frac{1}{n})$ the compound lottery in which lottery L_j is played with probability p_j . From the *independence axiom* of expected utility theory that should be familiar to most undergraduates in economics (see Part I for details), we know that since $L_i^* \succeq \widehat{L}_i$ for all $\widehat{L}_i \in \mathcal{L}_i$, it follows that

$$\left(L_1, \frac{1}{n}; \dots; L_i^*, \frac{1}{n}; \dots; L_n, \frac{1}{n} \right) \succeq \left(L_1, \frac{1}{n}; \dots; \widehat{L}_i, \frac{1}{n}; \dots; L_n, \frac{1}{n} \right) \text{ for all } \widehat{L}_i \in \mathcal{L}_i. \quad (9)$$

In words, the independence axiom says that if $L_i^* \succeq \widehat{L}_i$ then the decision maker prefers any mixture of lotteries that gives L_i^* over an identical mixture that contains $\widehat{L}_i$ in its place. Thus, the subject will choose L_i^* in task t_i . We can show this to be true for all tasks $t_j \in T$. Thus, RLI is incentive compatible.

Appendix B: In lieu of a problem set

This section poses 50 problems, written in the same order as the topics in the book and tailored to the material presented in the introduction. If you are already familiar with a particular problem, just move on to the next one. The problems range from straightforward applications of the material in the introduction to slightly more challenging ones. I deliberately avoid giving too much structure to the problems, so that you can try to solve them by writing your own models in a free-spirited manner. If you wish to use the neoclassical model to solve a problem, then think carefully about the auxiliary assumptions that you might invoke. In particular, do ask yourself

if the auxiliary assumptions that you use are ad hoc or not. You will encounter the solutions to these problems, and many others, as you progress through the book; this will also give you an opportunity to check your initial responses.

1. Many taxi drivers quit too early on rainy days in New York when the effective wage rate is actually very high. In other words, why is it so hard to find a taxi on a rainy day in New York?
2. Why do owners of objects, humans or chimps, value them more than non-owners (under neoclassical economics, everyone should value objects at their opportunity cost)?
3. If people play Russian roulette, why are they likely to pay more to reduce the number of bullets from 1 to 0, as compared to from 4 to 3 (check that under expected utility these two choices should be equally attractive)?
4. From 1926 to about the mid 1980s, the annual real return on stocks has been about 7% with a standard deviation of 20%, while the annual real return on treasury bills has been less than 1%. In neoclassical economics, a coefficient of relative risk aversion of about 30 can explain these findings, but the actual coefficient is around 1. This is the equity-premium puzzle. How can you explain the equity-premium puzzle?
5. A decision maker has initial wealth, w . Suppose that at all levels of wealth, he prefers to keep his wealth rather than play the lottery L_1 : win \$11 or lose \$10 with equal chance, for any w . Then, under expected utility theory, the decision maker will prefer the lottery L_1 to the lottery L_2 : lose \$100 and win an infinite amount with equal probabilities (*Rabin's paradox*). Thus risk aversion over small stakes, under expected utility theory, implies implausible risk aversion over large stakes. Does expected utility correctly encapsulate the risk attitudes of an individual? Can you think of modifications to expected utility that will explain Rabin's paradox?
6. For an amateur tax evader, the actual probability of audit is 1–3% and the penalty for being caught is (i) return of evaded taxes, plus (ii) fine at the rate of 1–2 times evaded tax. A quick back of the envelope calculation will show that this implies a return on tax evasion of about 96–98%. If you are an expected utility maximizer and have a coefficient of relative risk aversion of about 70 or above, you will pay your taxes; but, empirically, the coefficient is about 1. Since there are very few assets with this return, why do people pay any taxes?
7. Why do most people not buy insurance against very low probability events such as natural hazards, even when insurance is better than actuarially fair and the losses of insurance firms are underwritten by the government?
8. In everyday conversations about risky decisions, people speak of *optimism*, *pessimism*, *disappointment*, and *regret* (think only of pure risk in a simple game against nature). Can neoclassical decision theory under risk account for these emotions?
9. In recessions, why do firms typically prefer to lay-off workers rather than cut wages?
10. Why do most people find a cut in the nominal wage of 5% under zero inflation to be more unfair relative to a nominal wage increase of 2% under 7% inflation? What are the implications for macroeconomics?
11. In experiments on redistributive taxes, why do people often choose a smaller, more equally distributed, cake as compared to one in which they get a larger share of a very unequally distributed cake?
12. Should we generally expect any difference in outcomes relative to the neoclassical case if a minority of players have other-regarding preferences (i.e., also care about payoffs of others in addition to their own)? You may think of optimal contracts (static and finitely repeated)

between a principal and two agents in a production task where both agents are essential. One of the agents has self-regarding preferences, but the other has other-regarding preferences; the principal does not know who's who.

13. Why are individuals often willing to punish third parties for observed norm violations between other players even when their own material payoffs have not been affected?
14. Why do workers often respond to higher wage offers of firms by working harder, even in "static games," where they could take the money and run? Why, in these static problems, may firms choose to offer high wages that exceed the opportunity cost of hiring workers? How might you distinguish the predictions of this model from the model of efficiency wages?
15. Why may firms sometimes choose to offer non-enforceable bonus contracts to workers in preference to enforceable incentive contracts, even in static problems? And why might the choice of effort by workers under bonus contracts be relatively higher as compared to that under incentive contracts?
16. Why do moral suasion, trust, and giving workers a goal/sense of purpose or a particular company identity, often outperform monetary incentives?
17. Why do people not lie maximally even when their behavior is guaranteed to remain completely anonymous? And why do we teach children not to lie, act morally rather than opportunistically, and help those who are less fortunate, rather than the neoclassical prescriptions about human behavior (maximize your payoffs/utility subject to technological constraints, but ignore any moral or ethical considerations relative to payoff maximization)?
18. Consider Akerlof and Kranton's (2005, p. 9) description of the following initiation process at the US West Point military academy. "On plebes' first day . . . they strip down to their underwear. Their hair is cut off. They are put in uniform. They then must address an older cadet, with the proper salute . . . must stand and salute and repeat, and stand and salute and repeat, until they get it exactly right, all the while being reprimanded for every tiny mistake." How would an economist brought up on the theory of incentives and organizations make sense of this initiation ceremony?
19. Why does the law make a distinction between murder and manslaughter, assigning much lower punishments for manslaughter for the same harm to the victim?
20. Most people would prefer one apple today to two apples tomorrow, but they prefer two apples in 51 days to one apple in 50 days. How can you explain this preference reversal?
21. Discounting over an entire time interval $[t, \bar{t}]$ in one go, turns out to be larger, relative to discounting over n successive sub intervals $[t, t_1], [t_1, t_2], \dots, [t_{n-1}, \bar{t}]$; this is known as subadditive discounting. Under exponential discounting, the two answers should be identical. What modifications do you think are needed to the exponential discounting model to explain this empirical finding?
22. Why does the data (e.g., for the US) show a sharp drop in consumption at retirement?
23. Why do many people simultaneously hold illiquid assets and credit card debt?
24. Why do many people procrastinate so much?
25. Why do many people pay more for an annual gym membership when they could save money on a pay-as-you-go basis?
26. In normal form games in which a Nash equilibrium can be found with more than 2–3 steps of iterated elimination of dominated strategies, the experimental evidence often shows that the outcome is not a Nash equilibrium. Does this evidence cause you to have any reservations about equilibrium concepts in classical game theory or a desire to modify them? If so, how?

27. Why do we observe far more cooperation in a one-shot prisoner's dilemma games (about 60% of the time) relative to the prediction of classical game theory that predicts no cooperation under the assumption that people have self-regarding preferences? Bear in mind that the prisoner's dilemma game is possibly the most widely used game in the social sciences as a metaphor for human cooperation (or the lack of it), so this is not an unimportant result that can be ignored.
28. Empirical evidence shows that in centipede games, the backward induction outcome (play down at the first node) is played less than 10% of the time. In 6 node centipede games, in a majority of the cases, players move across to at least the 4th node. How can you explain this finding? Does your explanation have testable implications?
29. If you are unconvinced by the experimental method, can you come up with a few "stringent" non-experimental tests of classical game theory using real-world data?
30. Why do bargaining negotiations often stall with adverse consequences for both parties (union strikes, wars, family gridlocks over issues) even when issues of asymmetric information are not salient? This is particularly the case in conditions where the classical alternating offers bargaining game predicts an immediate bargaining solution without delay.
31. You live in Italy and as most folks who live there, you never tip cab drivers. However, one week you go abroad, and take a taxi to a friend's house who lives in the countryside in a country where there is a norm for tipping taxi drivers. A 70-year-old meek looking and frail taxi driver delivers you safely to your destination. Would you honor the norm of tipping him or just walk away? Suppose you answered that you would pay the tip. Is your behavior consistent with classical game theory? If not, then which feature of classical game theory could be altered to explain your tipping behavior?
32. Why are winners of common value auctions often "cursed" in the sense that they make far less money than they anticipated?
33. Eyetracking data from a three-round, two-player bargaining game whose structure is hidden from view, but searchable by using mouse clicks, reveals the following. Most subjects search for payoffs and the size of the cake to be divided in each round, forward from the first round rather than backward from the third round. Subjects trained in backward induction do search backwards more often. Does this in any way make you uneasy about equilibrium concepts in classical game theory, or would you simply discount this evidence?
34. Why do interrogators often (successfully) conduct around-the-clock interrogation of suspects?
35. Why might many people end up marrying or proposing marriage/seeking divorce without sufficient deliberation, or buy consumer durables in haste? Can you think of any legal interventions that take such human behavior into account? Would such legal interventions be necessary for the typical individual in the neoclassical framework?
36. Why are most people happy to buy consumer durables on installments, yet prefer to prepay for a holiday?
37. Why do smokers and alcoholics often pay money at rehab clinics to get rid of their addictions? Recall that the typical model of addiction in neoclassical economics assumes that people choose to get rationally addicted, taking account of the relevant costs and benefits now, and in the future.

38. Why do many cigarette smokers report an increase in happiness following an increase in excise duty on cigarettes (based on US and Canadian data)?
39. A town is served by two hospitals. In the larger hospital, about 45 babies are born each day, and in the smaller hospital, about 15 babies are born each day. As you know, about 50% of all babies are boys. However, the exact percentage varies from day to day. For a period of 1 year, each hospital recorded the days on which more than 60% of the babies born were boys. Which hospital do you think recorded more such days? 53 students in the sample said that both hospitals are equally likely to have recorded such days and 21 students each chose the larger and the smaller hospital, respectively. Are the students behaving like the agents in neoclassical economics (sometimes known as *econs*)? If not, what does their behavior reveal?
40. Why do sales at lotto stores that have sold a winning ticket soar in the immediate weeks following the lotto win (this positive effect on sales persists for up to 40 weeks following the lotto win)? How can you test the hypothesis that “the winning store just produced more interest among the local population to buy more lotto tickets, so that this finding is perfectly consistent with neoclassical economics”?
41. Why do so many mergers fail, yet we often observe waves of mergers from time to time?
42. Why do people find it more difficult to make a choice when the set of choices expands (people choose easily among 3 types of jams, but often struggle to choose among 27 types of jams)?
43. There is much cross-country variation in organ donation rates in European countries (98% in Austria, but 12% in Germany; 99.9% in France, but 17.7% in UK; 85.9% in Sweden, but 4.25% in Denmark). It turns out that in countries with high organ donation rates, people are automatically enrolled in the organ donation program, but can opt out if they wish. The situation is exactly the reverse in countries with low organ donation rates, where people can opt in if they wish. Can this empirical fact be explained under neoclassical economics? If you answer yes, then be careful in stating your auxiliary assumptions and think of the evidence for these assumptions.
44. There are only two cab companies in the city, Green and Blue; 85% of the cabs are Green. There was an accident last night. A witness comes forward to testify that the cab involved in the accident was Blue. In similar conditions, the reliability of the witness is 80%, i.e., the probability that he gets it wrong is 20%. What is the probability that the actual cab involved in the accident was Blue? The median and modal response was 80%. Are the students behaving like the textbook economic agents in neoclassical economics? If not, what does their behavior reveal?
45. Why do buyers of a new car often find that the particular model they drive suddenly appears more common on the roads?
46. Why do marketing people play on alternative ways of framing information that has identical information-content? Why for instance, may swimsuit models be placed next to sports cars in advertisements, when the main role of advertisement in the neoclassical framework is to convey information to potential buyers?
47. Why do sales drop if publicly known sales taxes are displayed on price stickers rather than being added at the check-out counter?
48. Why might the financial market price skewness in asset returns?
49. Why are so many financial crises accompanied by no-news (i.e., no information related to fundamental values)?

50. Consider the probability of success that entrepreneurs assign to their startups. In one empirical study, only 5% of startup entrepreneurs believe that their odds are any worse than comparable enterprises and a third believe that their success is assured. Based on French data, 56% expect 'development' and only 6% of startup entrepreneurs expect "difficulty"; three years on, the respective figures are 38% and 17%. Empirically, only half of all startups survive beyond three years and the high failure rate among startups is widely reported in the popular press. How can we square these figures with the supposed rationality of participants in corporate finance? How should we react to this sort of evidence?

REFERENCES

- Abdellaoui, M. A. (2000). Parameter-free elicitation of utility and probability weighting functions. *Management Science* 46(11): 1497–512.
- Abeler, J. and Nosenzo, D. (2015). Self-selection into laboratory experiments: pro-social motives versus monetary incentives. *Experimental Economics* 18(2): 195–214.
- Ackert, L. F., Martinez-Vazquez, J., and Rider, M. (2007). Social preferences and tax policy design: some experimental evidence. *Economic Inquiry* 45(3): 487–501.
- Ainslie, G. W. (1992). *Picoeconomics*. Cambridge: Cambridge University Press.
- Akerlof, G. A., and Kranton, R. E. (2005). Identity and the economics of organizations. *Journal of Economic Perspectives* 19(1): 9–32.
- Allais, M. (1953). La psychologie de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école Américaine. *Econometrica* 21: 503–46.
- al-Nowaihi, A. and Dhami, S. (2006a). A simple derivation of Prelec's probability weighting function. *Journal of Mathematical Psychology* 50(6): 521–24.
- al-Nowaihi, A. and Dhami, S. (2006b). A note on the Loewenstein-Prelec theory of intertemporal choice. *Mathematical Social Sciences* 52(1): 99–108.
- al-Nowaihi, A. and Dhami, S. (2015). Evidential equilibria: heuristics and biases in static games of complete information. Forthcoming in *Games*.
- al-Nowaihi, A., Bradley, I., and Dhami, S. (2008). A note on the utility function under prospect theory. *Economics Letters* 99(2): 337–9.
- Andersen, S., Ertaç, S., Gneezy, U., Hoffman, M., and List, J. A. (2011). Stakes matter in ultimatum games. *American Economic Review* 101(7): 3427–39.
- Anderson, J., Burks, S. V., Carpenter, J. et al. (2013). Self-selection and variations in the laboratory measurement of other-regarding preferences across subject pools: evidence from one college student and two adult samples. *Experimental Economics* 16: 170–89.
- Andreoni, J. and Bernheim, B. D. (2009). Social image and the 50–50 norm: a theoretical and experimental analysis of audience effects. *Econometrica* 77(5): 1607–36.
- Ariely, D., Bracha, A., and Meier, S. (2009). Doing good or going well? Image motivation and monetary incentives in behaving prosocially. *American Economic Review* 99(1): 544–55.
- Ariely, D., Loewenstein, G., and Prelec, D. (2006). Tom Sawyer and the construction of value. *Journal of Economic Behavior and Organization* 60: 1–10.
- Arkes, H. R., Dawes, R. M., and Christensen, C. (1986). Factors influencing the use of a decision rule in a probabilistic task. *Organizational Behavior and Human Decision Processes* 37: 93–110.
- Ashraf, N., Camerer, C. F., and Loewenstein, G. (2005). Adam Smith, behavioral economist. *Journal of Economic Perspectives* 19(3): 131–45.
- Ayton, P. and Fischer, I. (2004). The hot hand fallacy and the gambler's fallacy: two faces of subjective randomness? *Memory and Cognition* 32: 1369–78.
- Babcock, L. and Loewenstein, G. (1997). Explaining bargaining impasse: the role of self-serving biases. *Journal of Economic Perspectives* 11(1): 109–26.
- Babcock, L. and Pogarsky, G. (1999). Damage caps and settlement: a behavioral approach. *Journal of Legal Studies* 28: 341–70.
- Barberis, N. (2013). The psychology of tail events: progress and challenges. *American Economic Review* 103(3): 611–16.
- Barberis, N. and Thaler, R. H. (2003). A survey of behavioral finance. In G. M. Constantinides, M. Harris, and R. M. Stulz (eds.), *Handbook of the Economics of Finance, Volume I*. 1st edition Amsterdam: Elsevier, pp. 1053–128.
- Bardsley, N., Cubitt, R., Loomes, G., Moffatt, P., Starmer, C., and Sugden, R. (2010). *Experimental Economics: Rethinking the Rules*. Princeton, NJ: Princeton University Press.
- Bar-Hillel, M. (1980). The base rate fallacy in probability judgements. *Acta Psychologica* 44: 211–33.
- Bar-Hillel, M. and Wagenaar, W. A. (1991). The perception of randomness. *Advances in Applied Mathematics* 12: 428–54.

- Barmettler, F., Fehr, E., and Zehnder, C. (2012). Big experimenter is watching you! Anonymity and prosocial behavior in the laboratory. *Games and Economic Behavior* 75(1): 17–34.
- Battigalli, P. and Dufwenberg, M. (2009). Dynamic psychological games. *Journal of Economic Theory* 144(1): 1–35.
- Becker, G. M., Degroot, M. H., and Marschak, J. (1964). Measuring utility by a single-response sequential method. *Systems Research* 9: 226–32.
- Bénabou, R. and Tirole, J. (2003). Intrinsic and extrinsic motivation. *Review of Economic Studies* 70(3): 489–520.
- Benartzi, S. and Thaler, R. H. (1995). Myopic loss-aversion and the equity premium puzzle. *Quarterly Journal of Economics* 110(1): 73–92.
- Ben-David, I., Graham, J., and Harvey, C. (2013). Managerial miscalibration. *Quarterly Journal of Economics* 128(4): 1547–84.
- Biais, B. and Weber, M. (2009). Hindsight bias, risk perception and investment performance. *Management Science* 55(6): 1018–29.
- Binmore, K. (1991). *Playing for Real: A Text on Game Theory*. Oxford: Oxford University Press.
- Birnbaum, M. H. (2008). New paradoxes of risky decision making. *Psychological Review* 115(2): 463–501.
- Blaug, M. (1992). *The Methodology of Economics, Or, How Economists Explain*. Cambridge and New York: Cambridge University Press.
- Bohm, P., Lindén, J., and Sonnegård, J. (1997). Eliciting reservation prices: Becker-DeGroot-Marschak mechanisms vs. markets. *Economic Journal* 107: 1079–89.
- Bolton, G. E. and Ockenfels, A. (1998). Strategy and equity: an ERC-analysis of the Guth-van Damme game. *Journal of Mathematical Psychology* 42(2): 215–26.
- Bolton, G. E. and Ockenfels, A. (2000). ERC: a theory of equity, reciprocity, and competition. *American Economic Review* 90(1): 166–93.
- Bolton, G. E., Zwick, R., and Katok, E. (1998). Dictator game giving: rules of fairness versus acts of kindness. *International Journal of Game Theory* 27(2): 269–99.
- Bonner, S. E., Hastie, R., Sprinkle, G. B., and Young, S. M. (2000). A review of the effects of financial incentives on performance in laboratory tasks: implications for management accounting. *Journal of Management Accounting Research* 13: 19–64.
- Bosch-Domènech, A., Montalvo, J. G., Nagel, R., and Satorra, A. (2002). One, two, (three), infinity, . . . newspaper and lab beauty-contest experiments. *American Economic Review* 92(5): 1687–701.
- Bowles, S. and Gintis, H. (2011). *A Cooperative Species: Human Reciprocity and Evolution*. Princeton, NJ: Princeton University Press.
- Braga, J. and Starmer, C. (2005). Preference anomalies, preference elicitation and the Discovers Preference Hypothesis. *Environmental and Resource Economics* 32: 55–89.
- Breitmoser, Y. (2012). Strategic reasoning in p-beauty contests. *Games and Economic Behavior* 75(2): 555–69.
- Brown, J. N. and Rosenthal, R. W. (1990). Testing the minimax hypothesis: a re-examination of O'Neill's game experiment. *Econometrica* 58(5): 1065–81.
- Brown, M., Falk, A., and Fehr, E. (2004). Relational contracts and the nature of market interactions. *Econometrica* 72(3): 747–80.
- Brown, M., Falk, A., and Fehr, E. (2012). Competition and relational contracts: the role of unemployment as a disciplinary device. *Journal of the European Economic Association* 10(4): 887–907.
- Bruhin, A., Fehr-Duda, H., and Epper, T. (2010). Risk and rationality: uncovering heterogeneity in probability distortion. *Econometrica* 78(4): 1375–412.
- Bruni, L. and Sugden, R. (2007). The road not taken: how psychology was removed from economics, and how it might be brought back. *Economic Journal* 117: 146–73.
- Camerer, C. F. (2000). Prospect theory in the wild: evidence from the field. In D. Kahneman and A. Tversky (eds.), *Choices, Values and Frames*. Cambridge: Cambridge University Press, pp. 288–300.
- Camerer, C. F. (2003a). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton University Press.
- Camerer, C. F. (2003b). Behavioural studies of strategic thinking in games. *Trends in Cognitive Sciences* 7(5): 225–31.
- Camerer, C. F. (2015). The promise and success of lab-field generalizability in experimental economics: a critical reply to Levitt and List. In G. R. Fréchette and A. Schotter (eds.), *Handbook of Experimental Economic Methodology*. Oxford: Oxford University Press. pp. 249–95.

- Camerer, C. F., Babcock, L., Loewenstein, G., and Thaler, R. H. (1997). Labor supply of New York City cabdrivers: one day at a time. *Quarterly Journal of Economics* 112(2): 407–41.
- Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics* 119(3): 861–98.
- Camerer, C. F. and Hogarth, R. M. (1999). The effects of financial incentives in experiments: a review and capital-labor-production framework. *Journal of Risk and Uncertainty* 19(1–3): 7–42.
- Camerer, C. F. and Loewenstein, G. (2004). Behavioral economics: past, present, future. In C. F. Camerer, G. Loewenstein, and M. Rabin, (eds.), *Advances in Behavioral Economics*. New York: Russell Sage, pp. 3–51.
- Cameron, L. A. (1999). Raising the stakes in the ultimatum game: experimental evidence from Indonesia. *Economic Inquiry* 37(1): 47–59.
- Caplin, A., Dean, M., and Martin, D. (2011). Search and satisficing. *American Economic Review* 101(7): 2899–922.
- Carpenter, J. P. and Seki, E. (2011). Do social preferences increase productivity? Field experimental evidence from fishermen in Toyama Bay. *Economic Inquiry* 49(2): 612–30.
- Casscells, W., Schoenberger, A. and Graboys, T. B. (1978). Interpretation by physicians of clinical laboratory results. *New England Journal of Medicine* 299: 999–1000.
- Chamberlin, E. H. (1948). An experimental imperfect market. *Journal of Political Economy* 56(2): 95–108.
- Chapman, G. B. and Bornstein, B. H. (1996). The more you ask for, the more you get: Anchoring in personal injury verdicts. *Applied Cognitive Psychology* 10: 519–40.
- Charness, G. and Rabin, M. (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics* 117(3): 817–69.
- Chateauneuf, A. and Wakker, P. P. (1999). An axiomatization of cumulative prospect theory for decision under risk. *Journal of Risk and Uncertainty* 18(2): 137–45.
- Chen, M. K., Lakshminarayanan, V., and Santos, L. (2006). How basic are behavioral biases? Evidence from capuchin monkey trading behavior. *Journal of Political Economy* 114(3): 517–32.
- Cherry, T., Frykblom, P., and Shogren, J. (2002). Hardnose the dictator. *American Economic Review* 92(4) 1218–21.
- Cleave, B. L., Nikiforakis, N., and Slonim, R. (2012). Is there selection bias in laboratory experiments? The case of social and risk preferences. *Experimental Economics* 16(3): 372–82.
- Clotfelter, C. T. and Cook, P. J. (1993). The gambler's fallacy in lottery play. *Management Science* 39: 1521–25.
- Colander, D. and Kupers, R. (2014). *Complexity and the Art of Public Policy: Solving Society's Problems from the Bottom Up*. Princeton, NJ: Princeton University Press.
- Cooper, D. J. and Kagel, J. H. (2015). Other-regarding preferences: a selective survey of experimental results. In J. H. Kagel and A. E. Roth (eds.), *The Handbook of Experimental Economics, Volume 2*. Princeton University Press, forthcoming.
- Cosmides, L. and Tooby, J. (1996). Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty. *Cognition* 58: 1–73.
- Costa-Gomes, M., Crawford, V. P., and Broseta, B. (2001). Cognition and behavior in normal-form games: an experimental study. *Econometrica* 69(5): 1193–235.
- Costa-Gomes, M. A. and Crawford, V. P. (2006). Cognition and behavior in two-person guessing games: an experimental study. *American Economic Review* 96(5): 1737–68.
- Crawford, V. P. (2007). Let's talk it over: coordination via preplay communication with level-k thinking. Mimeo, UCLA Department of Economics.
- Crawford, V. P., Costa-Gomes, M. A., and Iriberri, N. (2013). Structural models of non-equilibrium strategic thinking: theory, evidence, and applications. *Journal of Economic Literature* 51(1): 5–62.
- Crawford, V. P., Gneezy, U., and Rottenstreich, Y. (2008). The power of focal points is limited: even minute payoff asymmetry may yield large coordination failures. *American Economic Review* 98(4): 1443–58.
- Crawford, V. P. and Iriberri, N. (2007a). Level-k auctions: can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica* 75(6): 1721–70.
- Crawford, V. P. and Iriberri, N. (2007b). Fatal attraction: salience, naïveté, and sophistication in experimental hide-and-seek games. *American Economic Review* 97(5): 1731–50.

- Crawford, V. P., Kugler, T., Neeman, Z., and Pauzner, A. (2009). Behaviorally optimal auction design: an example and some observations. *Journal of the European Economic Association* 7(2–3): 377–87.
- Cutler, D., Poterba, J., and Summers, L., (1991). Speculative dynamics. *Review of Economic Studies* 58: 529–46.
- Dana, J., Weber, R. A., and Kuang, J. X. (2007). Exploiting moral wriggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory* 33: 67–80.
- Davis, D. D. and Holt, C. A. (1993). *Experimental economics*. Princeton, NJ: Princeton University Press.
- De Bondt, W. and Thaler, R. H. (1985). Does the stock market overreact? *Journal of Finance* 40: 793–808.
- De Giorgi, E. G. and Legg, S. (2012). Dynamic portfolio choice and asset pricing with narrow framing and probability weighting. *Journal of Economic Dynamics and Control* 36(7): 951–72.
- De Long, J. B., Shleifer, A., Summers, L., and Waldmann, R. (1990). Noise trader risk in financial markets. *Journal of Political Economy* 98: 703–38.
- De Marchi, N. and Blaug, M. (eds.) (1991). *Appraising Economic Theories: Studies in the Methodology of Research Programmes*. Cheltenham: Edward Elgar.
- De Martino B., Kumaran D., Seymour B., and Dolan R. J. (2006). Frames, biases, and rational decision-making in the human brain. *Science* 313: 684–87.
- de Meza, D. and Webb, D. C. (2007). Incentive design under loss aversion. *Journal of the European Economic Association* 5(1): 66–92.
- De Quervain, D. Fischbacher, U. Treyer, V., et al. (2004). The neural basis of altruistic punishment. *Science* 305(5688): 1254–8.
- Dekel, E. and Lipman, B. L. (2010). How (not) to do decision theory. *Annual Review of Economics* 2: 257–82.
- Dekel, E. and Siniscalchi, M. (2015). Epistemic game theory. In H. P. Young and S. Zamir (eds.), *Handbook of Game Theory with Economic Applications, Volume 4*. Elsevier, pp. 619–702.
- DellaVigna, S. and Malmendier, U. (2004). Contract design and self-control: theory and evidence. *Quarterly Journal of Economics* 119(2): 353–402.
- Dhami, S. and al-Nowaihi, A. (2007). Why do people pay taxes: expected utility versus prospect theory. *Journal of Economic Behavior and Organization* 64(1): 171–92.
- Dhami, S. and al-Nowaihi, A. (2010a). Optimal income tax taxation in the presence of tax evasion: expected utility versus prospect theory. *Journal of Economic Behavior and Organization* 75(2): 313–37.
- Dhami, S. and al-Nowaihi, A. (2010b). A proposal to combine “prospect theory” and “cumulative prospect theory” University of Leicester. Discussion Papers in Economics 10/11.
- Dhami, S. and al-Nowaihi, A. (2010c). Existence of a Condorcet winner when voters have other-regarding preferences. *Journal of Public Economic Theory* 12(5): 897–922.
- Dhami, S. and al-Nowaihi, A. (2010d). Redistributive policy with heterogeneous social preferences of voters. *European Economic Review* 54(6): 743–59.
- Dhami, S. and al-Nowaihi, A. (2013). Stochastic dominance for Fehr-Schmidt preferences. University of Leicester, Discussion Papers in Economics. Working Paper No. 13/09.
- Diamond, P. and Koszegi, B. (2003). Quasi-hyperbolic discounting and retirement. *Journal of Public Economics* 87(9–10): 1839–72.
- Duesenberry, J. S. (1949). *Income, Saving, and the Theory of Consumer Behavior*. Cambridge, MA: Harvard University Press.
- Duflo, E., Kremer, M., and Robinson, J. (2011). Nudging farmers to use fertilizer: theory and experimental evidence from Kenya. *American Economic Review* 101(6): 2350–90.
- Dufwenberg, M., Heidhues, P., Kirchsteiger, G., et al. (2011). Other-regarding preferences in general equilibrium. *Review of Economic Studies* 78(2): 613–39.
- Dufwenberg, M. and Kirchsteiger, G. (2004). A theory of sequential reciprocity. *Games and Economic Behavior* 47(2): 268–98.
- Eckel, C. C. and Gintis, H. (2010). Blaming the messenger: notes on the current state of experimental economics. *Journal of Economic Behavior and Organization* 73(1): 109–19.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics* 75(4): 643–69.
- Englich, B. and Mussweiler, T. (2001). Sentencing under uncertainty: anchoring effects in the

- courtroom. *Journal of Applied Social Psychology* 31: 1535–51.
- Englich, B., Mussweiler, T., and Strack, F. (2006). Playing dice with criminal sentences: the influence of irrelevant anchors on experts' judicial decision makings. *Personality and Social Psychology Bulletin* 32(2): 188–200.
- Erev, I., Ert, E., and Roth, A. E. (2010). A choice prediction competition: choices from experience and from description. *Journal of Behavioral Decision Making* 23: 15–47.
- Ericson K. M. M., White J. M., Laibson, D., and Cohen, J. D. (2015). Money earlier or later? Simple heuristics explain intertemporal choices better than delay discounting does. *Psychological Science* 26(6): 826–33.
- Falk, A. and Fischbacher, U. (2006). A theory of reciprocity. *Games and Economic Behavior* 54(2): 293–315.
- Falk, A. and Heckman, J. J. (2009). Lab experiments are a major source of knowledge in the social sciences. *Science* 326: 535–38.
- Falk, A., Meier, S., and Zehnder, C. (2013). Do lab experiments misrepresent social preferences? The case of self-selected student samples. *Journal of the European Economic Association* 11(4): 839–52.
- Fang, H. M. and Silverman, D. (2009). Time-inconsistency and welfare program participation: evidence from the NLSY. *International Economic Review* 50(4): 1043–77.
- Farmer, J. D. and Foley, D. (2009). The economy needs agent-based modelling. *Nature* 460: 685–6.
- Fehr, E. and Fischbacher, U. (2005). Human altruism: proximate patterns and evolutionary origins. *Analyse & Kritik* 27(1): 6–47.
- Fehr, E. and Gächter, S. (1999). Cooperation and punishment in public goods experiments. University of Zurich, Institute for Empirical Research in Economics, Working Paper 10.
- Fehr, E. and Gächter, S. (2000a). Cooperation and punishment in public goods experiments. *American Economic Review* 90(4): 980–94.
- Fehr, E. and Gächter, S. (2000b). Fairness and retaliation: the economics of reciprocity. *Journal of Economic Perspectives* 14(3): 159–81.
- Fehr, E., Gächter, S., and Kirchsteiger, G. (1997). Reciprocity as a contract enforcement device: experimental evidence. *Econometrica* 65(4): 833–60.
- Fehr, E., Hart, O., and Zehnder, C. (2011). Contracts as reference points: experimental evidence. *American Economic Review* 101(2): 493–525.
- Fehr, E., Kirchsteiger, G., and Riedl, A. (1998). Gift exchange and reciprocity in competitive experimental markets. *European Economic Review* 42(1): 1–34.
- Fehr, E., Klein, A., and Schmidt, K. M. (2007). Fairness and contract design. *Econometrica* 75(1): 121–54.
- Fehr, E. and List, J. A. (2004). The hidden costs and returns of incentives: trust and trustworthiness among CEOs. *Journal of the European Economic Association* 2(5): 743–71.
- Fehr, E. and Schmidt K. (1999). A theory of fairness, competition and cooperation. *Quarterly Journal of Economics* 114(3): 817–68.
- Fehr, E. and Schmidt, K. (2006). The economics of fairness, reciprocity and altruism: Experimental evidence and new theories. In S.-C. Kolm and J. M. Ythier (eds.), *Handbook of the Economics of Giving, Altruism and Reciprocity, Volume 1*. Elsevier.
- Fehr, E. and Williams, T. (2015). The endogenous emergence of efficient sanctioning institutions. Mimeo, Department of Economics, University of Zurich.
- Fehr-Duda, H. and Epper, T. (2012). Probability and risk: foundations and economic implications of probability-dependent risk preferences. *Annual Review of Economics* 4: 567–93.
- Feynman, R. (1965). *The Character of Physical Law*. Cambridge, MA: MIT Press.
- Fischbacher, U. and Gächter, S. (2010). Social preferences, beliefs, and the dynamics of free riding in public goods experiments. *American Economic Review* 100(1): 541–56.
- Fischbacher, U., Gächter, S., and Fehr, E. (2001). Are people conditionally co-operative? Evidence from a public goods experiment. *Economics Letters* 71(3): 397–404.
- Fischhoff, B. (1975). Hindsight $\neq$ foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance* 1: 288–99.
- Fischhoff, B., Slovic, P., Lichtenstein S., Read S., and Combs, B. (1978). How safe is safe enough? A psychometric study of attitudes toward technological risks and benefits. *Policy Sciences* 9: 127–52.

- Fouraker, L. E. and Siegel, S. (1963). *Bargaining Behavior*. New York: McGraw-Hill.
- Frascarda, J. (1999). Cognition, emotion and other inescapable dimensions of human experience. *Visible Language* 33: 74–87.
- Fréchette, G. R. (2015). Laboratory experiments: professionals versus students. In G. R. Fréchette and A. Schotter (eds.), *Handbook of Experimental Economic Methodology*. Oxford: Oxford University Press, pp. 360–90.
- Frederick, S., Loewenstein, G., and O'Donoghue, T. (2002). Time discounting and time preferences: a critical review. *Journal of Economic Literature* 40(2): 351–401.
- Friedman, M. (1953). *The Methodology of Positive Economics*. Chicago: University of Chicago Press.
- Froot, K. and E. Dabora (1999). How are stock prices affected by the location of trade? *Journal of Financial Economics* 53: 189–216.
- Fryer, R. G., Levitt, S. D., List, J. A., and Sadoff, S. (2012). Enhancing the efficacy of teacher incentives through loss aversion: a field experiment. NBER. Working Paper 8237.
- Fudenberg, D. and Levine, D. K. (2006). A dual self model of impulse control. *American Economic Review* 96(5): 1449–76.
- Gächter, S., Renner, E., and Sefton, M. (2008). The long-run benefits of punishment. *Science* 322: 1510.
- Galinsky, A. D. and Mussweiler, T. (2001). First offers as anchors: the role of perspective-taking and negotiator focus. *Journal of Personality and Social Psychology* 81(4): 657–69.
- Geanakoplos, J., Pearce, D., and Stacchetti, E. (1989). Psychological games and sequential rationality. *Games and Economic Behavior* 1(1): 60–79.
- Gigerenzer, G. and Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review* 102: 684–704.
- Gigerenzer, G. and Selten, R. (2001). Rethinking rationality. In G. Gigerenzer and R. Selten (eds.), *Bounded rationality: The Adaptive Toolbox. Dahlem Workshop Report*. Cambridge, MA: MIT Press, pp. 1–12.
- Gigerenzer, G., Todd, P. M., and the ABC Research Group (1999). *Simple Heuristics That Make Us Smart*. New York: Oxford University Press.
- Gilboa, I. (2009). *Theory of Decision Under Uncertainty*. Cambridge: Cambridge University Press.
- Gilboa, I., Postlewaite, A., Samuelson, L., and Schmeidler, D. (2014). Economic models as analogies. *Economic Journal* 124: F513–F533.
- Gill, D. and Prowse, V. (2015). Cognitive ability, character skills, and learning to play equilibrium: a level-k analysis. Forthcoming in *Journal of Political Economy*.
- Gilovich, T., Griffin, D., and Kahneman, D. (2002). *Heuristics and Biases: the Psychology of Intuitive Judgment*. New York: Cambridge University Press.
- Gintis, H. (2009). *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences*. Princeton, NJ: Princeton University Press.
- Gintis, H. (2015). Modeling homo-socialis: a reply to critics. *Review of Behavior Economics* 2: 211–37.
- Glaser, M., Langer, T., Reynnders, J., and Weber, M. (2007a). Framing effects in stock market forecasts: the difference between asking for prices and asking for returns. *Review of Finance* 11: 325–57.
- Glaser, M., Langer, T., and Weber, M. (2007b). On the trend recognition and forecasting ability of professional traders. *Decision Analysis* 4: 176–93.
- Glimcher, P. W., Camerer, C. F., Fehr, E., and Pol-drack, R. A. (2009). *Neuroeconomics*. Academic Press, Elsevier Inc.
- Glimcher, P. W. and Fehr, E. (eds.) (2014). *Neuroeconomics*. Elsevier Inc.
- Gneezy, U. and Rustichini, A. (2000). Pay enough or don't pay at all. *Quarterly Journal of Economics* 115(3): 791–810.
- Gode, D. K. and Sunder, S. (1993). Allocative efficiency of markets with zero intelligence traders: market as a partial substitute for individual rationality. *Journal of Political Economy* 101(1): 119–37.
- Godfrey-Smith, P. (2006). The strategy of model-based science. *Biology and Philosophy* 21: 725–40.
- Godfrey-Smith, P. (2009). Models and fictions in science. *Philosophical Studies* 143: 101–16.
- Goldfarb, A. and Xiao, M. (2011). Who thinks about the competition? Managerial ability and strategic entry in US local telephone markets. *American Economic Review* 101(7): 3130–61.
- Goldfarb, A. and Yang, B. (2009). Are all managers created equal? *Journal of Marketing Research* 46(5): 612–22.
- Grether, D. M. and Plott, C. (1979). Economic theory of choice and the preference reversal

- phenomenon. *American Economic Review*. 69(4): 623–38.
- Grosskopf, B. and Nagel, R. (2008). The two-person beauty contest. *Games and Economic Behavior* 62(1): 93–9.
- Gruber, J. H. and Mullainathan, S. (2005). Do cigarette taxes make smokers happier? *B.E. Journals: Advances in Economic Analysis and Policy* 5(1): 1–45.
- Guala, F. (2008). Experimental economics, history of. In S. N. Durlauf and L. E. Blume (eds.), *The New Palgrave Dictionary of Economics*. 2nd edition. Palgrave Macmillan.
- Guilbault, R. L., Bryant, F. B., Brockway, J. H., and Posavac, E. J. (2004). A meta-analysis of research on hindsight bias. *Basic and Applied Social Psychology* 26: 103–17.
- Gul, F. and Pesendorfer, W. (2001). Temptation and self control. *Econometrica* 69(6): 1403–36.
- Guryan, J. and Kearney, M. S. (2008). Gambling at lucky stores: empirical evidence from state lottery sales. *American Economic Review* 98(1): 458–73.
- Güth, W., Schmittberger, R., and Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization* 3(4): 367–88.
- Haley, K. J. and Fessler, D. M. T. (2005). Nobody's watching? Subtle cues affect generosity in an anonymous economic game. *Evolution and Human Behavior* 26(3): 245–56.
- Hands, D. W. (1991). The problem of excess content: Economics, novelty, and a long Popperian tale. In M. Blaug and N. DeMarchi (eds.), *Appraising Economic Theories*. Aldershot: Edward Elgar. pp. 58–75.
- Hands, D. W. (2001). *Reflections Without Rules: Economic Methodology and Contemporary Science Theory*. Cambridge: Cambridge University Press.
- Hargreaves Heap, S., Rojo-Arjona, D., and Sugden, R. (2014). How portable is level-0 behavior? A test of level-k theory in games with non-neutral frames. *Econometrica* 82(3): 1133–51.
- Harsanyi, J. C. (1973). Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory* 2(1): 1–23.
- Hausman, D. (1992). *The Inexact and Separate Science of Economics*. Cambridge: Cambridge University Press.
- Heath, C. and Tversky, A. (1991). Preference and belief: ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty* 4: 5–28.
- Henrich, J., Boyd, R., Bowles, S., et al. (2001). Cooperation, reciprocity and punishment in fifteen small-scale societies. *American Economic Review* 91: 73–8.
- Herrmann, B., Thöni, C., and Gächter, S. (2008). Antisocial punishment across societies. *Science* 319(5868): 1362–67.
- Herbst, D. and Mas, A. (2015). Peer effects on worker output in the lab generalize to the field. Forthcoming in *Science*.
- Hertwig, R. and Ortmann, A. (2001). Experimental practices in economics: a methodological challenge for psychologists? *Behavioral and Brain Sciences* 24(3): 383–403.
- Hey, J. D. (2014). Choice under uncertainty: empirical methods and experimental results. In M. J. Machina and W. K. Viscusi (eds.), *Handbook of the Economics of Risk and Uncertainty Volume 1*. Oxford: North Holland, pp. 809–50.
- Hey, J. D. and Orme, C. (1994). Investigating generalizations of expected utility theory using experimental data. *Econometrica* 62(6): 1291–326.
- Ho, T., Camerer, C. F., and Weigelt, K. (1998). Iterated dominance and iterated best response in experimental “p-beauty contests.” *American Economic Review* 88(4): 947–69.
- Hoffman, E., McCabe, K. A., and Smith, V. L. (1998). Behavioral foundations of reciprocity: experimental economics and evolutionary psychology. *Economic Inquiry* 36: 335–52.
- Hogarth, R. M., Gibbs, B. J., McKenzie, C. R. M., and Marquis, M. A. (1991). Learning from feedback: exactingness and incentives. *Journal of Experimental Psychology: Learning, Memory and Cognition* 17: 734–52.
- Hong, H. and Stein, J. (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *Journal of Finance* 54: 2143–84.
- Hong, H. and Stein, J. (2007). Disagreement and the stock market. *Journal of Economic Perspectives* 21: 109–28.
- Huck, S., Lünser, G. K., and Tyran, J.-R. (2012). Competition fosters trust. *Games and Economic Behavior* 76(1): 195–209.
- Ippolito, R. A. (1989). Efficiency with costly information: a study of mutual fund performance,

- 1965–1984. *Quarterly Journal of Economics* 104(1): 1–23.
- Jegadeesh, N. and Titman, S. (1993). Returns to buying winners and selling losers: implications for stock market efficiency. *Journal of Finance* 48: 65–91.
- Jehle, G. J. and Reny, P. J. (2011). *Advanced Microeconomic Theory*. 3rd edition. Prentice Hall.
- Kagel, J. H. and Roth, A. E. (1995). *The Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press.
- Kahneman, D. (2003). Maps of bounded rationality: psychology for behavioral economics. *American Economic Review* 93(5): 1449–75.
- Kahneman, D. (2011). *Thinking Fast and Slow*. New York, Farrar, Strauss, Giroux.
- Kahneman, D., Knetsch, J. L., and Thaler, R. H. (1986). Fairness as a constraint on profit seeking: entitlements in the market. *American Economic Review* 76(4): 728–41.
- Kahneman, D., Knetsch, J. L., and Thaler, R. H. (1991). Anomalies: the endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives* 5(1): 193–206.
- Kahneman, D. and Tversky, A. (1972). On prediction and judgment. *Oregon Research Institute Bulletin* 12(4).
- Kahneman, D. and Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica* 47(2): 263–91.
- Kahneman, D. and Tversky, A. (1984). Choices, values, and frames. *The American Psychologist* 39: 341–50.
- Kahneman, D. and Tversky, A. (1996). On the reality of cognitive illusions: a reply to Gigerenzer's critique. *Psychological Review* 103: 582–91.
- Kahneman, D. and Tversky, A. (2000). *Choices, Values and Frames*. Cambridge: Cambridge University Press.
- Kessler, J. B. and Vesterlund, L. (2015). The external validity of laboratory experiments: the misleading emphasis on quantitative effects. In G. R. Fréchette and A. Schotter (eds.), *Handbook of Experimental Economic Methodology*. Oxford: Oxford University Press: Oxford. pp. 391–406.
- Kim, E. H., Morse, A., and Zingales, L. (2006). What has mattered to economics since 1970. *Journal of Economic Perspectives* 20, 189–202.
- Kirchler, E. and Maciejovsky, B. (2002). Simultaneous over- and underconfidence: evidence from experimental asset markets. *Journal of Risk and Uncertainty* 25: 65–85.
- Köszegi, B. and Rabin, M. (2006). A model of reference-dependent preferences. *Quarterly Journal of Economics* 121(4): 1133–65.
- Kothiyal, A., Spinu, V., and Wakker, P. P. (2014). An experimental test of prospect theory for predicting choice under ambiguity. *Journal of Risk and Uncertainty* 48(1): 1–17.
- Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Laibson, D. (1994). Essays in hyperbolic discounting. MIT, PhD Dissertation.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 112(2): 443–78.
- Laibson, D. (1998). Life-cycle consumption and hyperbolic discount functions. *European Economic Review* 42(3–5): 861–71.
- Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In I. Lakatos and A. Musgrave (eds.), *Criticism and the Growth of Knowledge*. Cambridge: Cambridge University Press.
- Lakonishok, J., Shleifer, A., and Vishny, R. W. (1992). The impact of institutional trading on stock prices. *Journal of Financial Economics* 32: 23–44.
- Lambdin, C. G. and Shaffer, V. A. (2009). Are within-subjects designs transparent? *Judgement and Decision Making* 4(7): 554–66.
- Lee, C. M. C., Shleifer, A., and Thaler, R. H. (1991). Investor sentiment and the closed-end fund puzzle. *Journal of Finance* 46: 76–110.
- Levitt, S. D. and List, J. A. (2007). What do laboratory experiments measuring social preferences reveal about the real world? *Journal of Economic Perspectives* 21(2): 153–74.
- Levitt, S. D. and List, J. A. (2008). Homo economicus evolves. *Science* 319: 909–10.
- Levitt, S. D., List, J. A., and Reiley, D. H. (2010). What happens in the field stays in the field: professionals do not play minimax in laboratory experiments. *Econometrica* 78(4): 1413–34.
- Lichtenstein, S. and Slovic, P. (1973). Response-induced reversals of preference in gambling: an extended replication in Las Vegas. *Journal of Experimental Psychology* 101(1): 16–20.
- Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M., and Combs, B. (1978). Judged frequency of lethal events. *Journal of Experimental*

- Psychology: Human Learning and Memory* 4: 551–78.
- Lipsey, R. G. (1979). *An Introduction to Positive Economics*. 5th edition. London: Weidenfield and Nicholson.
- List, J. A. (2003). Does market experience eliminate market anomalies? *Quarterly Journal of Economics* 118(1): 41–71.
- List, J. A. (2006). The behavioralist meets the market: measuring social preferences and reputation effects in actual transactions. *Journal of Political Economy* 114(1): 1–37.
- Locke, E. A. and Latham, G. P. (1990). *A Theory of Goal Setting and Task Performance*. Englewood Cliffs: Prentice-Hall.
- Loewenstein, G. F. (1999). Experimental economics from the vantage-point of behavioural economics. *Economic Journal* 109(453): 25–34.
- Loewenstein, G. F. and Prelec, D. (1992). Anomalies in intertemporal choice: evidence and an interpretation. *Quarterly Journal of Economics* 107(2): 573–97.
- Loewenstein, G. F., Weber, E. U., Hsee, C. K., & Welch, E. S. (2001). Risk as feelings. *Psychological Bulletin* 127: 267–86.
- Lord, C. G., Ross, L., and Lepper, M. R. (1979). Biased assimilation and attitude polarization: the effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology* 37(11): 2098–109.
- McNeil, B. J., Pauker, S. G., Sox, H. C., Jr., and Tversky, A. (1982). On the elicitation of preferences for alternative therapies. *New England Journal of Medicine* 306: 1259–62.
- Mak, V., Gisches, E. J., and Rapoport, A. (2015). Route vs. segment: an experiment on real-time travel information in congestible networks. *Production and Operations Management* 24(6): 947–60.
- Malmendier, U. and Tate, G. (2005). CEO overconfidence and corporate investment. *Journal of Finance* 60: 2661–700.
- Malmendier, U. and Tate, G. (2008). Who makes acquisitions? CEO overconfidence and the market's reaction. *Journal of Financial Economics* 89: 20–43.
- Markowitz, H. (1952). The utility of wealth. *Journal of Political Economy* 60(2): 151–58.
- Mas-Collel, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. New York: Oxford University Press.
- Meehl, P. E. and Rose, A. (1955). Antecedent probability and the efficiency of psychometric signs, patterns or cutting scores. *Psychological Bulletin* 52: 194–216.
- Michael, R., Thaler, R. H., and Womack K. (1995). Price reactions to dividend initiations and omissions. *Journal of Finance* 50: 573–608.
- Mitchell, M. (2009). *Complexity: A Guided Tour*. Oxford: Oxford University Press.
- Mookherjee, D. and Sopher, B. (1994). Learning behavior in an experimental matching pennies game. *Games and Economic Behavior* 7(1): 62–91.
- Mookherjee, D. and Sopher, B. (1997). Learning and decision costs in experimental constant sum games. *Games and Economic Behavior* 19(1): 97–132.
- Mussweiler, T. and Schneller, K. (2003). What goes up must come down: how charts influence decisions to buy and sell stocks. *Journal of Behavioral Finance* 4(3): 121–30.
- Nagel, R. (1995). Unraveling in guessing games: an experimental study. *American Economic Review* 85(5): 1313–26.
- Nash, J. F. (1950). Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences of the United States of America* 36(1): 48–9.
- Neilson, W. S. (2006). Axiomatic reference-dependence in behavior toward others and toward risk. *Economic Theory* 28(3): 681–92.
- Northcraft, G. B. and Neale, M. A. (1987). Experts, amateurs, and real estate: An anchoring-and-adjustment perspective on property pricing decisions. *Organizational Behaviour and Human Decision Processes* 39: 84–97.
- Novemsky, N. and Kahneman, D. (2005). The boundaries of loss aversion. *Journal of Marketing Research* 42(2): 119–28.
- Odean, T. (1998). Are investors reluctant to realize their losses? *Journal of Finance* 53: 1775–98.
- O'Donoghue, T. and Rabin, M. (2001). Choice and procrastination. *Quarterly Journal of Economics* 116(1): 121–60.
- O'Donoghue, T. and Rabin, M. (2006). Optimal sin taxes. *Journal of Public Economics* 90(10–11): 1825–49.
- O'Neill, B. (1987). Nonparametric test of the minimax theory of two-person zerosum games. *Proceedings of the National Academy of Sciences of the United States of America* 84(7): 2106–9.

- Page, T., Putterman, L., and Unel, B. (2005). Voluntary association in public goods experiments: reciprocity, mimicry and efficiency. *Economic Journal* 115(506): 1032–53.
- Palacios-Huerta, I. (2003). Professionals play minimax. *Review of Economic Studies* 70(2): 395–415.
- Palacios-Huerta I. and Volij, O. (2008). Experientia docet: professionals play minimax in laboratory experiments. *Econometrica* 76(1): 71–115.
- Pelphs, E. and Pollak, R. A. (1968). On second best national savings and game equilibrium growth. *Review of Economic studies* 35(2): 185–99.
- Pogarsky, G. and Babcock, L. (2001). Damage caps, motivated anchoring and bargaining impasse. *Journal of Legal Studies* 30: 143–59.
- Pope, D. and Schweitzer, M. (2011). Is Tiger Woods loss averse? Persistent bias in the face of experience, competition, and high stakes. *American Economic Review* 101(1): 129–57.
- Popper, K. (1934). *Logik der Forschung* (Hutchinson & Company published the translation by Karl Popper titled *The Logic of Scientific Discovery* in 1959).
- Popper, K. (1963). *Conjectures and Refutations: The Growth of Scientific Knowledge*. London: Routledge and Kegan Paul.
- Prelec, D. (1998). The probability weighting function. *Econometrica* 66(3): 497–527.
- Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior and Organization* 3(4): 323–43.
- Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review* 83(5): 1281–302.
- Rabin, M. (2002). Inference by believers in the law of small numbers. *Quarterly Journal of Economics* 117: 775–816.
- Rabin, M. and Schrag, J. (1999). First impressions matter: a model of confirmatory bias. *Quarterly Journal of Economics* 114(1): 37–82.
- Rachlinski, J. J. (1998). A positive psychological theory of judging in hindsight. *University of Chicago Law Review* 65: 571–625.
- Rapoport, A. and Budescu, D. V. (1997). Randomization in individual choice behavior. *Psychological Review* 104: 603–17.
- Read, D. (2005). Monetary incentives, what are they good for? *Journal of Economic Methodology* 12(2): 265–76.
- Redelmeier, D. A. and Shafir, E. (1995). Medical decision making in situations that offer multiple alternatives. *Journal of the American Medical Association* 273(4): 302–5.
- Roiser, J. P., De Martino, B., Tan, G. C. Y., et al. (2009). A genetically mediated bias in decision making driven by failure of amygdala control. *Journal of Neuroscience* 29(18): 5985–91.
- Rotemberg, J. J. (2008). Minimally acceptable altruism and the ultimatum game. *Journal of Economics and Organizational Behavior* 66: 457–76.
- Roth, A. E. (1987). Laboratory experimentation in economics. In T. Bewley (ed.), *Advances in Economic Theory, Fifth World Congress*. Cambridge: Cambridge University Press, pp. 269–99.
- Roth, A. E. (1988). Laboratory experimentation in economics: a methodological overview. *Economic Journal* 98: 974–1031.
- Roth, A. E. (1995). Introduction to experimental economics. In J. H. Kagel and A. E. Roth (eds.), *The Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press, pp. 3–109.
- Roth, A. E. (2015). Is experimental economics living up to its promise? In G. R. Fréchet and A. Schotter, (eds.), *Handbook of Experimental Economic Methodology*. Oxford: Oxford University Press, pp. 13–40.
- Roth, B. and Voskort, A. (2014). Stereotypes and false consensus: how financial professionals predict risk preferences. *Journal of Economic Behavior and Organization* 107: 553–65.
- Rubinstein, A. (2006). Dilemmas of an economic theorist, *Econometrica* 74(4): 865–83.
- Rydval, O. and Ortmann, A. (2004). How financial incentives and cognitive abilities affect task performance in laboratory settings: an illustration. *Economics Letters* 85: 315–20.
- Samuelson, P. A. (1937). A note on measurement of utility. *Review of Economic Studies* 4(2): 155–61.
- Samuelson, P. A. and Nordhaus, W. (1985). *Economics*. New York: McGraw Hill.
- Schotter, A. (2015). On the relationship between economic theory and experiments. In G. R. Fréchet and A. Schotter (eds.), *Handbook of Experimental Economic Methodology*. Oxford: Oxford University Press, pp. 58–85.
- Selten, R. (2001). What is bounded rationality? In G. Gigerenzer and R. Selten (eds.), *Bounded rationality: The Adaptive Toolbox*. Dahlem Workshop Report. Cambridge, MA: MIT Press, pp. 1–12.

- Selten, R. (1998). Features of experimentally observed bounded rationality. *European Economic Review* 42(3–5): 413–36.
- Shachat, J. M. (2002). Mixed strategy play and the minimax hypothesis. *Journal of Economic Theory* 104(1): 189–226.
- Shefrin, H. M. and Thaler, R. H. (1988). The behavioral life-cycle hypothesis. *Economic Inquiry* 26(4): 609–43.
- Shiller, R. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* 71: 421–36.
- Shleifer, A. (1986). Do demand curves for stocks slope down? *Journal of Finance* 41: 579–90.
- Shleifer, A. and Vishny, R. (1997). The limits of arbitrage. *Journal of Finance* 52: 35–55.
- Siegel, S. and Fouraker, L. E. 1960. *Bargaining and Group Decision Making*. New York: McGraw-Hill.
- Simon, H. A. (1978). Rational decision-making in business organizations. Nobel Memorial Lecture, December 8, 1978.
- Simon, H. A. (2000). Bounded rationality in social science: today and tomorrow. *Mind and Society* 1(1): 25–39.
- Slonim, R. and Roth, A. E. (1998). Learning in high stakes ultimatum games: an experiment in the Slovak Republic. *Econometrica* 66(3): 569–96.
- Slonim, R. L., Roth A. E., and Erev, I. (2003). Regularities in play of zero-sum games with unique mixed-strategy equilibria. Mimeo, Case Western Reserve University.
- Slonim, R., Wang, C., Garbarino, E., and Merrett, D. (2013). Opting-in: participation bias in economic experiments. *Journal of Economic Behavior and Organization* 90: 43–70.
- Slovic, P., Finucane, M., Peters, E., and MacGregor, D. G. (2002). The affect heuristic. In T. Gilovich, D. Griffin, and D. Kahneman (eds.), *Heuristics and biases: The Psychology of Intuitive Judgment*. New York: Cambridge University Press, pp. 397–420.
- Slovic, P., MacGregor, D. G., Malmfors, T., and Purchase, I. F. H. (1999). *Influence of Affective Processes on Toxicologists' Judgments of Risk (Report No. 99–2)*. Eugene, OR: Decision Research.
- Slovic, P., Monahan, J., and MacGregor, D. G. (2000). Violence risk assessment and risk communication: the effects of using actual cases, providing instruction, and employing probability versus frequency formats. *Law and Human Behavior* 24: 271–96.
- Smith, V. L. (1962). An experimental study of competitive market behavior. *Journal of Political Economy* 70: 111–37.
- Smith, V. L. (1976). Experimental economics: induced value theory. *American Economic Review* 66: 274–9.
- Smith, V. L. (1991). Rational choice: the contrast between economics and psychology. *Journal of Political Economy* 99: 877–97.
- Smith, V. L. (2001). From old issues to new directions in experimental psychology and economics. *Behavioral and Brain Sciences* 24(3): 428–9.
- Smith, V. L. and Walker, J. M. (1993). Monetary rewards and decision cost in experimental economics. *Economic Inquiry* 31: 245–61.
- Solan, L., Rosenblatt, T., and Osherson D. (2008). False consensus bias in contract interpretation. *Columbia Law Review* 108(5): 1272–300.
- Stahl, D. O. and Wilson, P. (1995). On players' models of other players: theory and experimental evidence. *Games and Economic Behavior*. 10: 218–54.
- Stanovich, K. E. (2012). On the distinction between rationality and intelligence: Implications for understanding individual differences in reasoning. In K. Holyoak and R. Morrison (eds.), *The Oxford Handbook of Thinking and Reasoning*. New York: Oxford University Press, pp. 343–65.
- Stanovich, K. E., and West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate [Target article and commentaries]. *Behavioral and Brain Sciences* 23: 645–726.
- Stanovich, K. E. and West, R. F. (2007). Natural Myside bias is independent of cognitive ability. *Thinking & Reasoning* 13(3): 225–47.
- Starmer, C. (2000). Developments in non-expected utility theory: the hunt for a descriptive theory of choice under risk. *Journal of Economic Literature* 38(2): 332–82.
- Starmer, C. and Sugden, R. (1989). Probability and juxtaposition effects: an experimental investigation of the common ratio effect. *Journal of Risk Uncertainty* 2(2): 159–78.
- Strack, F. and Mussweiler, T. (1997). Explaining the enigmatic anchoring effect: mechanisms of

- selective accessibility. *Journal of Personality and Social Psychology* 73: 437–46.
- Strotz, R. H. (1955–1956). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* 23(3): 165–80.
- Tanaka, T., Camerer, C. F., and Nguyen, Q. (2010). Risk and time preferences: linking experimental and household survey data from Vietnam. *American Economic Review* 100(1): 557–71.
- Terrell, D. (1994). A test of the gambler's fallacy: evidence from pari-mutuel games. *Journal of Risk and Uncertainty* 8: 309–17.
- Tesfatsion, L. and Judd, K. L. (eds.) (2006). *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*. North-Holland.
- Tetlock, P. E. (1999). Theory-driven reasoning about possible pasts and probable futures in world politics: are we prisoners of our preconceptions? *American Journal of Political Science* 43(2): 335–66.
- Tetlock, P. E. (2006). *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton, NJ: Princeton University Press.
- Thaler, R. H. (1980). Toward a positive theory of consumer choice. *Journal of Economic Behavior and Organization* 1: 39–60.
- Thaler, R. H. (1981). Some empirical evidence of dynamic inconsistency. *Economics Letters* 8(3): 201–7.
- Thaler, R. H. (2015). *Misbehaving: The Making of Behavioral Economics*. New York: W. W. Norton.
- Tom, S. M., Fox, C. R., Trepel, C., and Poldrack, R. A. (2007). The neural basis of loss aversion in decision-making under Risk. *Science* 315: 515–18.
- Tversky, A. and Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology* 5(2): 207–32.
- Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: heuristics and biases. *Science* 185: 1124–30.
- Tversky, A. and Kahneman, D. (1980). Causal schemas in judgments under uncertainty. In M. Fishbein (ed.), *Progress in Social Psychology*. Hillsdale, NJ: Erlbaum, pp. 49–72.
- Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4): 297–323.
- Tyran, J.-R. and Sausgruber, R. (2006). A little fairness may induce a lot of redistribution in a democracy. *European Economic Review* 50(2): 469–85.
- Varian, H. (1992). *Microeconomic Analysis*. New York: W. W. Norton.
- Veblen, T. (1899). *The Theory of the Leisure Class. An Economic Study of Institutions*. Random House.
- Viscusi, K. W., Huber, J., and Bell, J. (2008). Estimating discount rates for environmental quality from utility-based choice experiments. *Journal of Risk and Uncertainty* 37(2–3): 199–220.
- von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Wakker, P. P. (2010). *Prospect Theory for Risk and Ambiguity*. Cambridge: Cambridge University Press.
- Wakker, P. P. and Deneffe, D. (1996). Eliciting von Neumann Morgenstern utilities when probabilities are distorted or unknown. *Management Science* 42(8): 1131–50.
- Walker, M. and Wooders, J. (2001). Minimax play at Wimbledon. *American Economic Review* 91(5): 1521–38.
- Weyland, K. (1996). Risk taking in Latin American economic restructuring: lessons from prospect theory. *International Studies Quarterly* 40(2): 185–207.
- Wimsatt, C. W. (2007). *Re-Engineering Philosophy for Limited Beings*. Cambridge, MA: Harvard University Press.
- Wooders, J. (2010). Does experience teach? Professionals and minimax play in the lab. *Econometrica* 78(3): 1143–54.
- Zizzo, D. J. (2010). Experimenter demand effects in economic experiments. *Experimental Economics* 13(1): 75–98.

PART 1

Behavioral Economics of Risk, Uncertainty, and Ambiguity

INTRODUCTION TO PART 1

In the simplest economic problem, a single decision maker plays a game against nature. Depending on the action chosen by the decision maker, nature induces a probability distribution over a set of possible outcomes. The probability distribution may be *objective* or *subjective*.

If the probability distribution of the relevant outcomes can be estimated objectively, say, from past data or using frequentist methods, then we have a situation of *risk*. For instance, the probability that one obtains a 3 and a 6 in two successive throws of an unbiased dice. However, in some cases, no objective probabilities may be available, e.g., the dice may be known to be biased, but different people may have different views on the degree of the bias. In this case, decision makers may be able to assign a *subjective probability distribution* over outcomes. This is a situation of *uncertainty*.

A third class of events arise when no objective probabilities are available, nor can any *reasonable* subjective probabilities (that are non-negative and sum up to one) be assigned either. In this case, decision makers may be able to assign reasonable subjective probabilities to outcomes arising from each *source* of potential information (readers familiar with the Ellsberg paradox could think of the two urns in the paradox as two sources). However, when the decision maker's subjective probabilities are elicited by mixing up the different sources, then no reasonable subjective probabilities can be assigned. Such a situation of *source-dependent uncertainty* is known as *ambiguity*.

The central question in decision theory, taken up in Part 1, is how does the decision maker choose his optimal action in the presence of risk, uncertainty, or ambiguity?

Chapter 1 considers the elements of neoclassical decision theory. We mainly focus on *expected utility theory* (EU). Our exposition of EU is necessarily brief because most readers will already be familiar with this material. We examine the preference foundations of EU and highlight the critical axioms that have been subject to empirical testing, particularly the *independence axiom*. We also discuss a celebrated result of Leonard J. Savage, that derives a preference foundation for choices under uncertainty. The resulting representation of preference looks similar to EU and is called *subjective expected utility* (SEU). This formulation showed that there is no essential problem in making choices under uncertainty, provided that we can assign reasonable subjective probabilities to the relevant events.

Most lay people would probably believe that probabilities have something to do with risk. However, under EU, risk attitudes are determined entirely by the shape of the utility function. Further, if the decision maker is risk averse over small stakes for all levels of income, then he will be unreasonably risk averse for large stakes. We also discuss various methods of eliciting the utility function under risk and uncertainty. In the rest of Chapter 1, we consider the main refutations of EU. We use the convenient device of a Marshak–Machina probability triangle to illustrate the nature of the refutations. It turns out that most refutations of EU lie along the edges of the probability triangle, but it is precisely along the edges that most important economic situations of interest also occur. Other refutations that we consider include, *preference reversals*, *description invariance*, and *failure of the reduction of compound lotteries axiom*.

EU has two critical features. These are: (i) *linearity in probabilities*, and (ii) *utilities that depend only on final levels of wealth*. In contrast to these features, empirical evidence suggests that

decision makers engage in non-linear probability weighting and their behavior differs fundamentally in the domain of gains and losses. In Chapter 2, we consider a range of behavioral decision theories that relax at least one or both of these features of EU or even introduce new features; these theories explain the available evidence much better.

A central feature of most behavioral alternatives to EU is *non-linear probability weighting*. This is incorporated through a *probability weighting function* that captures the way in which individuals mentally code objective or subjective probabilities. For aggregate data, one typically finds an *inverse S-shaped* probability weighting, i.e., low probabilities are overweighted and high probabilities are underweighted. At the individual level, however, there is a degree of heterogeneity: One observes convex, concave, and inverse S-shaped weighting.

Rank dependent utility (RDU) under risk and under uncertainty is the most satisfactory alternative to EU that incorporates just non-linear probability weighting. RDU replaces the (objective or subjective) probabilities in an EU functional by *decision weights* that are *cumulative transformations of probabilities*. Decision makers who use RDU never choose first order stochastically dominated options.

Original prospect theory (OPT) and *cumulative prospect theory* (PT) incorporate non-linear probability weighting and a utility function that incorporates *reference dependence*. By expressing all outcomes relative to the reference point, decision makers who follow OPT and PT partition the domain of outcomes into a gain and a loss domain. The utility function is concave in gains and convex in losses. Furthermore, losses bite more than equivalent gains, a phenomenon known as *loss aversion*. Human behavior turns out to be fundamentally different in the two domains. PT can be readily extended to the case of uncertainty and ambiguity. Using the essential insight of cumulative transformations of probabilities from RDU, Kahneman and Tversky altered OPT, which was published in 1979, to PT, which was published in 1992. This ensured that decision makers would not choose stochastically dominated options. However, the psychological richness in the editing phase of OPT was also lost in this process.

Under EU, attitudes to risk are determined entirely by the shape of the utility function. However, under RDU, OPT, PT (and other theories that use non-linear weighting), attitudes to risk depend *jointly* on the shapes of the probability weighting function and the utility function. Attitudes to risk are even richer in PT as compared to RDU. One observes a *fourfold classification of risk* that can explain the simultaneous occurrence of insurance and gambling as well as many other phenomena.

In OPT and PT, the reference point is typically specified to be the *status-quo*, or a *fair/just* entitlement, all of which are eminently reasonable in situations where the individual faces no risk. However, under certain circumstances, the endowment of an individual may be stochastic (e.g., owning a stock that can either go up or down in value with some associated probability). In this case, the reference point itself may be stochastic. This has led to the development of *third generation prospect theory* (PT³) that can incorporate stochastic reference points. Another popular approach is the *Kőszegi–Rabin* framework that allows the reference point to be *stochastic and endogenous* in the sense that it requires a rational expectations property.

PT is the main, though not the only, decision theory under consideration in Chapter 2. The evidence suggests that it is the most satisfactory decision theory under risk and uncertainty.¹ We also discuss the elicitation of the various components of PT, give preference foundations for the theory, and note its limitations. We next consider a range of other behavioral decision theories that take account of factors that are not directly present in RDU, OPT, PT. These theories

¹ PT also applies to ambiguity, as we shall see in Chapter 4.

include *regret theory*, the *theory of disappointment aversion*, the role of *salience* and *context* in making choices among risky alternatives, *case-based decision theory*, and *configural weights models*. Existing theories find it difficult to explain human behavior for events of *very low probability*. In many situations of economic interest, one observes that a fraction of decision makers ignore events of very low probability, while others give it more than its objective salience. We consider a modification of PT, *composite prospect theory* (CPT) that can account for such behavior.

The final topic in this chapter considers the interaction between risk and time preferences, which has been a very active area of research in recent times. We shall continue with this topic in Part 3 of the book.

Chapter 3 considers applications of behavioral theories of decision making. An important application is to the disparity between *willingness to accept* (WTA) of sellers who are endowed with an object and *willingness to pay* (WTP) for the same item by buyers. The disparity between WTA and WTP arises because sellers who are endowed with the object experience loss aversion when they part with it. However, loss aversion does not apply to parting with money, giving rise to the disparity. We also consider a range of other possible explanations of exchange disparities and factors that might weaken or enhance it. Loss aversion is a fundamental trait of human nature that we share with close primate relatives such as chimpanzees. This suggests that it precedes the evolutionary separation of humans and chimpanzees and, thus, we may be hardwired to exhibit it. An important implication of these studies is that there might be *no risk aversion beyond loss aversion*. In other words, economists picking out risk aversion, could actually be picking out loss aversion.

Some of the other applications that we consider in Chapter 3 are as follows. In a two-horse race between EU and PT, we show that PT successfully explains the qualitative and quantitative tax evasion puzzles, but EU fails to do so. *Myopic loss aversion*, the empirically supported tendency to use very short horizons in evaluating gains and losses, explains several phenomena, e.g., the *equity premium puzzle* (high returns on equities relative to bonds that cannot be explained by the riskiness of the two assets, or by risk aversion). We show that under PT, *Rabin's puzzle* that arises under EU (risk aversion over small stakes implies implausible risk aversion over larger stakes) can be explained. We show that once *goals* and *targets* are recognized as reference points, PT is able to account for the relevance of goals in influencing human behavior that cannot be explained by EU and RDU. PT can also explain the backward bending labour supply curve of New York taxi drivers. A range of applications of the finding of an inverse S-shaped probability weighting function are also given, although some of these, such as the *psychology of tail events*, are postponed to Part 7. Finally, we apply PT to issues of optimal choice of contracts, the formation of firms, and the renegotiation of long-term contracts. These topics form a part of the exciting new literature on behavioral contract theory. Further developments in this literature based on other-regarding preferences are considered in Part 2 of the book.

Chapter 4, the final chapter in Part 1, considers models of *ambiguity*. The discussion in the chapter is organized around two main sections. Section 4.3 considers *neoclassical models of ambiguity* and Section 4.4 considers *behavioral models of ambiguity*. This is not a watertight separation of the literature into two convenient categories, but it helps in organizing it. The main distinction is that most, but not all, of the neoclassical models of ambiguity reduce to EU under pure risk, while the behavioral models of ambiguity in Section 4.4 do not. The evidence refuting EU under risk should also pose uncomfortable questions for some of the neoclassical models of ambiguity. The empirical evidence for neoclassical models of ambiguity aversion is at best mixed, hence, our discussion of these models is brief. On the other hand, while the behavioral models are extremely promising, these developments are at a relatively early stage. In order to maintain a common

theme across the chapter, we focus on the ability of alternative models to resolve the Ellsberg paradox.

Some of the neoclassical models that we consider are *Choquet expected utility*, *multiple priors models*, *two-stage recursive models* that do not require the reduction of compound lotteries, and models of *second order beliefs* such as the *smooth ambiguity aversion model*. Among the behavioral models, we consider *support theory* and *source-based uncertainty*. Promising new empirical evidence suggests that prospect theory outperforms most of the standard neoclassical models but more empirical testing and replication is required before we could assert this with more confidence. This strengthens the case that prospect theory applies equally well to risk, uncertainty, and ambiguity, a conclusion that is well represented in the title of Peter P. Wakker's book *Prospect Theory: For Risk and Ambiguity*.

CHAPTER 1

The Evidence on Human Choice under Risk and Uncertainty

1.1 Introduction

Suppose that a decision maker takes some action that leads to an *objective probability distribution* over a set of possible outcomes. This is a situation of *risk*. In this chapter, we briefly consider the elements of the neoclassical approach to decision making in such a situation, *expected utility theory* (EU), and the empirical evidence for EU. Most readers of this book would be familiar with EU. For this reason, no more than an outline of the basic elements is provided.¹

Section 1.2 outlines the basic framework of EU, with particular emphasis on its preference foundations (von Neumann and Morgenstern, 1944). We also consider *attitudes to risk* that, for EU, are determined entirely by the shape of the utility function. Although most people may believe that probabilities have something to do with risk, they play no role in EU in determining risk attitudes. We also outline the basic definitions of *stochastic dominance* for later use. The axioms that lead to an EU representation of preferences are also known as the *axioms of rationality*. Ultimately the validity of these axioms is an empirical matter.

In contrast to risk, under *uncertainty* no objective probabilities that are universally agreed upon can be identified. Section 1.3 considers the basic model of uncertainty in economics, *subjective expected utility* or SEU (de Finetti, 1937; Savage, 1954). These results form the basis of the *Bayesian approach to uncertainty*.

Section 1.4 considers a topic of considerable applied interest. Namely, the problem of eliciting the utility function of an EU decision maker under risk (McCord and de Neufville, 1986), and uncertainty (Wakker and Deneffe, 1996).

Section 1.5, the heart of the chapter, describes the violations of EU. Section 1.5.1 focuses on the evidence arising from the violation of the *independence axiom*, one of the main axioms of rationality. These violations take two forms, the *common ratio effect*, and the *common consequence effect* (Allais, 1953). The common ratio violation takes place when scaling down the probabilities in two-outcome lotteries by a common factor that leads to a reversal in preferences. The common

¹ The interested reader may wish to consult any advanced undergraduate or graduate text in economics; see, for instance, Varian (1992) and Mas-Collel et al. (1995). A more formal treatment can be found in Fishburn (1982b).

consequence violation takes place when two pairs of lotteries each have a common consequence, yet the decision maker chooses a different lottery in each pair.

Section 1.5.2 explains a particularly useful device to study the implications of violations of the axioms of rationality: the *Marshall–Machina probability triangle*. The independence axiom that lies at the heart of EU implies that the indifference curves in the probability triangle are parallel straight lines. However, one observes both *fanning-in* and *fanning-out* of indifference curves, which violates the independence axiom. Furthermore, empirical evidence suggests that the violations of EU are more pronounced along the edges of the probability triangle as compared to the interior of the triangle. The violations in many cases arise from non-linear weighting of probabilities that is particularly pronounced close to the edges of the probability triangle.

Section 1.5.3 briefly considers alternative theories that explain the evidence in the probability triangle by relaxing the independence axiom; for instance *weighted utility theory* (Chew and MacCrimmon, 1979). Theories that have the best chance of explaining the evidence incorporate both *fanning-in* and *fanning-out* of indifference curves in the probability triangle. An example is the theory of *disappointment aversion* (Gul, 1991) that we consider in more detail in the next chapter.

Section 1.5.4 considers Rabin's (2000a) *calibration theorem*. Consider an EU decision maker who has an increasing and concave utility function. Rabin showed that if this decision maker turns down a moderate stake gamble (say, 50–50 chance of losing 10 and gaining 11), then the decision maker will turn down the following equiprobable gamble: lose 100 and win an infinite amount. Clearly this is absurd, but it is a logical outcome under EU. We shall see in Chapter 3 that this paradox does not arise under prospect theory.

EU assumes that the choice among pairs of mathematically equivalent lotteries does not depend on how the information in the lotteries is framed, also known as the assumption of *description invariance*. Section 1.5.5 considers a classic problem from Tversky and Kahneman (1981) that powerfully illustrates a violation of description invariance. This problem also illustrates that decision makers are risk loving (respectively, risk averse) when they perceive themselves to be in losses (respectively, gains), a finding of considerable importance that we study more closely in the subsequent chapters.

Section 1.5.6 considers the widely documented phenomenon of *preference reversal* (Lichtenstein and Slovic, 1971). Suppose that we have two lotteries: A *P*-bet has a high probability of winning a low prize and a *\$*-bet has a low probability of winning a high prize. When decision makers have to directly choose between the two bets, they prefer the *P*-bet. However, at the same time they exhibit a higher certainty equivalent for the *\$*-bet. We also consider attempts to explain preference reversals.

Section 1.5.7 closes the chapter by demonstrating that the reduction of compound lotteries axiom typically does not hold. This has far reaching implications for the validity of EU and will play an important role in providing one potential explanation of the Ellsberg paradox in Chapter 4.

1.2 The elements of classical decision theory

Let $X = \{x_1, x_2, \dots, x_n\}$ be a fixed, finite, set of real numbers such that $x_1 \leq x_2 \leq \dots \leq x_n$. One may interpret X as the set of all possible levels of wealth of the decision maker. We define a *lottery*, or *gamble*, as

$$L = (x_1, p_1; x_2, p_2; \dots; x_n, p_n), \quad (1.1)$$

where $p_1, p_2, \dots, p_n$, are the respective probabilities corresponding to the outcomes $x_1, x_2, \dots, x_n$, such that $p_i \in [0, 1]$ and $\sum_{i=1}^n p_i = 1$. The lottery $(x_i, 1)$ denotes an outcome x_i received with certainty, so we sometimes simply denote it as x_i . The lottery in (1.1) is sometimes also known as a *simple lottery*. Given two simple lotteries, L_1 and L_2 , a lottery of lotteries, say, $(L_1, p; L_2, 1 - p)$, where $p \in (0, 1)$, is known as a compound lottery; the special case, $(L, p; 0, 1 - p)$ is sometimes denoted by (L, p) .

A situation of *risk* arises when the probabilities are objectively known. In contrast, when probabilities are not objectively known, then we have a situation of *uncertainty*. We focus on risk first.

Let the set of all lotteries of the form (1.1) be denoted by $\mathcal{L}$. Suppose that the decision maker must choose a lottery from some set $\mathcal{L}' \subset \mathcal{L}$. The fundamental question in decision theory is: Which lottery will the decision maker choose? This is also the simplest possible setting. It is “as if” the decision maker is simply playing a game against nature. Nature chooses the set of lotteries and the decision maker simply picks one or more elements from the set. In game theory, by contrast, the decision maker, in addition, has to take account of the behavior of others.

1.2.1 Preference foundations of expected utility theory (EU)

The most widely used theory under risk in economics is *expected utility theory* (EU).² Many behavioral models in decision theory build on the EU model and the lessons learnt from the empirical anomalies, relative to the EU predictions. We first briefly review the foundations of EU.

The decision maker is endowed with a binary preference relation over $\mathcal{L}$, denoted by $\succeq$; it is interpreted as “at least as good as.” So, for $L_1, L_2 \in \mathcal{L}$, $L_2 \succeq L_1$ means that the lottery L_2 is “at least as good as” the lottery L_1 , or L_2 is “weakly preferred” to L_1 . The notations $L_2 \succeq L_1$ and $L_1 \preceq L_2$ are equivalent, just as $5 < 7$ or $7 > 5$, and one may use either of the two; similar comments apply to strict preference, e.g., $L_2 \succ L_1$ and $L_1 \prec L_2$ are equivalent.

We now strive to represent preferences in a more convenient form. Consider the following conditions, or axioms, on preferences. We assume that all lotteries below belong to the set $\mathcal{L}$, and a typical member of this set is given in (1.1).

Axiom 1.1 (Order) *Order requires the following two conditions.*

- a. *Completeness: For all lotteries L_1, L_2 , either $L_2 \succeq L_1$ or $L_2 \preceq L_1$.*
- b. *Transitivity: For all lotteries L_1, L_2, L_3 , $L_3 \succeq L_2$ and $L_2 \succeq L_1 \Rightarrow L_3 \succeq L_1$.*

Two further binary relations between lotteries can be defined in terms of $\preceq$:

$$\left\{ \begin{array}{ll} \text{Indifference:} & L_1 \sim L_2 \Leftrightarrow L_2 \succeq L_1 \text{ and } L_1 \succeq L_2. \\ \text{Strict preference:} & L_2 \succ L_1 \Leftrightarrow \text{it is not the case that } L_1 \succeq L_2. \end{array} \right.$$

Axiom 1.2 (Best and worst): $x_n \succ x_1$ (i.e., $(x_n, 1) \succ (x_1, 1)$).

Axiom 1.3 (Continuity): *For each lottery, L , there is a $p \in [0, 1]$ such that $L \sim (x_1, 1 - p; x_n, p)$.*

Axiom 1.4 (Independence): *For all lotteries L_1, L_2, L , and all $p \in [0, 1]$, $L_2 \succeq L_1 \Leftrightarrow (L_2, p; L, 1 - p) \succeq (L_1, p; L, 1 - p)$.*

² For the preference foundations of EU, the interested reader can consult Kreps (1990), Varian (1992), and Mas-Colell et al. (1995). Kreps (1990) already makes a strong case for exploring alternatives to EU. For a more advanced and definitive treatment of EU, see Fishburn (1982b).

Axiom 1.5 (*Reduction, or law of compounding lotteries*): Let $p_1, p_2, p \in [0, 1]$. Let $L_1 \sim (x_i, 1 - p_1; x_j, p_1)$ and $L_2 \sim (x_i, 1 - p_2; x_j, p_2)$. Then

$$\begin{aligned} (L_1, p; L_2, 1 - p) &\sim ((x_i, 1 - p_1; x_j, p_1), p; (x_i, 1 - p_2; x_j, p_2), 1 - p) \\ &\sim (x_i, (1 - p_1)p + (1 - p_2)(1 - p); x_j, pp_1 + (1 - p)p_2). \end{aligned}$$

From Axiom 1.2, a decision maker strictly prefers receiving the highest outcome with certainty to receiving the lowest outcome with certainty. The continuity property in Axiom 1.3 asserts that every lottery is equivalent to a simple lottery obtained by probabilistically mixing the highest outcome, x_n , and the lowest outcome, x_1 . The independence axiom, Axiom 1.4, is the most crucial. It requires that if a lottery, L_2 , is preferred to a lottery, L_1 , then mixing each with a third lottery, L , (with the same positive mixing probability, p) does not alter the preference.

“Rationality” has a precise meaning in decision theory that is given in the next definition.³

Definition 1.1 (*Axioms of rationality*): The following axioms: Order, best and worst, continuity, independence, and reduction, collectively, are termed as the axioms of rationality. If a binary relation $\succeq$ on $\mathcal{L}$ satisfies the axioms of rationality, then it is called a rational ordering of $\mathcal{L}$.

Definition 1.2 (*Representation*): Let $\succeq$ be a binary relation on $\mathcal{L}$ and let $U : \mathcal{L} \rightarrow \mathbb{R}$ such that for all $L_1, L_2 \in \mathcal{L}$, $L_2 \succeq L_1$ if, and only if, $U(L_2) \geq U(L_1)$. Then U is said to represent $\succeq$ and we say that $\succeq$ is induced by U .

Definition 1.3 (*von Neumann–Morgenstern expected utility function*): A utility function $U : \mathcal{L} \rightarrow \mathbb{R}$ has the “expected utility form” (or the von Neumann–Morgenstern form) if we can assign a real number, $u(x_i)$, to each possible outcome, x_i , such that for the lottery L in (1.1), we can write

$$U(L) = \sum_{i=1}^n p_i u(x_i). \quad (1.2)$$

We now state the main result, the existence of a von Neumann–Morgenstern expected utility representation.

Proposition 1.1 (*von Neumann and Morgenstern, 1944*): Suppose that the binary relation, $\succeq$, on the set of lotteries $\mathcal{L}$, satisfies the axioms of rationality. Then there is a function, $U : \mathcal{L} \rightarrow \mathbb{R}$, such that, for all $L_1, L_2, L \in \mathcal{L}$,

- (a) $U(x_1) \leq U(L) \leq U(x_n)$ and $U(x_1) < U(x_n)$,
- (b) $L_2 \succeq L_1$ if, and only if, $U(L_2) \geq U(L_1)$,
- (c) for all $p \in [0, 1]$, the utility function is of the von Neumann–Morgenstern form

Proposition 1.1(a) shows that a constant utility function will not suffice. Proposition 1.1(b) shows that a utility function exists that represents the underlying preferences and Proposition 1.1(c) shows that it takes the familiar form under EU that is linear in probabilities.

Among the set of available lotteries, $\mathcal{L}$, the *EU criterion* requires that the decision maker chooses the one that gives the highest expected utility. The converse of Proposition 1.1 also holds. Namely that if a function U with the properties listed in Proposition 1.1 exists, then $\preceq$ satisfies the axioms of rationality. An important property of the EU functional is that it is

³ Notice that rationality does not imply other literal meanings of the word such as welfare or efficiency maximizing choices or preferences that are self regarding or other regarding, etc.

invariant with respect to positive affine transformations. Thus, the utility function $\widehat{U}(L) = \sum_{i=1}^n p_i [a + bu(x_i)]$, where $a \in \mathbb{R}$ and $b > 0$, represents the same preferences as does U in (1.2).

The axioms of rationality appear intuitive and plausible. A violation of the axioms of rationality is not a form of irrationality or madness. The axioms are no more than assumptions about human behavior, and their reasonableness must be tested by empirical evidence. A large number of rigorous tests have rejected the axioms of rationality. The challenge for behavioral economics is to produce a behavioral decision theory that is, on the one hand, in conformity with the evidence and, on the other hand, sufficiently tractable.

Remark 1.1 *EU has two main features:*

- (i) *Linearity in probabilities. Formally, this means that $EU(L_1, p; L_2, 1 - p) = pEU(L_1) + (1 - p)EU(L_2)$. This assumption is fundamental to many areas in economics. For instance, it is required in the proof of the existence of a mixed strategy Nash equilibrium. However, extensive empirical evidence shows that it is violated. Most non-EU decision theories relax the linearity assumption in some manner.*
- (ii) *In most applications of EU, utility typically depends on the final wealth level in each state of the world. Redefining outcomes under EU relative to a reference outcome (or any outcome) is not inconsistent with the axioms of EU. Doing so would partition the domain of outcomes into gains (outcomes greater than the reference point) and losses (outcomes lower than the reference point). However, EU offers no insights about human behavior in the domain of gains and losses. By contrast, empirically observed behavior of decision makers is fundamentally different in the domain of gains and losses. Indeed, differences in behavior in gains and losses are not only an essential driver of many important results in other decision theories such as prospect theory, but there is empirical confirmation of these predictions.*

1.2.2 Attitudes to risk under EU

The EU criterion contrasts with the earlier *expected value maximization criterion* (EV), which evaluates the lottery $L = (x_1, p_1; x_2, p_2; \dots; x_n, p_n)$ as $EV(L) = \sum_{i=1}^n p_i x_i$.

Consider a decision theory in which the utility function $\overline{U}$ is able to represent preferences over lotteries, e.g., under EU, $\overline{U} = U$. In other cases, $\overline{U}$ could be the utility representation under rank dependent utility or under prospect theory.

Definition 1.4 (*Certainty equivalent*): Consider the lottery $(C_L, 1)$, where $C_L \in X \subset \mathbb{R}$. If a solution exists to the equation $\overline{U}(C_L) = \overline{U}(L)$, then C_L is known as the *certainty equivalent* of the lottery L .

Notice that the definition of a certainty equivalent holds for any decision theory that admits a utility representation of preferences over lotteries, although the value of certainty equivalent under alternative decision theories is likely to differ. We now state the meaning of risk aversion; an individual who would rather accept a sure amount that is lower than the expected value of the lottery, in order to avoid playing the lottery, is said to be risk averse.

Definition 1.5 (*Risk aversion*): A decision maker is “risk neutral” if $C_L = EV(L)$, “risk averse” if $C_L < EV(L)$ and “risk loving” if $C_L > EV(L)$.

In this section, we only consider an expected utility representation of preferences ($\overline{U} = U$). The next proposition states a well-known result that you are asked to prove in one of the exercises.

Proposition 1.2 *Under expected utility theory, the properties of risk neutrality, risk aversion, and risk seeking are equivalent, respectively, to a linear, concave, and convex utility function.*

Risk aversion arising from a concave utility function underpins many important economic phenomena. For instance, the study of insurance, the design of contract schemes in principal-agent problems, the optimal rates of taxes, and determination of wage rates.

The EV criterion restricts decision makers to be risk neutral, hence, lotteries that have the same EV but appear different in their risk characteristics are equally preferred by the EV criterion. Daniel Bernoulli showed in 1738 that a famous paradox, the *St. Petersburg paradox*, cannot be resolved by the EV criterion but can be resolved by the EU criterion (you are asked to prove this in one of the exercises). The EU criterion can solve the St. Petersburg paradox because it allows for risk aversion. However, risk preferences are incorporated in EU through the shape of the utility function alone, rather than through probabilities. By contrast, many psychologists would typically argue that risk should be incorporated through probabilities rather than the shape of the utility function.

We can also compare the desirability of two lotteries that have the same outcomes but different probability distributions over them. Consider lotteries of the form $L = (x_1, p_1; x_2, p_2; \dots; x_n, p_n)$, where the outcomes $x_1 < x_2 < \dots < x_n$ belong to the set of real numbers. Under any distribution function P over this set of incomes, the probability of outcome x_i is denoted by $p_i \geq 0$, $\sum_{i=1}^n p_i = 1$. The *cumulative probability distribution* is given by $P_0 = 0$, $P_j = \sum_{i=1}^j p_i$. Denote by Π , the set of all such distributions over X . The *cumulative of the cumulative distribution* is given by $\tilde{P}_0 = 0$, $\tilde{P}_j = \sum_{i=1}^j P_i$, $j = 1, 2, \dots, n$. Let $\tilde{\Pi}$ be the set of all such distributions. The average, or mean, of $x_1, x_2, \dots, x_n$, under the distribution $P \in \Pi$ is $\mu_P = \sum_{i=1}^n p_i x_i$. Let u be the utility function for sure outcomes $(x_i, 1)$.

Notation 1.1 *Denote by $\mathbf{u}$, the class of all non-decreasing utility functions.*

We state here the basic results on *first order stochastic dominance* and *second order stochastic dominance* in the case of discrete outcomes; the continuous analogues can be picked up from any advanced microeconomics book. We shall need the results for the discrete case to prove several Propositions in the book.

Definition 1.6 (*First order stochastic dominance*): Let $P, Q \in \Pi$. Then P first order stochastically dominates Q ($P \succeq_1 Q$) if $P_j \leq Q_j$ for $j = 1, 2, \dots, n$. If, in addition, the inequality is strict for some j , then P strictly first order stochastically dominates Q ($P \succ_1 Q$).

Definition 1.7 We say that the outcome levels are equally spaced if for some positive real number, δ , $x_{i+1} - x_i = \delta$ for $i = 1, 2, \dots, n - 1$.⁴

Definition 1.8 (*Second order stochastic dominance*): Suppose that incomes are equally spaced (Definition 1.12).⁵ Let $P, Q \in \Pi$. Then P second order stochastically dominates Q

⁴ This is not restrictive because we can always introduce extra income levels, each with probability zero, to achieve equal spacing. For example, the incomes $y_1 = 5, y_2 = 7, y_3 = 11$, are not equally spaced. Let $P_1 = \frac{1}{3}$, $P_2 = \frac{2}{3}$, and $P_3 = 1$. Consider $y_1 = 5, y_2 = 7, y_3 = 9, y_4 = 11$, which are equally spaced, and $Q_1 = \frac{1}{3}$, $Q_2 = \frac{2}{3}$, $Q_3 = \frac{2}{3}$, and $Q_4 = 1$. Both P and Q describe the same reality.

⁵ Often "equal spacing" is not stated as part of the definition of second order stochastic dominance. But then it has to be stated in the Propositions. We find it more convenient to state it in the definition. The equal spacing assumption is not needed in the continuous case.

$(P \succeq_2 Q)$ if $\tilde{P}_j \leq \tilde{Q}_j$ for $j = 1, 2, \dots, n$ and if $\mu_P \geq \mu_Q$.⁶ If, in addition, one of these inequalities is strict, then P strictly second order stochastically dominates Q ($P \succ_2 Q$).

So far, the definitions of first order and second order dominance are purely statistical. We now explore their economic content. The results should be familiar to most advanced undergraduates in economics, so the proofs are omitted; readers interested in the continuous case can consult Mas-Collel et al. (1995).

Proposition 1.3 *Let $P, Q \in \Pi$. Then $P \succeq_1 Q$ if, and only if, for any $u \in \mathbf{u}$,*

$$\sum_{i=1}^{i=n} p_i u(x_i) \geq \sum_{i=1}^{i=n} q_i u(x_i), \quad (1.3)$$

and $P \succ_1 Q$ if, and only if, the inequality (1.3) is strict for all strictly increasing $u \in \mathbf{u}$.

Note that Proposition 1.18 holds for all $u \in \mathbf{u}$. Choosing the non-decreasing function $u(x) = x$ as a candidate function in Proposition 1.18, we get that if $P \succeq_1 Q$ ($P \succ_1 Q$), then $\mu_P \geq \mu_Q$ ($\mu_P > \mu_Q$).

Proposition 1.4 *Let $P, Q \in \Pi$. Then $P \succeq_2 Q$ if, and only if, for any concave $u \in \mathbf{u}$,*

$$\sum_{i=1}^{i=n} p_i u(x_i) \geq \sum_{i=1}^{i=n} q_i u(x_i), \quad (1.4)$$

and $P \succ_2 Q$ if, and only if, the inequality (1.4) is strict for all strictly increasing strictly concave $u \in \mathbf{u}$.

It is also possible to give an alternative definition of first order stochastic dominance with a discrete set of outcomes that holds fixed the probability distribution but alters the spread of the outcomes in X . Let $L_2 \succeq_1 L_1$ and $L_2 \succ_1 L_1$ denote that the lottery L_2 first order stochastically dominates L_1 , respectively, in a weak and strict sense.

Definition 1.9 *Let $L_1 = (x_1, p_1; \dots; x_n, p_n)$ and $L_2 = (y_1, p_1; \dots; y_n, p_n)$, then*

$$\begin{aligned} (a) \quad & x_i \leq y_i, i = 1, 2, \dots, n \Rightarrow L_2 \succeq_1 L_1, \\ (b) \quad & x_i \leq y_i, i = 1, 2, \dots, n \text{ and } x_i < y_i \text{ for some } i \Rightarrow L_2 \succ_1 L_1. \end{aligned} \quad (1.5)$$

Some additional properties of EU are taken up in the exercises at the end of Part 1.

1.3 Subjective expected utility theory (SEU)

Let there be n mutually exclusive events $E_1, E_2, \dots, E_n$ that form a partition of some relevant sample space, S . The objective probabilities of these events are not known. The decision maker assigns *subjective probabilities* $\mu(E_1), \mu(E_2), \dots, \mu(E_n)$ to these events such that $\mu(E_i) \geq 0$ for all i and $\sum_{i=1}^n \mu(E_i) = 1$. This is a situation of *uncertainty* because, although objective probabilities are not known, well-defined and consistent subjective probabilities exist. Now suppose that the individual needs to evaluate a special kind of lottery of the form

⁶ Often the condition $\mu_P \geq \mu_Q$ is not included as part of the definition of second order stochastic dominance. But then it needs to be included in the statement of the relevant Propositions. It is slightly simpler to include $\mu_P \geq \mu_Q$ as part of the definition of second order stochastic dominance.

$$f = (E_1, x_1; E_2, x_2; \dots; E_n, x_n), \quad (1.6)$$

where the decision maker receives the outcome x_i if the event E_i occurs, $i = 1, 2, \dots, n$. For pedagogical simplicity, let the outcomes be monetary, such as income or wealth. The object in (1.6) is also known as a *Savage act*. For any partition of a sample space, S , into a set of *mutually exclusive* events $\{E_1, E_2, \dots, E_n\}$, a Savage act, f , maps states of the world to outcomes, i.e., $f : S \rightarrow X$; we use the notation $f(E_i) = x_i$, $i = 1, \dots, n$.

One of the earliest justifications for the use of subjective utility was given in the work of de Finetti (1937). Denote the set of all Savage acts of the form in (1.6) by F , and define a binary relation over F by $\succeq$ which denotes “at least as good as.” The addition operator, $+$, applied to elements of the set F implies state by state addition of outcomes. For instance, let $f_1, f_2, f_3 \in F$:

$$f_1 = (E_1, x_1; E_2, x_2), f_2 = (E_1, y_1; E_2, y_2), f_3 = (E_1, z_1; E_2, z_2) \quad (1.7)$$

then,

$$f_1 + f_3 = (E_1, x_1 + z_1; E_2, x_2 + z_2), f_2 + f_3 = (E_1, y_1 + z_1; E_2, y_2 + z_2). \quad (1.8)$$

We now define the analogue of the independence axiom of EU under uncertainty, which is termed as *additivity*.

Definition 1.10

- (i) (*Additivity*): Suppose that we have Savage acts $f_1, f_2, f_3 \in F$ such that $f_1 \succeq f_2$. Then additivity implies that $f_1 + f_3 \succeq f_2 + f_3$.
- (ii) (*Monotonicity*): Let $f_1 = (E_1, x_1; E_2, x_2; \dots; E_n, x_n)$, $f_2 = (E_1, y_1; E_2, y_2; \dots; E_n, y_n)$. Suppose that $x_i \geq y_i$ for $i = 1, 2, \dots, n$. Then monotonicity implies that $f_1 \succeq f_2$.

Suppose that in addition to additivity, $\succeq$ is a weak order (transitive and complete), and *monotonicity* holds. Let f be as defined in (1.6) and $g = (E_1, y_1; E_2, y_2; \dots; E_n, y_n)$, where $f, g \in F$. Then (under a further continuity assumption) de Finetti (1937) obtained the following remarkable result (for details, see Wakker (2010) and Gilboa (2009)).

$$f \succeq g \Leftrightarrow \sum_{i=1}^n \mu(E_i) x_i \geq \sum_{i=1}^n \mu(E_i) y_i.$$

Thus, the individual acts so as to maximize the *expected value of a lottery*. Furthermore, the numbers $\mu(E_1), \mu(E_2), \dots, \mu(E_n)$, which may be interpreted as *subjective probabilities*, are *unique*.

In many contexts, expected value maximization might not be very attractive.⁷ Despite the drawbacks of the expected value criterion, de Finetti’s (1937) result provided the basis and initial impetus for using subjective probabilities in a formal manner. For small stake experiments, expected value might not necessarily be a bad approximation.

Let $f(E_i), g(E_i)$ denote the respective outcomes under the Savage acts f, g , when the event E_i occurs. Define a function $u : S \rightarrow R$ that it is unique up to a positive affine transformation;

⁷ Consider the binary Savage acts, f_1, f_2, f_3 in (1.7) and their addition, defined in (1.8). Let the individual be declared bankrupt if his income in any state falls below some monetary level b . Suppose that all monetary outcomes in (1.7), (1.8) are greater than b except for $x_1 + z_1$, which is below b . Then the individual might express the preference $f_2 + f_3 \succsim f_1 + f_3$ which violates additivity.

$u(f(E_i))$ has the interpretation of the utility of the outcome $f(E_i)$. The following functional has the interpretation of *subjective expected utility* from the Savage act f .

$$V(f) = \sum_{i=1}^n \mu(E_i)u(f(E_i)), \quad (1.9)$$

where $\mu(E_i) \geq 0 \forall i$ and $\sum_{i=1}^n \mu(E_i) = 1$. This is the sense in which EU can be plausibly extended to subjective beliefs. Risk can now be seen as a special case of uncertainty when there exists an objective probability p_i such that $p_i = \mu(E_i)$ for $i = 1, \dots, n$. When new information arrives, the decision maker engages in Bayesian updating to update his prior beliefs into a set of posterior beliefs that comprise the new, unique, additive probability measure.

In one of the most brilliant results in all of social science, Savage (1954) derived the following representation of preferences of the individual. For any two lotteries, f, g :

$$f \succeq g \Leftrightarrow \sum_{i=1}^n \mu(E_i)u(f(E_i)) \geq \sum_{i=1}^n \mu(E_i)u(g(E_i)), \quad (1.10)$$

where μ is a unique additive probability measure such that $\mu(E_i) \geq 0$ for all i and $\sum_{i=1}^n \mu(E_i) = 1$.⁸ Savage (1954) provides an axiomatic foundation of *subjective expected utility* (SEU), also sometimes known as the *Bayesian approach to uncertainty* because beliefs are updated in a Bayesian manner.⁹ In addition to the standard axioms of completeness and transitivity, the substantive new axiom in the Savage framework is the *sure thing principle* that is the analogue of the independence axiom under EU.¹⁰

Definition 1.11 (*Sure thing principle*): Consider the following pairwise choices between four Savage acts, $f_1, f_2, f_3, f_4 \in F$ that are defined over the events E_1, E_2, E_3 ; let y, x_1, x_2, x_3, x_4 be the possible outcomes.

$$\begin{aligned} f_1 &= (E_1, y; E_2, x_2; E_3, x_3), f_2 = (E_1, x_1; E_2, x_2; E_3, x_3), \\ f_3 &= (E_1, y; E_2, x_2; E_3, x_4), f_4 = (E_1, x_1; E_2, x_2; E_3, x_4). \end{aligned}$$

The sure thing principle holds if $f_1 \succ f_2 \Leftrightarrow f_3 \succ f_4$.

The remarkable thing about Savage's axiomatic framework is that the axioms make no assumptions about objects such as probabilities. However, the actual representation of preferences in (1.10) is similar to that under EU except that the (unique) objects $\mu(E_i)$ have the interpretation of subjective probabilities.

We now define *probabilistic sophistication*, a term that we require in Chapter 4.

Definition 1.12 Lotteries or Savage acts, as defined in (1.6), are known as *event contingent lotteries*.¹¹ By contrast, lotteries written in the standard form, $L = (x_1, p_1; \dots; x_n, p_n)$, such that $p_i \geq 0$ and $\sum_{i=1}^n p_i = 1$, are known as *probability contingent lotteries*.

Suppose that the decision maker assigns the probability p_i to the event E_i . $i = 1, \dots, n$. Then the decision maker may consider the lotteries $(x_1, p_1; \dots; x_n, p_n)$ and $(E_1, x_1; E_2, x_2; \dots; E_n, x_n)$

⁸ μ is an additive probability measure if $\mu(\cup_{i=1}^n E_i) = \sum_{i=1}^n \mu(E_i)$.

⁹ While Savage's preference foundations have received the most attention, there are others, such as Ramsey (1931), de Finetti (1937), and Anscombe and Aumann (1963).

¹⁰ There are two further axioms that allow one to make a distinction between beliefs of the individual and the individual's tastes. For a discussion of these axioms, see Karni (2003) and Gintis (2009, Section 1.5).

¹¹ One could give a more formal measure theoretic construction if required, but our discussion here is at a heuristic level only.

to be identical and treat them interchangeably; this is known as probabilistic sophistication and summarized in the next definition.

Definition 1.13 (*Probabilistic sophistication*): Consider Definition 1.12. The decision maker is said to be probabilistically sophisticated when event contingent lotteries can be replaced by probability contingent lotteries. The subjective probabilities in probability contingent lotteries constitute a probability measure. However, there is no presumption, under probabilistic sophistication, that in comparing two lotteries, the decision maker uses SEU. The decision maker is allowed to follow any of the non-expected utility theories. Furthermore, the decision maker is indifferent between two event contingent lotteries that lead to the same probability contingent prospect.¹²

In Chapter 4 we will show an example in which the *Ellsberg paradox* implies a violation of probabilistic sophistication.

1.4 Eliciting the utility function under EU

In this section, we consider the elicitation of the underlying utility function when the decision maker follows EU. We consider the two cases of known and unknown probabilities.

1.4.1 The case of known probabilities

Suppose that probabilities and outcomes are known. Consider any two outcomes such that $x_l < x_h$. Then, without loss of generality, we can choose the following convenient normalization.

$$u(x_l) = 0, u(x_h) = 1. \quad (1.11)$$

Example 1.1 Suppose that x_l, x_h are, respectively, the lowest and highest monetary outcomes. Suppose that there is a utility function v such that $v(x_l)$ and $v(x_h)$ are known real numbers. We wish to renormalize the utility function such that the utility from the lowest outcome is 0 and the utility from the highest outcome is 1. Recall that if v represents the decision maker's preferences under EU, then so does $u = a + bv$, $a \in \mathbb{R}$ and $b > 0$ (invariance to positive affine transformations). Solving the two simultaneous equations $u(x_l) = 0 = a + bv(x_l)$ and $u(x_h) = 1 = a + bv(x_h)$ gives the required values of $a = -bv(x_l)$ and $b = \frac{1}{v(x_h) - v(x_l)}$.

We now offer the decision maker, lotteries of the form

$$L = (x_l, 1 - p; x_h, p), \quad (1.12)$$

and compute the certainty equivalent, x , of the lottery. Under EU, and using (1.11), this is given by the solution to

$$u(x) = (1 - p)u(x_l) + pu(x_h) = p. \quad (1.13)$$

¹² The axiomatic justification of probabilistic sophistication can be found in Grant et al. (2008), Chew and Sagi (2006), Abdellaoui and Wakker (2005), Sarin and Wakker (2000), Epstein and LeBreton (1993), and Machina and Schmeidler (1992).

From (1.13) we can find the utility of the decision maker at the certainty equivalent as simply the known probability of the best outcome. The requirement that a solution exists to (1.13) is called *standard gamble solvability*. This procedure for eliciting the utility function is also known as the problem of eliciting standard gamble probability (SG) with respect to the outcomes x_l, x_h . Suppose that the decision maker's preferences are a *weak order* (complete and transitive) and the following condition, known as *standard gamble dominance*, holds:¹³

$$(x_l, 1 - p; x_h, p) \succ (x_l, 1 - q; x_h, q) \text{ for } p > q.$$

Then it can be shown that p in (1.13) is unique.¹⁴ When p is unique, then keeping fixed x_l and x_h , we can vary x , generating a corresponding series of values of p that reveal the decision maker's utility function.

Under SG, as shown in (1.13), the decision maker compares a risky choice (the RHS of the first equality) with a riskless choice (LHS of the first equality). McCord and de Neufville (1986) have criticized this approach and argued for eliciting the utility function from more general risky choices. In the SG method, for instance, if individuals are not really EU maximizers and evaluate probabilities non-linearly, then because high probabilities are underweighted, forcing probabilities to be linear will overestimate the concavity of the utility function.¹⁵ Furthermore, a strong *certainty effect* could magnify the problems created on account of the riskless choice in the SG method.¹⁶ In order to control for the certainty effect, McCord and de Neufville (1986) propose an improvement in the SG method by calculating the value of x such that

$$(x, \lambda; a, 1 - \lambda) \sim ((x_l, 1 - p; x_h, p), \lambda; a, 1 - \lambda), \quad (1.14)$$

where $a \in \mathbb{R}$ and $\lambda \in (0, 1)$. The indifference in (1.14) is arrived at by asking a series of questions from subjects in experiments that pit the two lotteries against each other and using the tipping point where the preferences flip from one direction to the other, as the point of indifference. This remark applies to all the indifference elicitation here and in Section 2.5.

For an EU maximizer, the indifference in (1.14) implies that

$$\lambda u(x) + (1 - \lambda) u(a) = \lambda [(1 - p)u(x_l) + pu(x_h)] + (1 - \lambda) u(a). \quad (1.15)$$

Using (1.11), we get from (1.15) that $u(x) = p$, which is identical to (1.13). When one applies EU, one expects the SG method and the method of McCord and de Neufville (1986) to give identical results. This is called *standard gamble equivalence*. In practice, however, the results differ, possibly because individuals in fact use non-linear weighting of probabilities.¹⁷

¹³ For $x_l = x_1$ and $x_h = x_n$ standard gamble dominance is equivalent to monotonicity.

¹⁴ See exercise 2.6.3 in Wakker (2010).

¹⁵ For instance, it has been shown that EU is violated for health outcomes (Rutten-van Mölken et al., 1995; Stalmeier and Bezembinder, 1999). One would then expect that measurements of utility of health outcomes based on the SG method would lead to biased estimates of utility. This was shown by Bleichrodt et al. (2007).

¹⁶ The certainty effect refers to relatively greater sensitivity to changes in probability around zero as one approaches the certainty of an outcome.

¹⁷ It can be shown (See Theorem 2.6.3 in Wakker, 2010, p. 62) that EU is equivalent to $\succeq$ being a weak order in the presence of the following conditions (defined above in this subsection). (1) Standard gamble solvability. (2) Standard gamble dominance. (3) Standard gamble equivalence.

1.4.2 The case of unknown probabilities

Now suppose that probabilities are unknown, i.e., they are not objectively given. This is a more difficult problem because both probabilities and the utility function are unknown. We now give a simple method of eliciting utilities and probabilities in this case.¹⁸ The basis of this elicitation is the *trade-off method* that is given in Wakker and Deneffe (1996); for the relevant preference foundation, see Section 2.6.1. This section will also help us to address the problem of eliciting probability weights and utilities under rank dependent utility and under cumulative prospect theory.

Suppose that there are two possible events E_1 and E_2 with unknown objective probabilities. The decision maker, however, forms subjective beliefs $\mu(E_1) \geq 0$ and $\mu(E_2) \geq 0$ such that $\mu(E_1) + \mu(E_2) = 1$, and evaluates the lottery $(E_1, x_1; E_2, x_2)$ as in *subjective expected utility* (SEU) (see Section 1.3), i.e.,

$$\mu(E_1)u(x_1) + \mu(E_2)u(x_2). \quad (1.16)$$

This form of evaluation can be extended to finitely many outcomes.

Suppose that an initial monetary reward y_0 is given. Consider the following sequence of lotteries that elicit a value $y_j, j = 1, 2, 3, 4$, in order to ensure the indifference relation

$$(E_1, y_j; E_2, a) \sim (E_1, y_{j-1}; E_2, b), \quad (1.17)$$

where $a, b \in$ and E_1, E_2 are given. Thus, y_1, y_2, y_3, y_4 are to be elicited using (1.17), over four iterations.

Using (1.16), (1.17) we get equalities of the form:

$$\mu(E_1)u(y_j) + \mu(E_2)u(a) = \mu(E_1)u(y_{j-1}) + \mu(E_2)u(b); j = 1, 2, 3, 4. \quad (1.18)$$

On rearranging, we get

$$\mu(E_1) [u(y_j) - u(y_{j-1})] = \mu(E_2) [u(b) - u(a)]; j = 1, 2, 3, 4. \quad (1.19)$$

Using the fact that the RHS in (1.19) is identical for all four iterations, we get the following system of equations.

$$u(y_4) - u(y_3) = u(y_3) - u(y_2) = u(y_2) - u(y_1) = u(y_1) - u(y_0). \quad (1.20)$$

In (1.20), we have eliminated the unknown probabilities. As in Example 1.1, normalize so that $u(y_0) = 0$ and $u(y_4) = 1$. This gives rise to the following system of three equations in three unknowns.

$$\begin{aligned} 1 - u(y_3) &= u(y_3) - u(y_2) \\ u(y_3) - u(y_2) &= u(y_2) - u(y_1) \\ u(y_2) - u(y_1) &= u(y_1) \end{aligned}$$

¹⁸ The interested reader can consult Wakker (2010, Ch.4) for a more detailed discussion.

Solving these equations simultaneously we get,

$$u(y_0) = 0; u(y_1) = \frac{1}{4}; u(y_2) = \frac{1}{2}; u(y_3) = \frac{3}{4}; u(y_4) = 1.$$

Recall that only y_0 is given, while y_1, y_2, y_3, y_4 are elicited from the choices made by the decision maker. By generating sufficiently many points such as these, one can recover the underlying utility function. By extending the number of iterations, and adjusting the value of y_0 , one can extend the domain over which the utility function is elicited.

Once the utility function is elicited, then the subjective probabilities of an event E in the sample space can be found. Suppose that the decision maker is asked for a value of x such that the following indifference relation holds.

$$x \sim (E, y; E^c, 0); i = 1, 2, \quad (1.21)$$

where E^c is the complement of E . Under SEU, the indifference relation in (1.21) implies that $u(x) = \mu(E)u(y)$ (because $u(0) = 0$). Since the utility function has already been elicited, $\mu(E) = u(x)/u(y)$ is the required answer. In actual practice, there are problems, however. Consistency might be violated (i.e., subjective probabilities might not add up to 1) and the utility function might not have been fitted (in the first stage) corresponding to the outcome x .

Another method for finding the relevant subjective probabilities is the *method of matching probabilities*. Here, the decision maker is asked to state the probability such that

$$(E, 1; E^c, 0) \sim (1, p; 0, 1 - p). \quad (1.22)$$

Notice the difference in the representation of the two lotteries under uncertainty (LHS) and certainty (RHS). In (1.22), p is said to be the matching probability. If EU holds for the lottery on the RHS, then $\mu(E) = p$.

1.5 Violations of expected utility theory

In this section we show that a great deal of evidence is inconsistent with the predictions of EU; indeed, it would not be unfair to say that EU is rejected by the evidence.¹⁹

1.5.1 Violations of the independence axiom

We first show a violation of the independence axiom, originally due to Allais (1953). Example 1.2, below, draws on Kahneman and Tversky (1979) who showed that the Allais paradox can be explained by assuming that decision makers overweight low probabilities and underweight high probabilities. In contrast, EU is linear in probabilities (Remark 1.1).

¹⁹ Strictly speaking, a *refutation* is an empirical finding that rejects any of the axioms of EU. A *violation* of EU occurs when the empirical evidence rejects an auxiliary assumption of EU. We use these terms interchangeably in this section. Readers can consult the excellent discussions in Kahneman and Tversky (1979), Kahneman and Tversky (2000), Starmer (2000), and Wakker (2010).

Example 1.2 (Allais paradox as a “common ratio violation.” Problems 3 and 4 in Kahneman and Tversky, 1979) Consider a decision maker with initial wealth, x_0 , and utility function u . Outcomes are in units of Israeli pounds.

Problem 3: (95 subjects) Choose between the two lotteries

$$b_1 = (x_0, 0; x_0 + 3000, 1) \text{ and } b_2 = (x_0, 0.2; x_0 + 4000, 0.8),$$

80% chose b_1 , while 20% chose b_2 .

Problem 4: (95 subjects) Choose between the two lotteries

$$b_3 = (x_0, 0.75; x_0 + 3000, 0.25) \text{ and } b_4 = (x_0, 0.8; x_0 + 4000, 0.2),$$

35% chose b_3 , while 65% chose b_4 .

Denoting by $\prec_{EU}$ the binary EU preference relation, the prediction of EU is,

$$b_1 \prec_{EU} b_2 \Leftrightarrow b_3 \prec_{EU} b_4. \quad (1.23)$$

The proof of (1.23) runs as follows. Since EU preferences are invariant to a positive affine transformation, we can assume $u(x_0) = 0$. Hence, the expected utility of the 2 pairs of lotteries in each problem is:

$$EU(b_1) = u(x_0 + 3000), EU(b_2) = 0.8u(x_0 + 4000). \quad (1.24)$$

$$EU(b_3) = 0.25u(x_0 + 3000), EU(b_4) = 0.2u(x_0 + 4000). \quad (1.25)$$

From (1.24) we get that

$$b_1 \prec_{EU} b_2 \Leftrightarrow u(x_0 + 3000) < 0.8u(x_0 + 4000).$$

Dividing both sides by 4, we get

$$0.25U(x_0 + 3000) < 0.2U(x_0 + 4000) \Leftrightarrow b_3 \prec_{EU} b_4.$$

Hence, $b_1 \prec_{EU} b_2 \Rightarrow b_3 \prec_{EU} b_4$. A similar argument shows that $b_3 \prec_{EU} b_4 \Rightarrow b_1 \prec_{EU} b_2$. Hence, (1.23) holds. However, since a majority of subjects chose b_1 over b_2 and b_4 over b_3 , at least some decision makers must have violated (1.23) and, hence, EU.

In Allais's original presentation to Leonard Savage, Paul Samuelson, and others in Paris in 1952 (in units of millions of 1952 Francs), $b_1 = (100, 1)$, $b_2 = (500, 0.98)$, $b_3 = (100; 0.01)$ and $b_4 = (4000; 0.0098)$. To quote from Allais: “The author of this memoir, who is perfectly familiar with Samuelson's arguments, unhesitatingly prefers lottery $[b_1]$ to lottery $[b_2]$ and lottery $[b_4]$ to lottery $[b_3]$ and he thinks that most of his readers will share his preference; yet to the best of his knowledge, he believes that he is not irrational. He is perfectly aware that 2 chances in 100,000 is a non-negligible quantity, but his view is that this quantity does not offset for him the reduction in the possible gain from 500 to 100 million, whereas for him by contrast, the achievement of certainty by raising the chance of winning from 98% to 100% is well worth this reduction” (Allais and Hagen 1979, p. 102).

Savage (like Allais) initially chose b_1 over b_2 and b_4 over b_3 . However, after an EU calculation, Savage concluded that he had made a mistake. Savage's response was that EU is needed to guide us to correct decision making. Of course, both positions may be correct. Savage may be right to conclude that decision makers should be trained in EU. Allais may be right to conclude that, in fact, most decision makers do not follow EU.

The Allais paradox can be explained as follows. In problem 4, the decision maker codes the probability 0.2 as close to 0.25 (and 0.75 as close to 0.80). Hence, he prefers b_4 to b_3 because it

gives him a higher outcome (4000 instead of 3000). By contrast, in problem 3, the decision maker does not code the probability 0.8 as close to 1; the certain outcome is particularly salient (this is an example of the *certainty effect*). In other words, he *underweights* the probability 0.8. Hence, *non-linear weighting of probabilities* is a potential explanation of the Allais paradox but EU does not allow it. However, it is not simply a matter of arbitrarily using non-linear transformations of probabilities. We show below that unless carefully applied, non-linear weighting can also allow stochastically dominated options to be chosen.

When the Allais paradox first appeared, it was thought to be an isolated example for a long time. However, it was realized later that the Allais paradox was part of a family of paradoxes that took two forms, the *common ratio violation* form, and the *common consequence violation* form. Both forms violate the independence axiom of EU (Axiom 1.4) and have been widely replicated. There is some evidence that professional traders also exhibit Allais-type violations but to a lesser extent than students, who are the main subject pool of these experiments; see List and Haigh (2005).

Definition 1.14 (*Common ratio violation*) Consider two probabilities $p, P \in (0, 1]$ and two prizes z, Z ; typically, $p < P$ and $z < Z$. Let $\mu \in [0, 1]$. Recall our convention that (x, p) denotes the lottery $(x, p; 0, 1 - p)$. Consider a choice between the following pairs of lotteries.

$$b_1 = (z, P) \text{ and } b_2 = (Z, p), \quad (1.26)$$

$$b_3 = (z, \mu P) \text{ and } b_4 = (Z, \mu p). \quad (1.27)$$

The “common ratio violation” occurs if for $\mu > 0$ a decision maker prefers b_1 over b_2 and b_4 over b_3 (or b_2 over b_1 and b_3 over b_4). The name “common ratio” derives from the fact that the ratio of probabilities for the outcomes in each of the two choices is the same, i.e., P/p .

Example 1.2, above, is a special case of this violation if we assume that the returns on the asset are not integrated with existing wealth level x_0 . Empirical evidence is consistent with non-integration of assets or, at best, partial integration (Andersen et al., 2012). Fafchamps et al. (2014) find no evidence that subjects in experiments integrate asset returns in the lab with their real-life income/wealth, although winnings in one round may influence subsequent risk attitudes. One explanation is that subjects create different *mental accounts* for the lab and non-lab earnings; see Chapter 20 for mental accounting.

Assuming non-integration and using the notation of Definition 1.14, in Example 1.2 $p = 0.8$, $P = 1$, $z = 3000$, $Z = 4000$, and $\mu = 1/4$. For other experimental tests of the common ratio violation, see, for instance, Chew and Waller (1986) and MacCrimmon and Larsson (1979). Blavatskyy (2010) reports a reverse common ratio violation effect that takes place when the probability p in Definition 1.14 drops below 0.5. The pattern of violation of EU that is observed is then the reverse of the normal one that we observe with the Allais paradox.²⁰

We already know, from Example 1.2, that an observation of a common ratio violation is a refutation of EU. We now consider the question of which axiom (or axioms) might be responsible for the refutation.

Proposition 1.5 *The common ratio violation is a violation of the independence axiom, or the reduction axiom, or of both axioms.*

²⁰ Blavatskyy (2011) suggests an explanation based on the *stochastic expected utility theory* of Blavatskyy (2007) and the *perceived relative argument model* of Loomes (2008).

Proof: Assume that $b_1 \succ b_2$, so $(z, P) \succ (Z, p)$. Hence, $(z, P) \succeq (Z, p)$. From the independence axiom (Axiom 1.4) we get that $((z, P), \mu; 0, 1 - \mu) \succeq ((Z, p), \mu; 0, 1 - \mu)$, which, given our notation, can be denoted as $((z, P), \mu) \succeq ((Z, p), \mu)$. From the reduction axiom (Axiom 1.5) it follows that $((z, P), \mu)$ and $((Z, p), \mu)$ are equivalent to $(z, \mu P)$ and $(Z, \mu p)$, respectively. Hence, $(z, \mu P) \succeq (Z, \mu p)$, i.e., $b_3 \succeq b_4$. Thus, the observation of a common ratio violation, namely, $b_1 \succ b_2$ and $b_4 \succ b_3$ is a rejection of the independence axiom, given the other axioms of EU. Alternatively, we may view the observation of a common ratio violation as a rejection of the reduction axiom, given the other axioms of EU, an observation that is missed by many texts in microeconomics. ■

We now move on to another generic class of paradoxes, associated with the *common consequence violation*, which was also given by Allais (1953).

Definition 1.15 (*Common consequence violation*) Let A, B, C, D be lotteries and $\mu \in [0, 1]$. Suppose that we have the following four compound lotteries, a_1, a_2 and a_3, a_4 , that are compared in a pairwise manner, namely, a_1 with a_2 and a_3 with a_4 .

$$a_1 = (A, \mu; D, 1 - \mu); a_2 = (B, \mu; D, 1 - \mu), \quad (1.28)$$

$$a_3 = (A, \mu; C, 1 - \mu); a_4 = (B, \mu; C, 1 - \mu). \quad (1.29)$$

The “common consequence violation” occurs if a decision maker prefers a_1 over a_2 and a_4 over a_3 (or a_2 over a_1 and a_3 over a_4). The following restrictions are typically imposed on the lotteries A, B, C, D : (1) The lottery A gives a sure prize of x , i.e., $A = (x, 1)$. (2) The lottery B has some outcomes that are greater than x and others that are less than x . (3) Lottery D first order stochastically dominates lottery C .

The independence axiom implies either of the following two scenarios. Since D is common in the first pair in (1.28) and C is common in the second pair in (1.29), it follows that: (1) If A is preferred to B , then subjects must choose a_1 over a_2 and a_3 over a_4 . If, on the other hand, B is preferred to A , then subjects must prefer a_2 over a_1 and a_4 over a_3 . Thus, the observation of a “common consequence violation,” namely, the empirically observed preference of a_1 over a_2 and a_4 over a_3 , is a refutation of EU, and, in particular, of the independence axiom.

Example 1.3 (*Allais paradox as a common consequence violation. Problems 1 and 2 from Kahneman and Tversky, 1979*) Let a decision maker have initial wealth, x_0 , and utility function, u . All lotteries are in units of Israeli pounds.

Problem 1: (72 subjects) Choose between

$$a_1 = (x_0, 0; x_0 + 2400, 1) \text{ and } a_2 = (x_0, 0.01; x_0 + 2400, 0.66; x_0 + 2500, 0.33).$$

Problem 2: (72 subjects) Choose between

$$a_3 = (x_0, 0.66; x_0 + 2400, 0.34) \text{ and } a_4 = (x_0, 0.67; x_0 + 2500, 0.33).$$

82% of the subjects chose a_1 in problem 1 and 83% of the subjects chose a_4 in problem 2. Each of these preferences was significant at the 1% level.

Example 1.3 can be recast as a common consequence violation as defined in Definition 1.15. To see this, let $\mu = 0.34$ and define the following lotteries.

$$A = (x_0 + 2400, 1); B = \left(x_0 + 2500, \frac{33}{34}; x_0; \frac{1}{34}\right);$$

$$C = (x_0, 1); D = (x_0 + 2400, 1).$$

Then the pairwise comparison lotteries (1.28), (1.29) can be written, in full standard form, as follows.

$$a_1 = (A, 0.34; D, 0.66) = (x_0, 0; x_0 + 2400, 1; x_0 + 2500, 0), \quad (1.30)$$

$$a_2 = (B, 0.34; D, 0.66) = (x_0, 0.01; x_0 + 2400, 0.66; x_0 + 2500, 0.33), \quad (1.31)$$

$$a_3 = (A, 0.34; C, 0.66) = (x_0, 0.66; x_0 + 2400, 0.34; x_0 + 2500, 0), \quad (1.32)$$

$$a_4 = (B, 0.34; C, 0.66) = (x_0, 0.67; x_0 + 2400, 0; x_0 + 2500, 0.33). \quad (1.33)$$

1.5.2 *The probability triangle and violations of the axioms of rationality*

The lotteries, b_1, b_2, b_3, b_4 in (1.26), (1.27) can be written in the following alternative but equivalent notation.

$$b_1 = (0, 1 - P; z, P; Z, 0) \text{ and } b_2 = (0, 1 - p; z, 0; Z, p),$$

$$b_3 = (0, 1 - \mu P; z, \mu P; Z, 0) \text{ and } b_4 = (0, 1 - \mu p; z, 0; Z, \mu p).$$

Using Definition 1.14, we can illustrate the common ratio violation through a useful device, *the probability triangle* (also known as the *Marshak–Machina probability triangle*). Order the three outcomes in Definition 1.14 as $0 < z < Z$ and suppose that we denote their respective probabilities by p_1, p_2 and p_3 . Since p_1, p_3 are the probabilities of outcomes 0 and Z respectively, the probability of the outcome z is simply calculated as a residual, $p_2 = 1 - p_1 - p_3$. Hence, we need only two independent probabilities, p_1, p_3 . Expressing the lotteries, b_1, b_2, b_3, b_4 , in terms of the ordered pair (p_1, p_3) , corresponding to the outcomes 0, Z (probability of z being calculated as a residual) we get:

$$(p_1, p_3): b_1 = (1 - P, 0); b_2 = (1 - p, p); b_3 = (1 - \mu P, 0); b_4 = (1 - \mu p, \mu p) \quad (1.34)$$

The probability triangle corresponding to this case is shown in Figure 1.1. We measure p_1 along the horizontal axis and p_3 along the vertical, such that $p_1, p_3 \in [0, 1]$ (by definition, $p_2 = 1 - p_1 - p_3$). Along the *left edge* of the diagram, $p_1 = 0$, so we get combinations of p_2, p_3 such that $p_2 + p_3 = 1$. Along the *hypotenuse* of the triangle, $p_2 = 0$, so we get combinations of p_1, p_3 such that $p_1 + p_3 = 1$. Along the *lower edge* of the diagram $p_3 = 0$, so we get combinations of p_1, p_2 such that $p_1 + p_2 = 1$.

The lotteries in (1.26), (1.27), expressed in terms of the ordered pair (p_1, p_3) in (1.34), can be represented in the probability triangle. We now draw the indifference curves corresponding to these lottery within the probability triangle. Consider the generic lottery $(0, p_1; z, 1 - p_1 - p_3; Z, p_3)$; we would like to find all combinations of (p_1, p_3) that (for fixed outcomes) generate identical levels of utility, $\bar{u}$, i.e.,

$$p_1 u(0) + (1 - p_1 - p_3) u(z) + p_3 u(Z) = \bar{u}. \quad (1.35)$$

From (1.35), the equation of the indifference curves in the probability triangle is,

$$p_3 = \frac{\bar{u} - u(z)}{u(Z) - u(z)} + \frac{u(z) - u(0)}{u(Z) - u(z)} p_1. \quad (1.36)$$

The outcomes 0, z , Z and $\bar{u}$ are held fixed, hence, (1.36) is the equation of a straight line. Since $0 < z < Z$, and u is a non-decreasing function, it follows that the slope, $[u(z) - u(0)] / [u(Z) - u(z)]$ is constant and positive. Hence, under EU, the indifference curves are all parallel, positively sloping, straight lines. Furthermore, for a fixed p_1 , higher points in the probability triangle (corresponding to higher values of p_3) represent higher values of $\bar{u}$, and lie on a higher indifference curve. The reader can now prove the following two propositions to gain useful practice with the probability triangle (part of the second proposition summarizes some results that we have already proved above).

Proposition 1.6 *The axioms of rationality (see Definition 1.1) translate into various features of the probability triangle. Completeness (see Axiom 1.1a) implies that any two points in the probability triangle lie either on the same indifference curve or different indifference curves. Transitivity (Axiom 1.1b) ensures that no two indifference curves intersect. Continuity (Axiom 1.3) ensures that indifference curves are not thick lines. Independence (Axiom 1.4) implies that indifference curves are parallel straight lines.*

Proposition 1.7 *In the probability triangle, the following properties hold:*

- (i) *The indifference curves under EU are parallel, positively sloping, straight lines.*
- (ii) *Upward shifts (in the north-west direction) of the indifference curves increase utility.*
- (iii) *Steeper indifference curves are associated with greater risk aversion.*

Using the fact that $P > p$, $\mu \in [0, 1]$, and the machinery developed above, the lotteries in (1.34) can now be expressed in the probability triangle in Figure 1.1. Two sets of indifference curves are shown, bold and dotted. The bold indifference curves correspond to EU; these are parallel, straight lines. From Figure 1.1, it is immediate that under EU, decision makers should




Figure 1.1 Illustration of the common ratio effect in a probability triangle.

either choose (1) b_1 over b_2 and b_3 over b_4 (this is the case shown in Figure 1.1), or (2) b_2 over b_1 and b_4 over b_3 (this would be obtained by making the indifference curves in Figure 1.1, flatter). A similar illustration can be given for the common consequence violation.

The actual evidence in the case of the common ratio violation is that the majority of decision makers choose b_1 over b_2 and b_4 over b_3 . These choices provide a clue to the behavior of decision makers relative to the predictions of EU. In order to explain the evidence, as one moves towards the south east corner of the probability triangle, the indifference curves must flatten out, i.e., the decision maker becomes less risk averse (see Proposition 1.7(iii)). This is called the *fanning-out hypothesis* and it is able to explain the common consequence violation and the common ratio violation. By contrast, the independence axiom in EU forces the indifference curves to be parallel, hence, they cannot fan out. Figure 1.1 shows two dotted indifference curves that demonstrate fanning-out.

1.5.3 Some attempts to relax the independence axiom

A number of theoretical models, underpinned by preference foundations, are able to incorporate fanning-out. These models relax the independence axiom, which leads to Allais type violations. We briefly mention some of these models here to give a flavor of the arguments.²¹ One model that generates fanning-out but retains the linearity of indifference curves is the *weighted utility theory* of Chew and MacCrimmon (1979).²² In this case, the utility arising from the lottery $L = (x_1, p_1; x_2, p_2; \dots; x_n, p_n)$ is given by²³

$$U(L) = \frac{\sum_{i=1}^n p_i \phi(x_i) u(x_i)}{\sum_{i=1}^n p_i \phi(x_i)},$$

where $\phi : \mathbb{R} \rightarrow \mathbb{R}_+$. Weighted utility theory uses a weak form of the independence axiom, called *weak independence*; Starmer (2000) raises concerns about the intuitive appeal of this axiom.

Definition 1.16 (*Weak Independence*) Suppose that for lotteries $L_1, L_2 \in \mathcal{L}$, $L_1 \prec L_2$. Then, for each $p \in (0, 1)$, there is a unique $q \in (0, 1)$ such that for all $L_3 \in \mathcal{L}$,

$$qL_1 + (1 - q)L_3 \prec pL_2 + (1 - p)L_3.$$

Another theory that allows for fanning-out is the theory of *disappointment aversion*; see for instance, Bell (1985) and Loomes and Sugden (1986).²⁴ In this formulation, for a lottery $L = (x_1, p_1; x_2, p_2; \dots; x_n, p_n)$ such that $x_1 < x_2 < \dots < x_n$, the decision maker derives the utility

$$D(L) = \sum_{i=1}^n p_i [u(x_i) + D(u(x_i) - \bar{U})]. \quad (1.37)$$

In (1.37), $\bar{U}$ is the utility that the decision maker expects from the lottery, say, the certainty equivalent of the lottery. $u(x_i)$ is the utility that actually arises when the outcome is x_i . If $u(x_i) - \bar{U} > 0$,

²¹ The interested reader can consult Wakker (2010), Sugden (2004), and Starmer (2000) for more details.

²² For various axiomatic foundations of weighted utility theory, see Chew and MacCrimmon (1979), Chew (1983), and Fishburn (1983).

²³ Two other models that retain linearity of the indifference curves but also permit fanning-out are the *implicit expected utility model* of Dekel (1986) and the *implicit weighted utility model* of Chew (1989).

²⁴ See Section 2.10.2 for a more detailed exposition of disappointment aversion.

the decision maker experiences *elation*. *Disappointment* is experienced if $u(x_i) - \bar{U} < 0$. *Disappointment aversion* is built into the model by assuming that $D(x)$ is concave for $x < 0$ and convex for $x > 0$.

It turns out that linearity of indifference curves in the probability triangle is equivalent to the *betweenness axiom*, which weakens the independence axiom by requiring the indifference curves to be straight lines, not necessarily parallel.

Definition 1.17 (*Betweenness*): Given two lotteries a_1, a_2 , if $a_2 \prec a_1$, then *betweenness* implies that for all $p \in (0, 1)$, $a_2 \prec (a_1, p; a_2, 1 - p) \prec a_1$.

To visualize betweenness, think of straight line indifference curves in a probability triangle. Suppose that a_1 lies on a higher indifference curve than a_2 . Then all linear combinations of a_1, a_2 lie on an indifference curve no higher than the indifference curve on which a_1 lies and no lower than the indifference curve on which a_2 lies. If the binary relation is one of indifference, betweenness implies that if $a_1 \sim a_2$, then for all $p < 1$, $a_1 \sim (a_1, p; a_2, 1 - p) \sim a_2$. This is also immediately seen by looking at one of the straight line indifference curves in the probability triangle. If the lotteries a_1, a_2 lie on the same indifference curve, then the decision maker is also indifferent between all linear combinations of a_1, a_2 because they too lie on the same indifference curve. However, while betweenness implies (and is implied by) linear indifference curves, it does not impose the restriction that the indifference curves be parallel in the probability triangle.

Now suppose that we relax betweenness and consider the two cases: concave indifference curves (i.e., quasi-convex preferences), and convex indifference curves (i.e., quasi-concave preferences). These cases are shown in Figures 1.2 and 1.3, respectively. In each diagram, the dotted line connecting the two lotteries, a_1, a_2 , among which the decision maker is indifferent, represents all possible linear combinations of a_1, a_2 , i.e., lotteries of the form $(a_1, p; a_2, 1 - p)$. When the indifference curves are concave, then any lottery of the form $(a_1, p; a_2, 1 - p)$ lies on a lower indifference curve. Hence, the decision maker is averse to randomizing among a_1, a_2 . The converse applies in Figure 1.3 where the indifference curves are convex, and the decision maker prefers to randomize.




Figure 1.2 Concave indifference curves.




Figure 1.3 Convex indifference curves.

The possibilities shown in Figures 1.2 and 1.3 are not the only ones. For instance, Chew et al. (1991) propose a model that is quadratic in probabilities. In this model, the utility from a lottery $L = (x_1, p_1; x_2, p_2; \dots; x_n, p_n)$ is written as follows.

$$Q(L) = \sum_{i=1}^n \sum_{j=1}^n p_i p_j v(x_i, x_j). \quad (1.38)$$

The utility function in (1.38) is generated by a weakened form of the independence axiom called *mixture symmetry*, which we define next.

Definition 1.18 (*Mixture asymmetry*) *Mixture symmetry requires that for all lotteries, a_1, a_2 , if $a_1 \sim a_2$, then for all $p < 1$, $(a_1, p; a_2, 1-p) \sim (a_2, p; a_1, 1-p)$.*

In the probability triangle, the indifference curves generated by (1.38) can exhibit both concavity and convexity over different domains. In other words, these indifference curves exhibit both fanning-in and fanning-out. The *theory of disappointment* in Gul (1991) incorporates both these features.

We now turn briefly to the empirical evidence.²⁵ Empirical evidence for fanning-out is, at best, mixed. The results do not change if we use real or hypothetical gambles (Camerer, 1995, p. 634). In a careful empirical study, Chew and Waller (1986) reported violations of the independence axiom in the common consequence and the common ratio forms. They also found that betweenness is not consistent with their data. Other studies have also subsequently cast doubt on the betweenness axiom. It is not easy, however, to discern a pattern to these violations; for a review of nine studies see Camerer and Ho (1994). Fanning-out is found when the gambles are unattractive but not when they are attractive; in the probability triangle, these are respectively the southeast and northwest corners (Camerer, 1989).

Fanning-in has been found in several parts of the probability triangle. For instance, along the *lower edge* (Starmer and Sugden, 1989a; Battalio et al., 1990), northwest corner (Conlisk, 1989), lower edge and southeast corner (Prelec, 1990), and moving up along the left edge (Starmer, 1992).

²⁵ For more detailed surveys, see Camerer (1995), Starmer (2000), Sugden (2004), and Wakker (2010).

While it is difficult to establish a pattern, it would seem that there is sufficient evidence of fanning-in. This, however, casts doubt on fanning-out as an explanation of the common consequence and the common ratio violations.

In the light of this evidence, theories which incorporate both fanning-in and fanning-out appear to have a better chance of explaining the data. This includes theories such as those proposed by Gul (1991) and Neilson (1992); however, the lack of evidence supporting the betweenness axiom, which both these papers assume, is a cause for concern.

Violations of EU are less pronounced, but still significant, in the interior of the probability triangle as compared to the edges (Conlisk, 1989; Camerer, 1992; Harless, 1992). This is most likely caused by non-linear weighting of probabilities and the certainty effect. However, a reduction in violations towards the interior of the probability triangle does not rescue EU. We shall see below that many economically important events involve events that occur towards the boundary of the probability triangle.

There is also evidence that some salient aspects of the behavior of human decision makers under risk might be shared with that of other animals. For instance, it has been found that rats also violate betweenness. They have indifference curves that fan-in and fan-out and that exhibit the common ratio violation.²⁶ There are also evolutionary reasons why animals may neglect events of very low probability.²⁷ Hence, many violations of EU arise from a shared evolutionary history of mammals.

1.5.4 *Attitudes to risk for small and large stakes: Rabin's paradox*

In a fundamental result, Rabin (2000a) showed that under EU risk aversion over small stakes implies implausibly high risk aversion over large stakes. A more accessible treatment of the issues can be found in Rabin (2000b).²⁸ We first state the result informally. Suppose that the utility function of a decision maker who follows EU is increasing and concave. Then if the decision maker turns down a moderate stakes gamble, then he/she will always turn down a large stakes gamble.

Let a decision maker have initial wealth, w . Define three lotteries, L_0, L_1, L_2 as follows. $L_0 = (w, 1)$, $L_1 = (w - 10, 0.5; w + 11, 0.5)$, and $L_2 = (w - 100, 0.5; w + m, 0.5)$, where m is any positive level of wealth. Consider assumption A1.

A1. Suppose that for all levels of w the decision maker expresses the preference

$$L_1 \prec L_0. \quad (1.39)$$

Under EU, (1.39) implies that

$$u(w) > 0.5u(w - 10) + 0.5u(w + 11), \quad (1.40)$$

²⁶ See Battalio et al. (1985), Kagel et al. (1990), and MacDonald et al. (1991).

²⁷ I am grateful to Herbert Gintis for bringing this to my attention. Evolutionary dynamics show that very small probabilities will be ignored because (a) the rate of false positives will be high and (b) the cost of maintaining the neural machinery to detect and react to low probability events is too high; see Heiner (1983).

²⁸ Wakker (2005) not only gives a particularly lucid account of these results but also discusses the context of subsequent remarks and some confusions in the literature.

where $u(w)$ is the utility from outcome w . Since the expected value of lottery L_1 is higher, yet $L_1 \prec L_0$, thus, the decision maker is risk averse over small stakes.²⁹ Under EU this implies that the utility function u must be concave. Rabin's (2000a) central insight is to show that for arbitrarily large values of m , the decision maker must then have the preferences

$$L_2 \prec L_1, \quad (1.41)$$

i.e., by turning down a lottery that offers infinite wealth with probability 0.5, the decision maker expresses an unreasonable amount of risk aversion. We now give a heuristic proof of Rabin's result.

Rewriting (1.40) we get

$$u(w + 11) - u(w) < u(w) - u(w - 10). \quad (1.42)$$

At this stage, the reader may find it useful to visualize an increasing, concave, utility function u at three values: $w - 10$, w , $w + 11$. The following sequence of steps can then be constructed:

$$\begin{aligned} u'(w + 11) &< \frac{u(w + 11) - u(w)}{11} \quad (\text{from concavity of } u), \\ \frac{u(w + 11) - u(w)}{11} &< \frac{u(w) - u(w - 10)}{11} \quad (\text{since (1.42) holds}), \\ &= \frac{10}{11} \frac{[u(w) - u(w - 10)]}{10}, \\ &\leq \frac{10}{11} u'(w - 10) \quad (\text{from concavity of } u). \\ \Rightarrow u'(w + 11) &< \frac{10}{11} u'(w - 10). \end{aligned} \quad (1.43)$$

In other words, over an interval of length 21, marginal utility drops by $\frac{10}{11}$ of the original value. Thus, each unit of money between w and $w + 11$ is valued, on average, at most $\frac{10}{11}$ of each unit of money, on average, between $w - 10$ and w .

Table 1.1 Drop in marginal utility for successive intervals in Rabin's calibration theorem.

Interval size	21	210	2100	21000	210000
Final to initial marginal utility	0.90909	0.38554	73×10^{-6}	4.0487×10^{-42}	1.1834×10^{-414}

Table 1.1 shows this drop as the interval size increases successively (recall that m can be as large as we want it to be). Successive numbers in the second row of table 1.1 are derived from the sequence $(\frac{10}{11})^1, (\frac{10}{11})^{10}, (\frac{10}{11})^{100}, \dots$

The problem is that under EU, the marginal utility drops off too rapidly. Consequently, and this is crucial to Rabin's results, for extremely large rewards, the marginal utility of the large reward

²⁹ The expected value of the lottery L_1 is $w + 0.5$ but the certainty equivalent is less than w , hence, the decision maker is risk averse.

relative to a relatively small loss becomes so miniscule that the lottery has little attraction for the decision maker, hence, accounting for the preferences shown in (1.41). Thus, if an EU decision maker is risk averse for small stake gambles (which is reasonable), then the same decision maker is absurdly risk averse for large stake gambles (which is unreasonable). By contrast, there is no puzzle under prospect theory (see Section 3.6).

Example 1.4 Let us now consider a concrete example to illustrate Rabin's paradox under EU.³⁰ If the utility function, u , is not bounded above, then, clearly, $u(\infty) = \infty$, hence, $\frac{1}{2}u(w + \infty) = \infty$. It follows that for any real numbers, w and x , i.e., for any finite level of wealth, w , and any finite loss, x ,

$$EU(L_2) = \frac{1}{2}u(w - x) + \frac{1}{2}u(w + \infty) = \infty > u(w).$$

Thus, any decision maker who follows EU will accept the gamble L_2 . Hence, the utility function must be bounded above. Consider the utility function

$$u(w) = 1 - e^{-w}, \quad (1.44)$$

Thus, $u'(w) > 0$ and $u''(w) < 0$. Under the utility function in (1.44), the preference $L_1 \prec L_0$ for any level of initial wealth, w , implies that

$$\frac{1}{2}(1 - e^{-(w-10)}) + \frac{1}{2}(1 - e^{-(w+11)}) < 1 - e^{-w} \Leftrightarrow 22026 > 1$$

Hence, assumption A1 holds for the utility function in (1.44). We now wish to show that for this utility function, it is also the case that $L_2 \prec L_1$, i.e., the decision maker turns down a lottery that offers infinite wealth with a probability of 0.5. The preference $L_2 \prec L_1$ for any level of initial wealth, w , implies that

$$\frac{1}{2}u(w - 100) + \frac{1}{2}u(w + \infty) < \frac{1}{2}u(w - 10) + \frac{1}{2}u(w + 11).$$

For the utility function in (1.44) we have

$$u(w + \infty) = 1 - e^{-(w+\infty)} = 1 \text{ and } u(w - 100) = 1 - e^{-(w-100)},$$

It follows that the preference $L_2 \prec L_1$ implies

$$\frac{1}{2}(1 - e^{-(w-100)}) + \frac{1}{2}(1 - e^{-(w+\infty)}) < \frac{1}{2}(1 - e^{-(w-10)}) + \frac{1}{2}(1 - e^{-(w+11)})$$

or $e^{100} > e^{10} + e^{-11}$, which is true. Hence, under Assumption A1, an EU decision maker exhibits the absurd preference $L_2 \prec L_1$ (Rabin's paradox).

1.5.5 Violations of description invariance

In the typical applications of EU, alternative methods of describing lotteries are assumed not to have any effect on choices. This is known as *description invariance* or the absence of *framing effects*. Description invariance is not required by the axioms of EU. Human behavior is state, time, and context dependent, irrespective of the decision theory that is used. However, as we shall see below, in Example 1.5, a particular violation of description invariance served as a building

³⁰ I am grateful to Ali al-Nowaihi for suggesting this example.

block of a non-EU theory, prospect theory, that we study in more detail below. This example also illustrates the relative risk attitudes of subjects when they perceive that they are in a gain situation as compared to a loss situation.

Example 1.5 (Tversky and Kahneman, 1981): *Imagine that the US is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows:*

Positive Framing: If program A is adopted, 200 people will be saved. If program B is adopted, there is a one-third probability that 600 people will be saved and a two-thirds probability that no people will be saved. Which of the two programs would you favor?

Negative Framing: If program C is adopted, 400 people will die. If program D is adopted, there is a one-third probability that nobody will die and a two-thirds probability that 600 people will die. Which of the two programs would you favor?

Under positive framing, the solution is framed in terms of lives saved (a gain situation). The result was that 72% of the subjects chose A, the safe program, and 28% chose the risky program B. For the case of negative framing, i.e., lives lost (a loss situation), 22% voted for C and 78% for D, the risky program. The options in each frame are identical. In each frame, two hundred lives are saved for sure in the first option (options A and C) and two hundred lives are expected to be saved in the second option (options B and D). While the aim of Kahneman and Tversky was to show that individuals are risk averse for gains and risk loving for losses, an issue of critical importance in prospect theory (see Chapter 2), the framing of the same question leads to completely different choices.

In this book, we shall review overwhelming evidence of framing effects, which shows that description invariance is not supported.³¹ Among the main decision theories, only prospect theory is able to handle framing effects by altering the *reference point*.

1.5.6 Preference reversals

The basic preference reversal problem is that under one elicitation method, lottery A is preferred to lottery B, while under another method, B is preferred to A; this result also holds for experienced players (Lichtenstein and Slovic, 1971).³² Consider two probabilities $p < P$, two outcomes $z < Z$, and the following pair of lotteries.

$$P\text{-bet} = (0, 1 - P; z, P), \quad \$\text{-bet} = (0, 1 - p; Z, p). \quad (1.45)$$

The names *P*-bet and *\$*-bet come from the fact that the first lottery has a higher probability of winning, while the second has a higher prize. The decision maker must make a choice between the two bets in two different tasks.

1. In the first task, decision makers directly choose between the *P*-bet and the *\$*-bet.
2. In the second task, decision makers are asked to assign certainty equivalents to the *P*-bet and the *\$*-bet, respectively, C_P and $C_\$$.

³¹ It is likely that future decision theories will take greater account of the choice heuristics used by decision makers. We consider some models of choice heuristics elsewhere in the book. Rubinstein's (1998) similarity relation is one example of a model that takes account of the procedural elements of choice.

³² For the literature circa late 1980s see Tversky and Thaler (1990).

The typical empirical finding in a large number of studies is that the P -bet is chosen over the $\$$ -bet in the first task and $C_P < C_\$$ in the second task (Lichtenstein and Slovic, 2006).³³ These findings of *preference reversals* contradict EU (or equivalently, the axioms of rationality).

The empirical findings from tasks 1 and 2 imply that: $\$ \text{-bet} \sim C_\$ > C_P \sim P \text{-bet} > \$ \text{-bet}$, which is a failure of transitivity. Theories that allow for violation of transitivity such as regret theory (Loomes and Sugden, 1982) can potentially explain these findings. Other attempts to explain preference reversals include abandoning the reduction principle (Segal, 1988), and abandoning the independence axiom (Holt, 1986; Karni and Safra, 1987). Tversky and Thaler (1990) argue that none of these attempts can fully account for preference reversals. Tversky et al. (1990) put the various explanations to the test and find that it is violations of *procedure invariance* that explain the finding of preference reversals. Gintis (2009, p. 238) expresses the view that the expected value of the lotteries used in preference reversal experiments is too close to permit firm conclusions to be drawn. Lichtenstein and Slovic (2006) offer a recent view on the subject.

If transitivity does not hold, then there can be no preference foundations for utility functions. For this reason, most economists may feel uneasy about abandoning transitivity. Insofar as intransitivity arises from changing the frame, context, or reference point, one may consider preferences over objects, conditional on a given frame, context, or reference point. Such conditional preferences would then be transitive. This idea is formally pursued in Sugden (2003). However, the resulting utility representation is linear in probabilities, as in subjective expected utility, hence, it would contradict the burgeoning evidence on non-linear probability weighting (Fehr-Duda and Epper, 2012).

We shall show later in Part 1 that preference reversals can also be explained using the *third generation of prospect theory* due to Schmidt et al. (2008). Preference reversals can also arise under *ambiguity* but they are of a different nature; see Trautmann et al. (2011).

1.5.7 *Is the reduction axiom supported by the evidence?*

Under EU, decision makers are able to reduce a compound lottery (where the outcomes themselves are lotteries) to a simple lottery of the form $L = (x_1, p_1; x_2, p_2; \dots; x_n, p_n)$ (see the *reduction axiom*, Axiom 1.5, above). However, empirical evidence is not supportive of the reduction axiom (Bernasconi and Loomes, 1992; Conlisk, 1989; Miao and Zhong, 2012). Harrison et al. (2012) find mixed support for the hypothesis. In simple one outcome choice problems there is support for the reduction axiom but in multichoice problems, the axiom does not hold.

Abdellaoui et al. (2014) find that the reduction axiom does not hold. They compute the certainty equivalent of compound lotteries, C_c , and the certainty equivalent of their associated reduced form lotteries, C_r . They term the difference between the two certainty equivalents, $C_r - C_c$, as the *compound risk premium*. If the difference is zero, the reduction axiom holds; if the difference is positive, the decision maker is *compound risk averse*; and if the difference is negative, the decision maker is *compound risk seeking*. They find that the difference is non-zero and most individuals are compound risk averse, i.e., they are more unwilling to play the compound lottery relative to the equivalent reduced lottery.

It also turns out that there are important links between the decision maker's attitudes to ambiguity aversion and the compound risk premium. Halevy (2007) found the presence of compound

³³ Also see Hershey and Schoemaker (1985) and Johnson and Schkade (1989).

risk premium as a necessary condition for non-neutral ambiguity attitudes. Indeed, he proposed that the failure of the reduction axiom is an important explanation of the Ellsberg paradox (see Chapter 4). Abdellaoui et al. (2014) find support for this view but their evidence suggests a weaker link between ambiguity aversion and compound risk premium. We postpone a discussion of models of ambiguity to Chapter 4.

CHAPTER 2

Behavioral Models of Decision Making

2.1 Introduction

Under expected utility theory (EU), decision makers weight probabilities linearly, and in each state of the world, they derive utility from the final level of wealth 1. In this chapter, we relax both these features of EU and consider the main behavioral decision theories under risk and uncertainty. Theories of ambiguity are considered in Chapter 4.

The theories considered in this chapter include behavioral decision theories such as *rank dependent utility* (RDU), *original prospect theory* (OPT), *cumulative prospect theory* (PT), and theories of *regret and disappointment aversion*. We also consider the more recent theories such as *third generation prospect theory*, *prospect theory with endogenous reference points*, and *composite prospect theory* (CPT). The behavioral decision theories explain the evidence much better than EU. The continued use of EU in large swathes of contemporary research in economics remains an enduring puzzle for most behavioral economists.

The linear weighting of probabilities under EU leads to many puzzles. For instance, Chapter 1 described the Allais paradox that is highly suggestive of *non-linear probability weighting*; evidence reviewed in this chapter gives further support to non-linear probability weighting. A *probability weighting function* reflects the mental representation of objective or subjective probabilities by a decision maker. It is central to almost all alternatives to EU. Section 2.2 introduces the basics, with particular emphasis on the *Prelec probability weighting function* (Prelec, 1998), although several other weighting functions are also outlined.

Empirical evidence shows heterogeneity in probability weighting functions for individual decision makers. One observes convex weighting, concave weighting, and *inverse S-shaped weighting* (Fehr-Duda and Epper, 2012). However, once we pool this information, at the aggregate level, the typical finding is that of an *inverse S-shaped* probability weighting, i.e., low probabilities are overweighted and high probabilities are underweighted. It turns out that one cannot simply replace probabilities under EU/SEU by probability weighting functions because this may lead to violation of first order stochastic dominance (Fishburn, 1978); we shall show this formally.

Section 2.3 introduces RDU under risk (Quiggin, 1982) and under uncertainty (Schmeidler, 1989). The idea under RDU is to replace the probabilities in an EU functional by *decision weights* that are *cumulative transformations of probabilities*. Decision makers who use RDU do not choose first order stochastically dominated options. RDU can explain several puzzles that EU cannot—the Allais paradox is a prominent example. Not surprisingly, preference foundations of RDU relax

the independence axiom in EU. One implication of inverse S-shaped probability weighting under RDU is that since low probabilities are overweighted, financial markets may price the skewness of asset returns; these issues are considered in more detail in Part 7 (Barberis, 2013).

Under EU, attitudes to risk, for instance, the degree of risk aversion, depend entirely on the shape of the utility function. However, under RDU, attitudes to risk depend *jointly* on the shapes of the probability weighting function and the utility function. This also holds for other behavioral decision theories that use non-linear probability weighting. For instance, under RDU, if the probability weighting function is convex throughout, then the decision maker can be shown to be *pessimistic* in the sense that he assigns relatively higher decision weights to lower outcome. A decision maker with a weighting function that is concave throughout, on the other hand, can be shown to be *optimistic* and places relatively higher decision weights on higher outcomes.

In what follows, we abbreviate the *original prospect theory* of Kahneman and Tversky (1979) by OPT and *cumulative prospect theory* of Tversky and Kahneman (1992) by PT. The main difference between these theories is the way that they account for non-linear probability weighting. Following convention, unless we wish to make a distinction between the two theories, we shall use the abbreviation PT for both theories.¹

Both EU and RDU assume that decision makers derive utility from *final levels* of variables of interest, such as wealth. These theories can account for utility that is derived from changes in outcomes relative to a *reference point*. However, since EU/SEU/RDU do not have insights about how behavior in the domain of gains and losses may differ, incorporating a reference point in these theories is not very illuminating.

In contrast, PT appeals to robust psychological evidence to argue that decision makers derive utility from *changes* in variables of interest relative to some *reference point*. While it is usually argued that the reference point under PT is exogenous, Kahneman and Tversky justified the status quo as a useful reference point. In other cases, a social norm of fairness, or a legal entitlement may serve as reference points. The reference point partitions the domain of outcomes into *gains* and *losses*. Furthermore, there are robust differences in behavior in gains and losses. For example, decision makers exhibit *loss aversion*: losses bite more than equivalent gains. Section 2.4 considers PT under risk when the reference point is a fixed, non-stochastic, entity. Section 2.4.1 shows that the extension of PT to uncertainty is straightforward.

The utility function of the decision maker under PT (unlike EU/SEU/RDU) is *concave in gains* and *convex in losses*. Thus, a central insight of PT, based on robust evidence, is that attitudes to risk are fundamentally different in the two domains. Section 2.4.2 shows that risk attitudes under PT are even richer than RDU, since they are jointly determined by the shapes of the utility function and the probability weighting function, and we distinguish between gains and losses. Under PT, one observes a *four fold pattern of attitudes to risk aversion* (Tversky and Kahneman, 1992). Although decision makers can choose stochastically dominated options under OPT they cannot do so under PT. The reason is that PT borrows its machinery on non-linear probability weighting from RDU in which decision makers never choose dominated options.

While PT is the most satisfactory and influential account of human behavior under risk and uncertainty (Wakker, 2010), it is work in progress. Section 2.4.4 considers some common criticisms of PT that are misplaced. In contrast, Section 2.9 considers more serious criticisms of PT that an even more complete theory of decision making may wish to address in due course.

¹ Daniel Kahneman won the Nobel prize for PT in 2003. Amos Tversky, who would almost surely have shared this prize, had sadly died by that time, and the Nobel prize cannot be conferred posthumously.

Section 2.5 considers the applied problem of eliciting the utility function and the probability weighting function under PT (Wakker and Deneffe, 1996; Bleichrodt and Pinto, 2000; Abdellaoui, 2000; Köbberling and Wakker, 2003; Abdellaoui et al., 2007). In this section, we also consider the various definitions of loss aversion, as well as a brief discussion of some of the parameter estimates under PT.

Section 2.6 considers the preference foundations of PT in three different subsections. Section 2.6.1 considers the preference foundations of PT based on the *trade-off consistency method* (Chateauneuf and Wakker, 1999; Köbberling and Wakker, 2003). The trade-off consistency method is one of the most profound and practical achievements in behavioral decision theory. This method is able to provide a unified account of the preference foundations of PT, RDU, and EU in terms of risk and uncertainty.² Next, we consider the preference foundations of the two most widely used components in applied PT: the *Prelec probability weighting function*, and the *power form of utility* that fits well the experimental data.

Section 2.6.2 considers the preference foundations of the Prelec function (al-Nowaihi and Dhami, 2006). Section 2.6.3 considers the preference foundations of the power form of utility (al-Nowaihi et al., 2008). In particular, we show that under the axiom of *preference homogeneity*, not only does the utility function take the power form, but it also implies loss aversion, identical powers of utility in the domain of gains and losses, and identical probability weighting functions in the domain of gains and losses.

A decision maker's endowment may be stochastic. For instance, houses, consumer durables, financial assets, and other personal belongings may suffer a loss in value with some probability, arising on account of, say, changes in market value, theft, accidents, and other natural hazards. Thus, expectations of the loss may influence one's reference point. We now consider two ways in which stochastic reference points have been incorporated into PT.

The first proposal, in Section 2.7, considers *third generation prospect theory*, PT³ (Schmidt et al., 2008). *Reference dependent expected utility* (Sugden, 2003) is a special case of this model. PT³ proposes a representation of preferences over *Savage acts*.³ The representation is similar to PT, except that the utility function is now defined over outcomes and a *stochastic reference point*. This framework can explain disparities between *willingness to pay* (WTP) and *willingness to accept* (WTA), sometimes also known as the *endowment effect* (see Chapter 3). It can also explain *preference reversals*.

The second proposal, in Section 2.8, considers the Köszegi–Rabin framework for *endogenous stochastic reference points* (Köszegi and Rabin, 2006, 2007, 2009). They distinguish between three equilibrium concepts—the two main concepts differ in the timing of actual decisions, prior to the resolution of uncertainty. Conditional on a stochastic reference point, the decision maker makes ex-ante choices for future states of the world, which, in equilibrium, are confirmed to be optimal, ex-post. The initial choices then become self-fulfilling; this is the sense in which reference points are endogenous.

We consider several illustrative examples of this framework to convey the insights of the two main equilibrium concepts. In particular, in Section 2.8.1 we consider the newsvendor problem (Herweg, 2013) and in Section 2.8.2, the backward bending labour supply curve. In Section 2.8.3, we consider the empirical evidence, which is consistent with the predictions of the Köszegi–Rabin framework (Abeler et al. 2011; Ericson and Fuster, 2011). On the one hand, the Köszegi–Rabin

² Our discussion is necessarily condensed. For a more detailed account, see the masterful treatment by one of the leaders in this area, Wakker (2010).

³ Given a set of uncertain events that form a partition of the sample space, a Savage act specifies the outcome (possibly monetary) for the decision maker, for each uncertain event.

framework appeals to those who were uneasy with the traditional justifications for a reference point in PT (e.g., status quo, social norms, and other legal or social entitlements). However, on the other hand, this framework has extremely high cognitive requirements, which sit uneasily with the main thrust of behavioral models. The authors are fully aware of these difficulties, and are at pains to highlight them, but no practical alternatives to the rational expectations requirement that could be used in applied work have been proposed. Whatever the merits of these arguments, empirical evidence must be the final arbiter. Future research must consider even more stringent tests of the framework; this is likely to be an area of continued importance.

Section 2.9 considers some limitations of PT. In Section 2.9.1, we consider the critique based on several papers by Michael Birnbaum and associates; for a summary, see Birnbaum (2008). These experiments consider the implications of splitting outcomes or probabilities that, under PT, are not predicted to alter the choice between lotteries. Yet, empirically, this is found to be violated. However, this evidence appears to reject most decision theory models. The associated theoretical alternatives suggested by this literature that go under the name of *configural weights models* are considered in a later section, however, these are not without problems either.

Section 2.9.2 considers an implication of PT, *gain-loss separability* (Wu and Markle, 2008). Suppose that we split two mixed lotteries,⁴ A and B , into their respective gain and loss components. If the decision maker prefers each of the components of A over the corresponding components of B , then, under PT, he should also prefer A to B . Yet this is found to be violated. However, these results need to be replicated with a wider range of probabilities and payoffs to judge not just how general they are, but also to discover the mechanisms at work.

Section 2.9.3 considers the *stake size effect*, that cannot be accounted for by any theory that separates utilities from probabilities.⁵ The stake size effect refers to the observed change in relative risk aversion as the stake size in the domain of gains goes up. In an important advance, it has been shown that the stake size effect arises not from the utility function but from the probability weighting function (Fehr-Duda et al., 2010). This suggests that the development of theories that are non-separable in probabilities and outcomes might be a fruitful area of research in the future.

Most non-stochastic decision theories of choice typically predict an unambiguous choice of one lottery over another. However, the typical experimental finding is a dispersion of choices across subjects for pairs of lotteries. One possible resolution is to introduce arbitrary noise or errors into the choices of subjects. Although not uncommon in empirical work, this is neither intellectually satisfying, nor well founded unless one has good reasons for preferring one noise structure to another. A second resolution, is to allow for the possibility that different individuals may follow different decision theories. This has led to the exploration of heterogeneity in preferences using a *mixture model approach* (Bruhin et al., 2010; Fehr-Duda et al., 2010; Conte et al., 2011). The main empirical finding is the presence of *preference heterogeneity*. For instance, in three separate datasets, Bruhin et al. (2010) find the respective shares of EU and RDU/PT preferences to be 20% and 80%. By contrast, the shares of SEU types are much higher in the absence of a mixture model approach that relies instead on noisy choices (Hey, 2014).

Section 2.10 considers a range of other behavioral theories of decision making. These theories offer additional insights into decision making that we have not considered so far. Section 2.10.1 considers *regret theory* (Bell, 1982; Loomes and Sugden, 1982; Fishburn, 1982a). The idea is that in any state of the world, the decision maker derives utility not only from his own choice but also from the forgone outcomes from other choices that were not made. Regret theory allows for the

⁴ Mixed lotteries contain outcomes in the domain of gains and losses.

⁵ These theories include EU, RDU, OPT, PT with fixed and stochastic reference points, PT³, CPT, and disappointment aversion.

violation of transitivity. One can also explain the Allais paradox using regret theory. However, the evidence is not always consistent with the pattern of intransitivity predicted by the theory (Starmer, 2000).

Section 2.10.2 considers the *theory of disappointment aversion* (Gul, 1991). This is a theory of endogenous reference points because the decision maker treats the certainty equivalent of the lottery under consideration as his reference point. All outcomes in the lottery that are less than the certainty equivalent are coded as loss outcomes, and all outcomes that are greater than the certainty equivalent are coded as gain outcomes. Non-linear probability weights are then used to form decision weights. However, the construction of decision weights is different from RDU and PT. The resulting model may, however, be observationally similar to RDU (Fehr-Duda and Epper, 2012).

Section 2.10.3 considers the role of *salience* and *context* in making choices among risky alternatives (Bordalo et al., 2012). In this case, the decision maker compares pairwise outcomes in each choice and assigns salience to the outcomes, using a salience function that depends on the relative and the absolute sizes of the outcomes. The outcomes are then ranked by salience and decision weights assigned to the outcomes. More salient outcomes receive higher decision weights. This theory can explain the Allais paradox, preference reversals and context dependent preferences.

In Section 2.10.4, we consider *case-based decision theory* (Gilboa and Schmeidler, 1995, 1996, 1997, 2001). The model is based on an extremely plausible idea. Namely, that in deciding on our current actions, we first identify the environment, or situation. Next, we look at past actions taken in various environments and compare the current environment–action pairs to all environment–action pairs in the past using a *similarity function*. Such a comparison may suggest the optimal action that we must take today. Under certain conditions, the long–run behavior of this process converges to one that would be chosen by an expected utility maximizer. However, the evidence strongly indicates that expected utility is refuted (see Chapter 1). A drawback of case-based decision theory is that it does not allow one to construct hypothetical cases. Indeed, it would seem that since no environment is likely to be an exact replica of another, one’s ability to construct hypothetical cases (perhaps also based on past environment–action pairs) are critical in constructing an even better theory. However, some emerging evidence is consistent with the theory (Pape and Kurtz, 2013; Bleichrodt et al., 2013).

Section 2.10.5 considers *configural weights models* that can explain the evidence that we outline in Section 2.9.1 (Birnbbaum, 2008). We consider a range of models in this class. These include the *rank affected multiplicative weights model*, the *transfer of attention in exchange model*, and the *gains decomposition utility model*. Unlike most other behavioral models that we consider, there are no preference foundations for these models. Further, the construction of decision weights in these models are likely to appear arbitrary to many economists. Perhaps for this reason, this class of models has not yet gained prominence in economics.

Section 2.11 considers the extremely important, but often neglected issue of the behavior of decision makers towards *events of very low probability*. In a range of situations, one observes that a fraction of decision makers ignore events of very low probability, which can lead to serious consequences for the decision maker.⁶ Examples, with well-documented empirical support include the following (al-Nowaihi and Dhami, 2010a). Reluctance to buy insurance against low probability natural hazards; running red traffic lights; driving and talking on mobile phones; and not wearing non-mandatory seat belts. Existing theories cannot explain human behavior in these cases. EU certainly cannot explain it but the problem is even worse when one uses a typical

⁶ Peter Wakker (personal communication) points out that Buffon (1777, p. 72) argued that all probabilities less than .0001 be treated as “morally” equal to zero.

probability weighting function (such as the Prelec function) in any of the non-EU theories such as RDU and PT. The reason is that the typical probability weighting function is too steep near the origin, so, such probabilities should be tremendously salient, rather than be ignored.

In other words EU and most non-EU theories (except OPT) predict that nobody should run red traffic lights, or refrain from wearing seat belts in vehicles, or drive and talk on mobile phones simultaneously, or refrain from buying insurance against even the most unlikely natural hazard. In Section 2.11.1, we argue that Kahneman and Tversky (1979) were well aware of these issues in OPT. In the *editing phase* in OPT, a fraction of the decision makers simply omit very low probability events from the lotteries that they are presented with. In the next step, the *evaluation phase*, decision makers choose among alternatives that have been edited in the first phase.

Our main explanation of human behavior for low probability events is based on *composite prospect theory* (CPT) (al-Nowaihi and Dhami, 2010a). In CPT, a fraction of the decision makers use a standard, inverse S-shaped, probability weighting function, e.g., the Prelec function. However, the remaining fraction use the *composite Prelec function* (CPF), which is an axiomatically founded modification of the Prelec function. Decision makers who use CPF, significantly underweight very low probability events (unlike the Prelec function). Although CPT can explain human behavior in all the above mentioned puzzles of human behavior for very low probability events, it awaits empirical testing.

The outcomes of most human decisions are *uncertain* and take *time* to materialize. Yet, it is usual practice to treat *risk preferences* and *time preferences* separately. An emerging literature makes a strong case to integrate the two. Section 2.12 considers the relation between risk and time preferences. We also give a complementary discussion of this topic in Part 4 of the book.

The usual method of integrating risk and time preferences is to offer decision makers a sequence of dated payoffs, but future payoffs materialize with a probability less than one. In Section 2.12.1, we use this method, and first consider evidence which suggests that neither EU, nor any of the mainstream behavioral alternatives can explain the experimental choices (Andreoni and Sprenger, 2012b). However, it is crucial to specify how one reduces lotteries that have a time dimension. Using a plausible method of reduction, the portfolio method, one can explain the experimental findings in Andreoni and Sprenger (2012b) if the decision maker follows RDU (Epper and Fehr-Duda, 2014a).

Section 2.12.2 introduces the concept of *intertemporal risk aversion*. Section 2.12.3 lays out the agenda for future research for the study of the interaction between risk and time preferences (Epper and Fehr-Duda, 2014b). Some features of this agenda are as follows. Decision makers appear more risk tolerant for events in the more distant future. The decision making process, for instance, whether dated lotteries are reduced in one go or sequentially reduced, influences choices. The timing of the resolution of uncertainty matters too. The extent of discounting depends on the degree of riskiness of lotteries. Finally, behavior differs whether lotteries are reduced first in the time dimension or the risk dimension. Suppose that an observer observes the choices of a decision maker when outcomes are delayed, and when the probability that they will materialize, falls over time. Sections 2.12.4 and 2.12.5 consider the resulting effects on the observed probability weighting function in terms of factors such as its elevation, subproportionality, and elasticity (Epper and Fehr-Duda, 2014b).

2.2 Probability weighting functions

Under EU, decision makers weight probabilities linearly (see Remark 1.1). In contrast, empirical evidence suggests non-linear weighting of probabilities. There is fairly widespread agreement

in the empirical literature on the shape of the non-linear weighting for probabilities that are bounded away from the end points of the interval $[0, 1]$. However, there are unresolved issues about the precise shape of non-linear weighting close to the end points of the probability interval $[0, 1]$, an issue that only the more recent literature has focused on. To separate these issues, and based on the empirical evidence, it is useful to consider the following two hypotheses, H1 and H2. Most behavioral decision theory proceeds under the assumption of H1. However, a more nuanced picture is given in H2.

H1: Decision makers overweight small probabilities and underweight large probabilities, no matter how small or how large they are.⁷

H2: There are two kinds of decision makers:

(H2a) A fraction $\mu \in [0, 1]$ of decision makers (i) ignore events of extremely low probability and, (ii) treat extremely high probability events as certain.⁸ For all other (non-extreme) probabilities they behave as in H1, i.e., they overweight small probabilities and underweight large ones.

(H2b) A fraction $1 - \mu$ of the decision makers behave as in H1 for all levels of probabilities.

Classical non-EU theories largely subscribe to H1 (or assume that $\mu = 0$ in hypothesis H2). In contrast, hypothesis H2 has been largely ignored despite a compelling case for it in Kahneman and Tversky's (1979) original prospect theory, and in the *bimodal perception of risks* model (Viscusi, 1998). We follow the chronological order of developments so, for the moment, we focus only on H1. In Section 2.11, we examine the evidence for H2 and outline al-Nowaihi and Dhami's (2010a) *composite prospect theory* (CPT) that is able to account for H2.

Recall that we have already conjectured that the Allais paradox can be resolved under non-linear probability weighting. As a compelling illustration of the importance of non-linear probability weighting, consider the following example from Kahneman and Tversky (1979 p. 283). Suppose that one is compelled to play Russian roulette using a revolver that can potentially fire six rounds, but all rounds are not loaded.⁹ One would be willing to pay much more to reduce the number of bullets from one to zero (elimination of risk) than from four to three (reduction, but not elimination of risk). However, in each case, the reduction in the probability of a bullet firing is $1/6$ and, so, under EU, the decision maker should be willing to pay the same amount. This suggests non-linear weighting of probabilities.

For any lottery L , let $E(L)$ be the expected value of the lottery and let C_L be the certainty equivalent. Then, we can define the *relative risk premium* as

$$RRP = \frac{E(L) - C_L}{|E(L)|}. \quad (2.1)$$

⁷ For the evidence see Kahneman and Tversky (2000), Starmer (2000), and Appendix A in Wakker (2001). For the necessary and sufficient conditions required for inverse S-shaped probability weighting functions, see Tversky and Wakker (1995). There is also emerging evidence of the neuro-biological foundations for such behavior (Berns et al., 2008); see Part 9 for a more detailed discussion.

⁸ In the context of the take-up of insurance for low probability natural hazards, the results from one set of experiments by Kunreuther et al. (1978) are consistent with $\mu = 0.8$.

⁹ Readers may remember the game of Russian roulette from the 1978 film, *The Deer Hunter*. In this potentially lethal game, fewer than 6 rounds in a revolver are loaded, the cylinder is spun, and the player places the muzzle against his/her head, and pulls the trigger.




Figure 2.1 Relative risk premium for the Student and the Representative population.

Source: Reproduced with permission of *Annual Review of Economics*, Volume 4, Helga Fehr-Duda and Thomas Epper, "Probability and risk: foundations and economic implications of probability-dependent risk preferences," © 2012 by Annual Reviews <http://www.annualreviews.org>

There are three possibilities: Risk aversion corresponds to $RRP > 0$, risk neutrality to $RRP = 0$, and risk seeking/loving to $RRP < 0$. Fehr-Duda and Epper (2012) compute the RRP for a student population recruited from two major Swiss Universities in 2006, and a representative population of Swiss German-speaking adults. Lotteries of the form $(x_1, 1 - p; x_2, p)$ are used, where $0 \leq x_1 < x_2$. The results are shown in Figure 2.1; relative risk premium, RRP , is plotted on the vertical axis and the probability of the highest outcome, p , is plotted on the horizontal axis. Clearly, the RRP is probability dependent. Thus, risk preferences of the decision maker depend on probabilities. This is inconsistent with EU because under EU, risk preferences arise solely through the shape of the utility function.

Definition 2.19 (*Probability weighting function*): By a probability weighting function, we mean a strictly increasing function $w : [0, 1] \xrightarrow{\text{onto}} [0, 1]$.

To fix ideas, readers may, for the moment, consider $w(p)$ as the subjective weight placed by decision makers on the objective probability, p (these ideas carry over to non-negative subjective probabilities that sum up to 1). The proof of the next proposition is left as an exercise.

Proposition 2.8 A probability weighting function has the following properties:

(a) $w(0) = 0$, $w(1) = 1$. (b) w has a unique inverse, w^{-1} , and w^{-1} is also a strictly increasing function from $[0, 1]$ onto $[0, 1]$. (c) w and w^{-1} are continuous.

In order to distinguish between the hypotheses H1 and H2, we will need to look at the behavior of alternative probability weighting functions for probabilities close to the endpoints of $[0, 1]$. In particular, we need to speak meaningfully about the steepness of the probability weighting function around $p = 0$ and $p = 1$. The next two definitions introduce the relevant concepts.




Figure 2.2 Non-parametrically estimated probability weighting functions for the Student and Representative population based on aggregate data.

Source: Reproduced with permission of *Annual Review of Economics*, Volume 4, Helga Fehr-Duda and Thomas Epper, "Probability and risk: foundations and economic implications of probability-dependent risk preferences," © 2012 by Annual Reviews <http://www.annualreviews.org>

Definition 2.20 (al-Nowaihi and Dhami, 2010a): A probability weighting function, $w(p)$, (i) infinitely overweights infinitesimal probabilities, if for $\gamma > 0$, $\lim_{p \rightarrow 0} \frac{w(p)}{p^\gamma} = \infty$, and (ii) infinitely underweights near-one probabilities if $\lim_{p \rightarrow 1} \frac{1-w(p)}{1-p} = \infty$.¹⁰

Definition 2.21 (al-Nowaihi and Dhami, 2010a): A probability weighting function, $w(p)$, (i) zero-underweights infinitesimal probabilities, if for $\gamma > 0$, $\lim_{p \rightarrow 0} \frac{w(p)}{p^\gamma} = 0$, and (ii) zero-overweights near-one probabilities in the sense that $\lim_{p \rightarrow 1} \frac{1-w(p)}{1-p} = 0$.

For aggregate data (and for non-extreme probabilities), the data typically suggests an *inverse S-shaped probability weighting function*. Figure 2.2 shows the estimated probability weighting function using non-parametric estimates corresponding to the data in Figure 2.1.

2.2.1 Prelec's probability weighting function

A large number of probability weighting functions have been proposed in the literature. The most popular of these functions is the Prelec (1998) function. The Prelec function is parsimonious, tractable, and has axiomatic foundations. Depending on parameter values, we will see that it is either consistent with H1 or the special case of H2 with $\mu = 1$. We now turn to a detailed

¹⁰ In particular, for $\gamma = 1$, Definition 2.20(i) implies that $\lim_{p \rightarrow 0} \frac{w(p)}{p} = \infty$.

examination of the Prelec weighting function, reminding readers that the comments we make here apply to most probability weighting functions (the so-called *standard weighting functions*, which we define, below).¹¹ Furthermore, an understanding of the Prelec function will be critical to appreciating the recent developments that take account of hypothesis H2; see Section 2.11.3.

Definition 2.22 (Prelec, 1998): By the *Prelec function*, we mean the probability weighting function $w(p) : [0, 1] \rightarrow [0, 1]$ given by

$$w(0) = 0, w(1) = 1, \quad (2.2)$$

$$w(p) = e^{-\beta(-\ln p)^\alpha}, 0 < p \leq 1, \alpha > 0, \beta > 0. \quad (2.3)$$

Sometimes the special case of $\beta = 1$ is known as *Prelec-I* and the case of $\beta \neq 1$ is known as *Prelec-II* to distinguish between the one-parameter and the two-parameter family of distributions.

The Prelec function only requires that $\alpha > 0, \beta > 0$. In order to distinguish between H1 and H2, it is critical whether α is smaller or larger than 1. For this reason, we make a distinction between the *Prelec function* and the *standard Prelec function*; Prelec (1998) prefers the latter for reasons we give below.

Definition 2.23 (Standard Prelec function): By the *standard Prelec probability weighting function*, we mean the Prelec function, defined in (2.2) and (2.3), but with $0 < \alpha < 1$.

Remark 2.2 For $\alpha = \beta = 1$ we get $w(p) = p$ as in EU. Since an infinite number of α, β values are possible and just one pair of values coincides with the special case of EU, it would be quite remarkable if human behavior could be represented by just that one pair of values. Indeed, the evidence clearly rules out this possibility. It is even more remarkable that many economists still stubbornly wish to cling on to EU.

A simple proof that is omitted gives rise to the following result.

Proposition 2.9 The Prelec function (Definition 2.22) satisfies the properties listed in Proposition 2.8.

The roles of α, β in the Prelec weighting function are highlighted next.

1. (Role of α) The parameter α controls the convexity/concavity of the Prelec function. If $\alpha < 1$, then the Prelec function is strictly concave for low probabilities but strictly convex for high probabilities. In this case, it is *inverse S-shaped*, as in the curve $w(p) = e^{-(-\ln p)^{0.5}}$ (i.e., $\alpha = 0.5, \beta = 1$), which is sketched as the *thick curve* in Figure 2.3. The converse holds if $\alpha > 1$, in which case the Prelec function is *S-shaped*. An example is the curve $w(p) = e^{-(-\ln p)^2}$ ($\alpha = 2, \beta = 1$), sketched in Figure 2.3 as the *light curve*. The straight line in Figure 2.3 is the 45° line $w(p) = p$ ($\alpha = \beta = 1$) corresponding to the case of EU.
2. (Role of β) Between the region of strict convexity ($w'' > 0$) and the region of strict concavity ($w'' < 0$), there is a point of inflexion ($w'' = 0$). The parameter β controls the location of the inflexion point, relative to the 45° line. Thus, for $\beta = 1$, the point of inflexion is at $p = e^{-1}$ and lies on the 45° line, as in Figure 2.3. However, if $\beta < 1$, then the point of inflexion lies above the 45° line, as in the curve $w(p) = e^{-0.5(-\ln p)^2}$ ($\alpha = 2, \beta = 0.5$).

¹¹ The interested reader can also consult al-Nowaihi and Dhami (2011) for more details.




Figure 2.3 Plots of $w(p) = e^{-(\ln p)^{0.5}}$ and $w(p) = e^{-(\ln p)^2}$.

For this example, the fixed point, $w(p^*) = p^*$, is at $p^* \simeq 0.14$ but the point of inflexion, $w''(\tilde{p}) = 0$, is at $\tilde{p} \simeq 0.20$.

Some useful results about the Prelec function are given next. The reader is asked to prove them in the exercises.

Proposition 2.10 For $\alpha = 1$, the Prelec probability weighting function (Definition 2.22) takes the form $w(p) = p^\beta$, is strictly concave if $\beta < 1$, but strictly convex if $\beta > 1$. In particular, for $\alpha = \beta = 1$, $w(p) = p$ (as under EU).

Proposition 2.11 (al-Nowaihi and Dhami, 2010a): Suppose that $\alpha \neq 1$. Then, the Prelec function has the following properties.

- (a) It has exactly three fixed points: 0, $p^* = e^{-\left(\frac{1}{\beta}\right)^{\frac{1}{\alpha-1}}}$, and 1. For $\beta = 1$, $p^* = e^{-1}$.
- (b) It has a unique inflexion point, $\tilde{p} \in (0, 1)$ at which $w''(\tilde{p}) = 0$.
- (c) If $\alpha < 1$, it is strictly concave for $p < \tilde{p}$ and strictly convex for $p > \tilde{p}$ (inverse S-shaped).
- (d) If $\alpha > 1$, it is strictly convex for $p < \tilde{p}$ and strictly concave for $p > \tilde{p}$ (S-shaped).
- (e) The inflexion point, $\tilde{p}$, lies above, on, or below, the 45° line, respectively, if β is less than, equal to, or greater than one.










Table 2.1, graphs the Prelec function $w(p) = e^{-\beta(-\ln p)^\alpha}$, for various values of α, β .

Corollary 2.1

- (i) Suppose that $\alpha \neq 1$. Then $\tilde{p} = p^* = e^{-1}$ (i.e., the point of inflexion, $\tilde{p}$, and the fixed point, p^* , coincide) if, and only if, $\beta = 1$.
- (ii) If $\beta = 1$, then:
 - (a) If $\alpha < 1$, then w is strictly concave for $p < e^{-1}$ and strictly convex for $p > e^{-1}$ (inverse S-shape; see the thick curve in Figure 2.3).
 - (b) If $\alpha > 1$, then w is strictly convex for $p < e^{-1}$ and strictly concave for $p > e^{-1}$ (S-shape; see the light curve in Figure 2.3).

For the thick curve in Figure 2.3 (and first row in Table 2.1), where $\alpha < 1$, the slope of $w(p)$ becomes very steep near $p = 0$. By contrast, for the thin curve in Figure 2.3 (and last row in Table 2.1), where $\alpha > 1$, the slope of $w(p)$ becomes shallow near $p = 0$; in this case low

Table 2.1 Representative graphs of the Prelec function.

	$\beta = \frac{1}{2}$	$\beta = 1$	$\beta = 2$
$\alpha = \frac{1}{2}$	 $w(p) = e^{-\frac{1}{2}(-\ln p)^{\frac{1}{2}}}$	 $w(p) = e^{-(-\ln p)^{\frac{1}{2}}}$	 $w(p) = e^{-2(-\ln p)^{\frac{1}{2}}}$
$\alpha = 1$	 $w(p) = p^{\frac{1}{2}}$	 $w(p) = p$	 $w(p) = p^2$
$\alpha = 2$	 $w(p) = e^{-\frac{1}{2}(-\ln p)^2}$	 $w(p) = e^{-(-\ln p)^2}$	 $w(p) = e^{-2(-\ln p)^2}$

Source: al-Nowaihi and Dhami (2011).

probabilities are underweighted. We now tie these observations with the notion of steepness of a probability weighting function at extreme probabilities (Definitions 2.20, 2.21).

Proposition 2.12

- (a) For $\alpha < 1$ and $\gamma > 0$, the Prelec (1998) function: (i) infinitely overweights infinitesimal probabilities, i.e., $\lim_{p \rightarrow 0} \frac{w(p)}{p^\gamma} = \infty$, and (ii) infinitely underweights near-one probabilities, i.e., $\lim_{p \rightarrow 1} \frac{1-w(p)}{1-p} = \infty$.
- (b) For $\alpha > 1$ and $\gamma > 0$, the Prelec function: (i) zero-underweights infinitesimal probabilities, i.e., $\lim_{p \rightarrow 0} \frac{w(p)}{p^\gamma} = 0$, and (ii) zero-overweights near-one probabilities, i.e., $\lim_{p \rightarrow 1} \frac{1-w(p)}{1-p} = 0$.

We have made a distinction between the *Prelec function* ($\alpha > 0$) and the *standard Prelec function* ($0 < \alpha < 1$). Prelec (1998) himself favors the case $0 < \alpha < 1$, and it is this case which is standard in non-EU theories such as rank dependent utility and prospect theory. According to Prelec (1998, p. 505), for the case $\gamma = 1$, the infinite limit in Proposition 2.12(a) captures the qualitative change as we move from impossibility to improbability. On the other hand, this contradicts hypothesis H2a, i.e., the observed behavior that many people simply ignore events of very low probability. These problems are avoided for $\alpha > 1$. However, for $\alpha > 1$, the Prelec function is S-shaped, see Proposition 2.11(d) and Figure 2.3, which conflicts with hypotheses H1

and H2b. This suggests that we need a richer account of decision making under risk for events when near-zero and near-one probabilities are involved; *composite prospect theory*, outlined in Section 2.11.3, is one attempt to give such an account.

When aggregate data is taken into account, most weighting functions appear to be inverse S-shaped. However, there is substantial heterogeneity at the individual level with the two most common shapes being convex and inverse S-shape.¹² Fehr-Duda and Epper (2012) give estimates of the Prelec function for two different populations—a student population and a representative population (see description, above). For the student population, $\alpha = 0.513$ and $\beta = 0.958$, and for the representative population, $\alpha = 0.423$ and $\beta = 0.868$.

Definition 2.24 (*Subproportionality and certainty effect*): A feature of most probability weighting functions, at least over a range $[\hat{p}, 1]$, $\hat{p} \geq 0$, is subproportionality, i.e.,

$$\frac{w(p)}{w(q)} > \frac{w(\lambda p)}{w(\lambda q)}, 0 < \lambda < 1, \hat{p} \leq q < p \leq 1. \quad (2.4)$$

If $\hat{p} = 0$ and $0 < q$, then w is subproportional over the entire probability range. The certainty effect arises when $p = 1$. In this case, we have

$$w(\lambda q) > w(\lambda)w(q).$$

Probability weighting functions that have the feature of subproportionality, can typically explain the common ratio effect. In addition, inverse S-shaped probability weighting functions that cut the diagonal from above can explain the common consequence effect and the simultaneous existence of insurance and gambling.¹³

There are other probability weighting functions also; we now briefly consider the main ones.

OTHER PROBABILITY WEIGHTING FUNCTIONS

1. The Tversky and Kahneman (1992) probability weighting function is

$$w(p) = \frac{p^\tau}{[p^\tau + (1-p)^\tau]^{\frac{1}{\tau}}}, 0.279 < \tau < 1. \quad (2.5)$$

This probability weighting function is not subproportional for small values of p , hence, it cannot account for the common ratio effect for small probabilities. A smaller value of τ induces a more pronounced departure from linearity. One potential drawback of this function is that the point of intersection with the diagonal is not independent of the curvature parameter, τ (recall that this problem does not arise with the Prelec function).

2. Karmakar (1979) proposed the following weighting function

$$w(p) = \frac{p^\tau}{p^\tau + (1-p)^\tau}, 0 < \tau < 1. \quad (2.6)$$

This function is not subproportional for low values of p but it intersects the diagonal at $p = 0.5$, which is independent of τ (unlike the Tversky–Kahneman function).

¹² See, for instance, Stott (2006), Gonzalez and Wu (1999), van de Kuilen and Wakker (2011), Fehr-Duda and Epper (2012).

¹³ Friedman and Savage (1948) proposed a utility function with concave and convex segments to explain the simultaneous occurrence of insurance and gambling but this does not explain why risk preferences of individuals are sensitive to probabilities at fairly different levels of payoffs.

3. The Goldstein and Einhorn (1987) linear-in-log-odds probability weighting function is given by

$$w(p) = \frac{\delta p^\eta}{\delta p^\eta + (1-p)^\eta}, p \in [0, 1], \delta \geq 0, \eta \geq 0. \quad (2.7)$$

In this specification, η is responsible for the curvature and a smaller value of η induces a more pronounced departure from the EU case, $w(p) = p$. The parameter δ controls elevation; an increase in δ increases the elevation of the weighting function. The standard EU case corresponds to $\delta = \eta = 1$.

Fehr-Duda and Epper (2012) argue that the one parameter weighting functions do not fit the data as well as the two-parameter family of functions. Gonzalez and Wu (1999) show that the Prelec function in (2.3) and the Goldstein–Einhorn function in (2.7) fit the data equally well.

2.2.2 Stochastic dominance under non-linear probability weighting

If hypotheses H1 and H2 are important, then it may seem natural to replace the linear weighting under EU with non-linear weighting. However, the next example shows that this is not straightforward.

Example 2.6 (Simple non-linear probability weighting): One alternative to EU is the decision weighted form of utility in which we define:

$$W(x_1, p_1; \dots; x_n, p_n) = \sum_{i=1}^n w(p_i) u(x_i), \quad (2.8)$$

where $w(p_i)$ is the probability weight of outcome x_i . Early attempts at formulating these kinds of utility functions include Edwards (1954, 1962) and Handa (1977). One problem with this utility function is that the choices generated by it might be stochastically dominated. This is shown in Example 2.7 (Fishburn, 1978).

Example 2.7 (Problems with non-linear weighting of probabilities): Consider the lottery $L = (x, p; y, 1-p)$, $0 < p < 1$. Let $w : [0, 1] \rightarrow [0, 1]$ be a probability weighting function. Let u be a utility function for sure outcomes and let $\tilde{U}$ be the utility function in (2.8). By definition,

$$\tilde{U}(L) = w(p) u(x) + w(1-p) u(y).$$

Let preferences induced by $\tilde{U}$ be denoted by $\prec_{\tilde{U}}$

- (a) For the special case, $x = y$, we get $\tilde{U} = [w(p) + w(1-p)] u(x)$. In general, for non-EU theories, $w(p) + w(1-p) \neq 1$. For probabilities bounded away from the end points of the interval $[0, 1]$, we typically have $w(p) + w(1-p) < 1$.¹⁴ Hence, $(x, p; x, 1-p) \prec_{\tilde{U}} (x, 1)$. But any “sensible” theory of risk should give $(x, p; x, 1-p) \sim_{\tilde{U}} (x; 1)$.
- (b) Take $y = x + \epsilon$, $\epsilon > 0$. Then, first order stochastic dominance for discrete outcomes (Definition 2.9) implies that $(x, 1) \prec_1 (x, p; x + \epsilon, 1-p)$. Now, $\tilde{U}(x, p; x + \epsilon, 1-p) = w(p) u(x) + w(1-p) u(x + \epsilon)$. Assuming u is continuous and $w(p) + w(1-p) < 1$, we get

¹⁴ For instance, for the Prelec probability weighting function, and for values of the parameters $\alpha = 0.5$ and $\beta = 1$ the statement $w(p) + w(1-p) < 1$ is true for all probabilities in the range $[0.03, 0.97]$.

$$\begin{aligned}\lim_{\epsilon \rightarrow 0} \tilde{U}(x, p; x + \epsilon, 1 - p) &= w(p) u(x) + w(1 - p) u(x) \\ &= [w(p) + w(1 - p)] u(x) < \tilde{U}(x).\end{aligned}$$

Hence, for sufficiently small $\epsilon > 0$, $\tilde{U}(x, p; x + \epsilon, 1 - p) < \tilde{U}(x)$. Thus, we have that $(x, 1) \prec_1 (x, p; x + \epsilon, 1 - p)$ but $\tilde{U}(x, p; x + \epsilon, 1 - p) < \tilde{U}(x, 1)$. Hence, first order stochastic dominance is violated. We discuss the approaches taken to deal with this problem in various non-EU theories, below.

2.3 Rank dependent utility theory (RDU)

To summarize the story so far, that is relevant for this section: Hypothesis H1, above, highlighted the tendency for decision makers to overweight small probabilities and underweight large probabilities. EU is unable to take account of H1 because it weights probabilities in a linear manner. Example 2.7 captures the early attempts to introduce non-linear weighting, by transformation of *point probabilities*. However, decision makers could choose stochastically dominated options in these theories, which was not attractive.

Rank dependent utility theory (RDU), relying on the work of Quiggin (1982, 1993), proposed the first satisfactory solution to the problem of non-linear probability weighting. Unlike Example 2.7, Quiggin suggested transformation of *cumulative probabilities*. EU and RDU are otherwise similar, except for cumulatively transformed probabilities. This was a major advance over EU, because it could resolve several paradoxes of EU, e.g., the Allais paradox. Further, decision makers who use RDU do not choose stochastically dominated options.

We now introduce *decision weights* that play a key role in most non-EU theories.

Definition 2.25 (*Decision weights*): Consider the lottery $(x_1, p_1; x_2, p_2; \dots; x_n, p_n)$, where $x_1 < x_2 < \dots < x_n$.¹⁵ Let w be the probability weighting function. For RDU, the decision weights, π_i , are defined as follows.

$$\begin{aligned}\pi_n &= w(p_n), \\ \pi_{n-1} &= w(p_{n-1} + p_n) - w(p_n), \\ &\dots \\ \pi_i &= w\left(\sum_{j=i}^n p_j\right) - w\left(\sum_{j=i+1}^n p_j\right), \\ &\dots \\ \pi_1 &= w\left(\sum_{j=1}^n p_j\right) - w\left(\sum_{j=2}^n p_j\right) = w(1) - w\left(\sum_{j=2}^n p_j\right) = 1 - w\left(\sum_{j=2}^n p_j\right).\end{aligned}$$

From Definition 2.25, we get that,

$$\pi_j \geq 0 \text{ and } \sum_{j=1}^n \pi_j = 1. \quad (2.9)$$

The logic and intuition for decision weights in Definition 2.25 is made clear later in this section.

Definition 2.26 Consider the lottery $L = (x_1, p_1; x_2, p_2; \dots; x_n, p_n)$, where $x_1 < x_2 < \dots < x_n$. Let w be the probability weighting function. Let π_i , $i = 1, 2, \dots, n$, be given by Definition 2.25. The decision maker's rank dependent utility from L , is given by

¹⁵ If any outcomes in the original lottery are equal, we sum up the associated probabilities and write lotteries in this format. For instance, the lottery $L = (4, 0.3; 4, 0.4; 6, 0.3)$ is equivalent to the lottery $(4, 0.7; 6, 0.3)$. If all outcomes are equal then $n = 1$.

$$RDU(L) = \sum_{i=1}^n \pi_i u(x_i). \quad (2.10)$$

From (2.10), an implication of models of RDU is that two outcomes that differ in rank but have the same probability will typically have different decision weights assigned to them.

2.3.1 Attitudes to risk under RDU

Under EU, attitudes to risk are captured by the shape of the utility function. For instance, a concave utility function implies and is implied by risk aversion. Typically, for psychologists, and many economists, the notion that risk is captured by the shape of the utility function alone, does not seem a palatable one. Indeed, it seems natural to associate the notion of riskiness also with probabilities. An important implication of RDU is that the utility function and the decision maker's perception of objective probabilities jointly determine attitudes to risk.¹⁶

Example 2.8 Suppose that a decision maker is offered the lottery $L = (0, 1 - p; 1, p)$, $0 < p < 1$, so the expected value is $EL = p$. Suppose that the decision maker follows RDU, the weighting function is convex, say, $w(p) = p^2$, and the utility function is linear, i.e., $u(x) = x$. Under EU, such a decision maker is risk neutral. Using Definition 2.26, under RDU, $RDU(L) = w(p) = p^2$. Since, $u(x) = x$, the certainty equivalent of the lottery under RDU, C_L , equals p^2 . Risk aversion requires that the certainty equivalent of a lottery is lower than the expected value of the lottery for all lotteries, hence,

$$C_L < EL \Leftrightarrow p^2 < p \Leftrightarrow p < 1,$$

which is true. In other words, the finding of risk aversion can equally well be explained by a linear utility function and a convex probability weighting function. In this example, the decision maker would be risk loving if $w(p) = \sqrt{p}$, which illustrates the importance of the probability weighting function in determining risk attitudes under RDU.

The discussion, below, shows that decision weights in Definition 2.25 have a simple, intuitive, explanation.

Example 2.9 (Curvature of weighting function and optimism, pessimism): The curvature of w is also sometimes interpreted in terms of the optimism or pessimism of the decision maker; these terms have no natural analogue in EU. To see this, assume that w is convex throughout. Suppose that we have the balanced risk lottery $L = (x_1, 0.5; x_2, 0.5)$; $x_1 < x_2$. Using Definition 2.26, under RDU

$$RDU(L) = u(x_1) [1 - w(0.5)] + u(x_2) w(0.5). \quad (2.11)$$

Since $w(0) = 0$, $w(1) = 1$ (see Definition 2.19) and w is assumed to be convex throughout, so $w(0.5) < 0.5$. Thus, in (2.11), the weight on $u(x_1)$ is higher than the weight on $u(x_2)$ despite identical objective probabilities. In other words, the decision maker is pessimistic in placing a higher weight on the lower outcome. The converse (optimism) would apply if w were concave throughout; in this case, the decision maker puts a higher decision weight on the higher outcome, despite both outcomes being equiprobable.

¹⁶ See Chateaufneuf and Cohen (1994) for an early theoretical demonstration of this idea.




Figure 2.4 An illustration of optimism and pessimism as the shape of the probability weighting function varies.

Let us delve a bit further into the relation between decision weights and probability weights. Consider the lottery $(30, 1/3; 60, 1/3; 90, 1/3)$. Using Definition 2.25, the decision weights corresponding to the three outcomes are

$$\pi(30) = 1 - w(2/3); \pi(60) = w(2/3) - w(1/3); \pi(90) = w(1/3). \quad (2.12)$$

Figure 2.4 plots a convex weighting function (LHS panel) and a concave weighting function (RHS panel). The magnitudes of the decision weights corresponding to each of the three outcomes, as calculated in (2.12), are marked on the vertical axis. For the convex case, the magnitudes of the decision weights decrease as the outcome increases, $\pi(90) < \pi(60) < \pi(30)$; this reflects *pessimism*. If such a weighting function is combined with a concave utility function, it enhances risk aversion. The converse happens when the weighting function is concave. Here larger outcomes are given greater decision weights, $\pi(90) > \pi(60) > \pi(30)$; this reflects *optimism*. Combined with a concave utility function, optimism reduces risk aversion.

Example 2.10 (*Elevation of probability weighting function and optimism*): Consider a decrease in the elevation of the probability weighting function, e.g., an increase in β in the Prelec function (Definition 2.22). This change is associated with a lower degree of optimism. To see this, consider the lottery $L = (x_1, 1 - p; x_2, p)$, $x_1 < x_2$. The rank dependent utility of this lottery is given by

$$RDU(L) = w(p) [u(x_2) - u(x_1)] + u(x_1).$$

Following a decrease in elevation, $w(p)$ decreases, which reduces RDU by reducing the decision weight that the decision maker puts on the higher outcome. A similar insight holds under prospect theory if all outcomes are in the domain of gains.

What is the shape of decision weights when there are several outcomes and we have an inverse S-shaped probability weighting function? Following simulations reported in Fehr-Duda and Epper (2012), the typical shape is shown in Figure 2.5. We have 21 ranked outcomes with a uniform probability distribution over them, plotted for two weighting functions—the Prelec-II function




Figure 2.5 Shapes of decision weights when there are multiple outcomes under RDU. Panel (a) shows the Prelec-II function and Panel (b) shows the Goldstein–Einhorn function. + denotes the uniform objective probability distribution. ● denotes the case $\alpha = \eta = 0.2$, $\beta = \delta = 1$. ▲ denotes the case $\alpha = \eta = 0.9$, $\beta = \delta = 1$. ■ denotes the case $\alpha = \eta = 0.2$, $\beta = \delta = 0.3$.

Source: Reproduced with permission of *Annual Review of Economics*, Volume 4, Helga Fehr-Duda and Thomas Epper, "Probability and risk: Foundations and economic implications of probability-dependent risk preferences," © 2012 by Annual Reviews <http://www.annualreviews.org>

and the Goldstein–Einhorn function (see Definition 2.22 and (2.7)); recall that $0 < \alpha < 1$ gives the inverse S-shape under the Prelec function, and a similar role is played by the parameter η in the Goldstein–Einhorn function. Extreme outcomes are accorded extreme decision weights, while the intermediate outcomes are accorded relatively flat weights. This enhances the salience of extreme outcomes, a feature that explains several puzzles in finance (see Section 3.9).

Remark 2.3 Several axiomatizations of RDU allow for additive separability between terms that multiply utility with probability weights, e.g., Abdellaoui (2002) and Wakker (1994). In order to ensure concavity or convexity of the weighting function, these contributions introduce additional global auxiliary conditions such as probabilistic risk aversion. See in particular, Wakker (2001) who allows for a preference axiomatization of convex capacities. This allows for a weighting function that is either concave or convex throughout. Diecidue et al. (2009) introduce local conditions that allow for probabilistic risk aversion to vary over the probability interval $[0, 1]$. For one set of parameter values, they can address hypothesis H1, while for the complementary set of parameter values, they can address hypothesis H2a. However, their proposal was not designed with a view to simultaneously addressing H1 and H2.

A decision maker who follows RDU never chooses stochastically dominated lotteries. We state here the formal result without proof.¹⁷ Let $x_1 < x_2 < \dots < x_n$ be a set of outcomes. Consider two lotteries, L_p, L_q , corresponding to two different distribution functions P, Q over the outcomes,

$$L_p = (x_1, p_1; x_2, p_2; \dots; x_n, p_n), \quad L_q = (x_1, q_1; x_2, q_2; \dots; x_n, q_n).$$

¹⁷ Formal proofs are surprisingly hard to find. The formal results below draw on unpublished joint research conducted by Ali al-Nowaihi and Sanjit Dhami.

Suppose that the decision weights (Definition 2.25) for the lotteries L_p, L_q are denoted respectively by π_{p_i} and π_{q_i} , $i = 1, 2, \dots, n$.

Let W be the set of probability weighting functions, i.e.,

$$W = \{w : [0, 1] \rightarrow [0, 1], w(0) = 0, w(1) = 1, w \text{ is non-decreasing}\}. \quad (2.13)$$

Definition 2.27 Under RDU, the lottery, L_p , is preferred to another lottery, L_q , written $L_q \preceq_{RDU} L_p$, if and only if

$$\sum_{i=1}^n \pi_{q_i} u(x_i) \leq \sum_{i=1}^n \pi_{p_i} u(x_i), \text{ for all } u \in \mathbf{u} \text{ and for all } w \in W,$$

where $\mathbf{u}$ is the class of all non-decreasing utility functions. If, in addition, the inequality is strict for some $u \in \mathbf{u}$, then we say that L_p is strictly preferred to L_q under RDU. This is written as $L_q \prec_{RDU} L_p$.

Definition 2.28 We define first order stochastic dominance under RDU, $\preceq_1^{RDU}$, as follows.

$$Q \preceq_1^{RDU} P \Leftrightarrow \text{for all } u \in \mathbf{u} \text{ and all } w \in W, L_q \preceq_{RDU} L_p. \quad (2.14)$$

Proposition 2.13 (a) $Q \preceq_1^{RDU} P \Leftrightarrow P_i \leq Q_i$, for all $i = 1, 2, \dots, n$.

(b) $Q \preceq_1^{RDU} P \Leftrightarrow Q \preceq_1^{EU} P$, where $\preceq_1^{EU}$ is the analogue of the first order stochastic dominance relation in Definition 2.28 when the decision maker follows expected utility.

2.3.2 RDU under uncertainty

Recall our discussion of subjective expected utility (SEU) under uncertainty in Section 1.3. We now consider RDU under uncertainty. Consider the Savage act (see Section 1.3 for the definition),

$$f = (E_1, x_1; E_2, x_2; \dots; E_n, x_n),$$

where $\cup_{i=1}^n E_i = S$, the sample space. Suppose also that the outcomes are ordered as $x_1 < x_2 < \dots < x_n$. We briefly describe how Schmeidler (1989) formulated a more general rank dependent theory to deal with uncertainty (and not just risk, which was Quiggin's (1982) contribution).

Definition 2.29 A subjective probability weighting function (or event weighting function), $\tilde{w} : E \rightarrow \mathbb{R}$ where $E \subseteq S$; $\tilde{w}$ has the following properties:

- (1) $\tilde{w}(\emptyset) = 0$, where $\emptyset$ is the empty set.
- (2) If $E_i \subset E_j$, then $\tilde{w}(E_i) \leq \tilde{w}(E_j)$.
- (3) $\tilde{w}(S) = 1$.

$\tilde{w}$ need not be additive, i.e., for mutually exclusive sets E_i, E_j , $\tilde{w}(E_i \cup E_j)$ need not equal $\tilde{w}(E_i) + \tilde{w}(E_j)$.

The objective function for rank dependent utility under uncertainty is defined as:

$$U(f) = \sum_{j=1}^n \pi_j u(x_j).$$

The decision weights under uncertainty, π_j , are defined as follows:

$$\pi_j = \tilde{w}(E_j \cup E_{j+1} \cup \dots \cup E_n) - \tilde{w}(E_{j+1} \cup E_{j+2} \cup \dots \cup E_n).$$

RDU under risk is a special case of Schmeidler's (1989) framework when the probability of event E_j equals its objective probability and probabilistic sophistication (Definition 2.13) holds.

The Ellsberg paradox (see Chapter 4) can be explained under RDU if one relaxes the assumption of probabilistic sophistication. The assumption of probabilistic sophistication allows us to write

$$\tilde{w}(E) = w(p(E)), \quad (2.15)$$

where w is the probability weighting function under risk and $p(E)$ denotes the objective probability of event E . The non-linearity of $\tilde{w}$ can now be seen immediately. Let E_i, E_j be mutually exclusive events, then, under probabilistic sophistication:

$$\begin{aligned} \tilde{w}(E_i \cup E_j) &= w(p(E_i \cup E_j)) \\ &= w(p_i + p_j) \\ &\neq w(p_i) + w(p_j) = \tilde{w}(E_i) + \tilde{w}(E_j). \end{aligned}$$

The first line is true by definition, the second follows because E_i and E_j are mutually exclusive events and the third because w is non-linear. In the discussion of RDU under risk, we showed how various shapes of the probability weighting function, w , can denote optimism or pessimism. Similar notions hold for the event weighting function, $\tilde{w}$. For instance, pessimism can be shown to be equivalent to the inequality $\tilde{w}(E_i \cup E_j) \geq \tilde{w}(E_i) + \tilde{w}(E_j)$ (see exercise 10.4.2 in Wakker, 2010).

The most interesting aspect of the formulation of rank dependent utility under uncertainty, aside from its generality, is that the empirical evidence on the shape of w carries over to the shape of $\tilde{w}$. Indeed, the inverse S-shape is even more vivid under $\tilde{w}$. We omit a discussion of the empirical findings. The interested reader can consult Hogarth and Kunreuther (1989), Kilka and Weber (2001), and Abdellaoui et al. (2005).

2.3.3 Drawbacks of RDU

RDU offers a distinct improvement over EU. Since EU is a special case of RDU when $w(p) = p$, RDU explains everything that EU does but the converse is false. RDU is able to explain the Allais paradoxes. An exercise asks you to do the relevant calculations. However, RDU cannot explain all the facts, as the next example shows.

Example 2.11 (A refutation of both EU and RDU): Consider Problems 11 and 12 from Kahneman and Tversky (1979).

Problem 11. In addition to whatever you own (x_0), you have been given 1000. You are now asked to choose between

$$A = (x_0 + 1000, 0.5; x_0 + 2000, 0.5) \text{ and } B = (x_0 + 1000, 0; x_0 + 1500, 1).$$

A is chosen by 16%, while B is chosen by 84%. Using Definition 2.26, under RDU, the evidence, thus, suggests that for most subjects,

$$w(0.5)u(x_0 + 2000) + [1 - w(0.5)]u(x_0 + 1000) < u(x_0 + 1500). \quad (2.16)$$

Problem 12: In addition to whatever you own (x_0), you have been given 2000. You are now asked to choose between

$$C = (x_0 + 1000, 0.5; x_0 + 2000, 0.5) \text{ and } D = (x_0 + 1500, 1; x_0 + 2000, 0)$$

C is chosen by 69% of the subjects, while *D* is chosen by 31%. Using Definition 2.26, under RDU, the empirical evidence indicates that for most subjects,

$$w(0.5)u(x_0 + 2000) + [1 - w(0.5)]u(x_0 + 1000) > u(x_0 + 1500). \quad (2.17)$$

Clearly, (2.16) contradicts (2.17), violating RDU. The main reason for this violation is that RDU does not take account of shifts in the “reference point” of the decision maker, when the endowment changes from $x_0 + 1000$ to $x_0 + 2000$. We shall see below in Example 2.14 that prospect theory, which takes account of such changes in reference points, can account for this puzzle.

RDU is unable to explain many important phenomena, whose explanation rests on features that are central to prospect theory such as *loss aversion* and *reference dependence*. Examples include explanations of the equity premium puzzle, tax evasion puzzles, disposition effect, backward bending labor supply curve, endowment effect, asymmetric price elasticities, and exchange asymmetries. These stylized facts are amenable to explanation based on prospect theory.

2.4 Prospect theory (PT)

Remark 1.1, above, listed the two main features of EU: linearity in probabilities and utility that does not differ in the domain of gains and losses. RDU relaxed the first feature by introducing non-linear probability weighting. Prospect theory goes a step further and relaxes both features of EU. There are two main versions of prospect theory.¹⁸ *Original prospect theory* (OPT) is due to Kahneman and Tversky (1979) and *cumulative prospect theory*, or simply *prospect theory* (PT) is due to Tversky and Kahneman (1992).¹⁹ Our focus here will be on PT. Section 2.11.1, below, gives a brief treatment of OPT.

Under PT (and OPT), the carriers of utility are not final levels of wealth, but deviations of actual wealth levels from a *reference point*, an idea that is well established in psychology. Evidence indicates that when exposed to external stimuli, such as temperature, brightness, and pain, individuals are more sensitive to *changes* rather than *levels* (Helson, 1964). In order to fix ideas on reference dependence, consider the following well-known experiment: Put your right hand in cold water and your left hand in warm water until both hands have had time to adjust to the temperature. Then, simultaneously take out both hands and dip them in lukewarm water. Although both hands experience the identical lukewarm temperature, the right hand feels warm and the left hand feels cold.

Strictly speaking, the view that the reference point is exogenous under PT needs to be qualified. Kahneman and Tversky were fully aware that the *status-quo* turns out to be a useful reference point in many applications. Other possible choices could be an *expected outcome*, a *fair outcome*, or a *legal entitlement*; clearly the context dictates the appropriate reference point. The reference point could also be state-dependent wealth, average wealth, desired wealth, or it could even be

¹⁸ See also *third generation prospect theory* in Section 2.7, and *composite prospect theory* in Section 2.11.3.

¹⁹ Starmer and Sugden (1989b) also introduced non-linear probability weights in a prospect theory framework. Although this is rarely mentioned, Tversky and Kahneman (1992) also extend prospect theory to ambiguity, using the framework suggested by Schmeidler (1989).

made endogenous by using the rational expectations of future wealth (see Section 2.8). For a discussion, see Kahneman and Tversky (2000), Köszegi and Rabin (2006), and Schmidt et al. (2008). Abeler et al. (2011) find that expectations of income serve as reference points and determine effort provision. Falk and Knell (2004) show how individuals might attempt to endogenize their reference points with a view to taking account of self-improvement and self-enhancement.

If the outcome turns out to be better (respectively worse) as compared to the reference point, then the decision maker is said to be in the *domain of gains* (respectively *losses*). Experimental and field evidence suggests that utility is evaluated differently in the domain of gains and losses. For instance, Example 1.5, due to Tversky and Kahneman (1981), demonstrated that decision makers are risk averse in the domain of gains and risk loving in the domain of losses. Using Australian data, Page et al. (2012) showed that following a natural disaster (floods), individuals who suffer a loss of at least 50% of their property values are more likely to accept risky gambles.

Extensive evidence, that we review below, also indicates that losses bite more than equivalent gains. This phenomenon, known as *loss aversion*, accounts for a large number of puzzles that could not be explained using EU and RDU. Finally, under PT (but not under OPT), decision makers transform probabilities non-linearly, in a manner similar to RDU, to produce decision weights. The difference from RDU is that, under PT, these transformations are applied separately in the domain of gains and losses.

We now give a more formal treatment of PT.²⁰ To take account of reference dependence, we first define lotteries in *incremental form*.

Definition 2.30 (*Lotteries in incremental form or prospects*) We say that a lottery is presented in incremental form if it is represented as:

$$L = (y_1, p_1; y_2, p_2; \dots; y_n, p_n), \quad (2.18)$$

where

$$y_i = x_i - x_0, \quad i = 1, 2, \dots, n, \quad (2.19)$$

is the increment (positive, negative, or zero) in wealth relative to reference wealth, x_0 . The reference wealth may be fixed at some given level, e.g., when wealth is simply the status-quo level, or it may be stochastic, as in the models described below in Sections 2.7 and 2.8. Lotteries represented in the incremental form are sometimes also known as “prospects.”

In PT, the *utility function* is termed as a *value function*. However, we speak of utility under PT in the normal sense that readers would understand the term in EU and RDU, and reserve the term value function for the objective function under PT.

Definition 2.31 (*Utility function under PT; Tversky and Kahneman, 1979*). Let $Y \subset \mathbb{R}$ be the set of wealth levels relative to a reference point, as in (2.19). A utility function, v , is a mapping $v : Y \rightarrow \mathbb{R}$ that satisfies the following:

1. v is continuous.
2. v is strictly increasing.
3. $v(0) = 0$ (reference dependence).
4. v is concave for $y \geq 0$ (declining sensitivity for gains).

²⁰ A feature of the ideas in OPT/PT is that they may appear self-evident. For an interesting account of the application of these ideas to Greek mythology, see Orwin (1994) and Ober (2012).

5. v is convex for $y \leq 0$ (declining sensitivity for losses).
6. $-v(-y) > v(y)$ for $y > 0$ (loss aversion).²¹

There are several alternative concepts of loss aversion in PT; these are discussed in Section 2.5.2. Tversky and Kahneman (1992) propose the following power form of the utility function.²²

$$v(y) = \begin{cases} y^\gamma & \text{if } y \geq 0 \\ -\lambda(-y)^\gamma & \text{if } y < 0 \end{cases}, \quad (2.20)$$

where γ and λ are constants such that $0 < \gamma < 1$, and $\lambda > 1$ is known as the *coefficient of loss aversion*. One may also impose separate powers in the domain of gains and losses, respectively, γ^+ and γ^- . Tversky and Kahneman (1992) estimated that $\gamma^+ \simeq \gamma^- = \gamma = 0.88$ and $\lambda \simeq 2.25$.

The properties in Definition 2.31, are illustrated in Figure 2.6, which is a plot of (2.20) for $\lambda = 2.5$, $\gamma = 0.5$.

Suppose that we have a total of $m + 1 + n$ outcomes such that there are m distinct outcomes in the domain of losses, n distinct outcomes in the domain of gains, and there is 1 reference outcome. Consider the following lottery in incremental form (Definition 2.30) that represents these outcomes and the associated probabilities

$$L = (y_{-m}, p_{-m}; y_{-m+1}, p_{-m+1}; \dots; y_{-1}, p_{-1}; y_0, p_0; y_1, p_1; y_2, p_2; \dots; y_n, p_n), \quad (2.21)$$

where y_i is defined in (2.19). The restriction on probabilities is given by

$$\sum_{i=-m}^n p_i = 1, p_i \geq 0, i = -m, -m+1, \dots, n. \quad (2.22)$$




Figure 2.6 The utility function under prospect theory.

²¹ Notice that this formulation makes sense because of properties 2 and 3 of the utility function.

²² For the axiomatic foundations of this utility function, see Section 2.6 below.

It is pedagogically more convenient to refer to the y_i values as “outcomes” rather than using the more cumbersome but accurate expression “outcomes relative to the reference point.” We shall follow this convention, below. The restriction on outcomes is given by

$$y_{-m} < y_{-m+1} < \dots < y_{-1} < y_0 = 0 < y_1 < y_2 < \dots < y_n, \quad (2.23)$$

where $y_0 = 0$ is the *reference outcome* (the reference point is x_0). Thus, in PT, outcomes are ordered from worst (y_{-m}) to best (y_n).

Definition 2.32 (*Set of lotteries*): Denote by $\mathcal{L}_P$, the set of all lotteries of the form given in (2.21), subject to the restrictions (2.22), (2.23).

In order to compute the decision weights for gains under PT, we apply the usual RDU calculations, as in Definition 2.25, starting at the positive extreme and using a probability weighting function, w^+ , for the domain of gains. However, for losses, we start at the negative extreme, and use a probability weighting function, w^- , for the domain of losses.

Remark 2.4 Abdellaoui (2000) and Abdellaoui et al. (2005) find that there is no significant difference in the curvature of the weighting function for gains and losses. For the Prelec function (see Definition 2.22), we know that the parameter α controls the curvature. However, the elevation (which in the Prelec function is controlled by the parameter β) can be different in the domain of gains and losses.

Definition 2.33 (Tversky and Kahneman, 1992). For PT, the decision weights, π_i , are defined as follows:

$$\begin{aligned} \pi_n &= w^+(p_n) \\ \pi_{n-1} &= w^+(p_{n-1} + p_n) - w^+(p_n) \dots \\ \pi_i &= w^+\left(\sum_{j=i}^n p_j\right) - w^+\left(\sum_{j=i+1}^n p_j\right) \dots \\ \pi_1 &= w^+\left(\sum_{j=1}^n p_j\right) - w^+\left(\sum_{j=2}^n p_j\right) \\ \pi_{-m} &= w^-(p_{-m}) \\ \pi_{-m+1} &= w^-(p_{-m} + p_{-m+1}) - w^-(p_{-m}) \dots \\ \pi_{-j} &= w^-\left(\sum_{i=-m}^{-j} p_i\right) - w^-\left(\sum_{i=-m}^{-j-1} p_i\right) \dots \\ \pi_{-1} &= w^-\left(\sum_{i=-m}^{-1} p_i\right) - w^-\left(\sum_{i=-m}^{-2} p_i\right) \end{aligned}$$

We now define the objective function of a decision maker who follows PT. This objective function is called the *value function*.

Definition 2.34 (*The value function*; Starmer and Sugden, 1989a; Tversky and Kahneman, 1992): The value of the lottery $L \in \mathcal{L}_P$ to the decision maker is given by

$$V(L) = \sum_{i=-m}^n \pi_i v(y_i). \quad (2.24)$$

Recall that under RDU, there is no notion of the domain of gains and losses, and the decision weights add up to one (see (2.9)). However, under PT this result may no longer hold. This is stated in the next proposition, which the reader is asked to prove.

Proposition 2.14 Under PT, if all outcomes are in the domain of gains or in the domain of losses, then the decision weights add up to one, otherwise they do not.

In most examples in this section, the Prelec (1998) probability weighting function will be used because it appears to be the one with the strongest empirical support. To quote from Stott (2006, p. 102): “the most predictive version of [cumulative prospect theory] has a power value curve, a single parameter risky weighting function due to Prelec (1998) and a Logit stochastic process.” Under PT, w^+ and w^- are allowed to differ, however, we will typically make the simplifying assumption that $w^+ = w^- = w$; see Prelec (1998).

2.4.1 A brief note on PT under uncertainty

PT can be easily extended to uncertainty along the lines described in Section 2.3.2, so we shall be brief. The main difference is that we need to work with lotteries in incremental form. Consider the following Savage act

$$L = (E_{-m}, y_{-m}; \dots; E_{-1}, y_{-1}; E_0, y_0; E_1, y_1; \dots; E_n, y_n), \quad (2.25)$$

where (2.23) holds, $y_0 = 0$ is the outcome relative to the reference point, and $E_{-m}, \dots, E_0, \dots, E_n$ are $m+n$ disjoint events whose union forms the sample space. Denote by $\bar{\mathcal{L}}_P$ the set of all lotteries of the form given in (2.25). Denote the event weighting function (Definition 2.29) for gains by $\tilde{w}^+$ and that for losses by $\tilde{w}^-$.

Definition 2.35 For PT under uncertainty, the decision weights, π_i , are defined as follows:

$$\begin{aligned} \pi_n &= \tilde{w}^+(E_n) \\ \pi_{n-1} &= \tilde{w}^+(E_{n-1} \cup E_n) - \tilde{w}^+(E_n) \dots \\ \pi_i &= \tilde{w}^+\left(\bigcup_{j=i}^n E_j\right) - \tilde{w}^+\left(\bigcup_{j=i+1}^n E_j\right) \dots \\ \pi_1 &= \tilde{w}^+\left(\bigcup_{j=1}^n E_j\right) - \tilde{w}^+\left(\bigcup_{j=2}^n E_j\right) \\ \pi_{-m} &= \tilde{w}^-(E_{-m}) \\ \pi_{-m+1} &= \tilde{w}^-(E_{-m} \cup E_{-m+1}) - \tilde{w}^-(E_{-m}) \dots \\ \pi_j &= \tilde{w}^-\left(\bigcup_{i=-m}^j E_i\right) - \tilde{w}^-\left(\bigcup_{i=-m}^{j-1} E_i\right) \dots \\ \pi_{-1} &= \tilde{w}^-\left(\bigcup_{i=-m}^{-1} E_i\right) - \tilde{w}^-\left(\bigcup_{i=-m}^{-2} E_i\right) \end{aligned}$$

As in the case under risk, since $v(y_0) = v(0) = 0$, the decision weight corresponding to the reference outcome, π_0 , can be chosen arbitrarily.

The objective function of the decision maker who uses PT under uncertainty is defined next.

Definition 2.36 (The value function): The value of the lottery $L \in \bar{\mathcal{L}}_P$ to the decision maker is given by

$$V(L) = \sum_{i=-m}^n \pi_i v(y_i). \quad (2.26)$$

2.4.2 Attitudes to risk under prospect theory

Under EU, $w(p) = p$, so the attitudes to risk are completely determined by the shape of the utility function. Under PT (and OPT), the utility function is concave in the domain of gains and convex in the domain of losses. From this, it might be tempting to conclude that a decision maker who uses PT is risk averse in the domain of gains and risk loving in the domain of losses. Such an inference would be incorrect. The reason is that in PT (and in most non-EU theories), attitudes

to risk are determined not only by the shape of the utility function but also by the shape of the probability weighting function, $w(p)$.

Example 2.12 Suppose that a decision maker who uses PT faces the lottery $L = (0, 1 - p; y, p)$ in incremental form, where $y > 0$. Using Definition 2.31, $v(0) = 0$. Let the utility function be given by $v(y) = \sqrt{y}$ and let the weighting function be the Prelec function, $w(p) = e^{-\beta(-\ln p)^\alpha}$, with $\alpha = 0.5$, $\beta = 1$. Using Definition 2.34,

$$V(L) = [1 - w(p)]\sqrt{0} + w(p)\sqrt{y} = w(p)\sqrt{y}.$$

The expected value of the lottery is $EL = py$. The certainty equivalent of the lottery, C_L , can be calculated as follows:

$$\begin{aligned} w(p)\sqrt{y} &= \sqrt{C_L}, \\ \Rightarrow C_L &= yw(p)^2. \end{aligned}$$

The decision maker is risk loving if $C_L > EL$ and risk averse if $C_L < EL$.

$$C_L \gtrless EL \Leftrightarrow yw(p)^2 \gtrless py \Leftrightarrow w(p)^2 \gtrless p. \quad (2.27)$$

For the special case of EU, $w(p) = p$. Since $p \in [0, 1]$, for EU, $w(p)^2 \leq p$ and so $C_L \leq E$. Thus, assuming a concave utility function under EU (in this case, $v(y) = \sqrt{y}$) is sufficient to guarantee risk aversion. However, for PT and for the Prelec function, $w(p)^2 > p$ for low p and $w(p)^2 < p$ for high p . In particular, for the parameter values chosen in this example,

$$w(p)^2 \gtrless p \Leftrightarrow -2(-\ln p)^{0.5} - \ln p \gtrless 0.$$

Thus, for $p^* = 0.0183$, $w(p^*)^2 = p^*$, for $p < 0.0183$, $w(p)^2 > p$ (i.e., the decision maker is risk loving despite a concave utility function), and for $p > 0.0183$, $w(p)^2 < p$ (i.e., the decision maker is risk averse).

The result in Example 2.12 is more general, and stated in the next remark.

Remark 2.5 Under PT, the decision maker has a strictly concave utility function in the domain of gains (see Definition 2.31). Nevertheless, for small probabilities in the domain of gains, the decision maker is risk loving, while for higher probabilities, the decision maker is risk averse.

We leave it to the reader to construct an analogous example in the domain of losses to illustrate the following more general remark.

Table 2.2 The fourfold classification of risk attitudes under prospect theory.

Probability	Domain of gains	Domain of losses
Low	$(100, 0.05; 0, 0.95)$	$(-100, 0.05; 0, 0.95)$
	$C = 14, E = 5$	$C = -8, E = -5$
High	$(100, 0.95; 0, 0.05)$	$(-100, 0.95; 0, 0.05)$
	$C = 78, E = 95$	$C = -84, E = -95$

Remark 2.6 *Despite having a convex utility function (see Definition 2.31) in the domain of losses, the decision maker is risk averse for low probabilities and risk loving for higher probabilities.*

The next example presents some experimental evidence supporting Remarks 2.5, 2.6. Put together, these remarks leads to a *fourfold classification of risk attitudes* under PT, which is one of the salient contributions of PT.

Example 2.13 *(The fourfold pattern of attitudes to risk; Tversky and Kahneman, 1992): Let C , E denote the certainty equivalent and the expected value, respectively of an associated lottery.*

Table 2.2 shows values of C , E for lotteries of the form $(x, p; 0, 1 - p)$, i.e., win $x > 0$ with probability p or 0 with probability $1 - p$ for low and high p in the domain of gains and losses. Comparing the values of C and E , we see that for low probabilities, there is risk seeking in the domain of gains but risk aversion in the domain of losses. For moderate or high probabilities, we see risk aversion in the domain of gains but risk seeking in the domain of losses. This pattern is hard to explain under EU or RDU, but is easily explained by PT. Thus, the attitudes to risk in PT are much richer than under EU or RDU.

In subsequent studies the fourfold pattern of risk has been tested using two kinds of methods. The first method is a *choice-based elicitation procedure* in which subjects are directly asked to choose between the lottery and its expected value. A risk averse decision maker should prefer the expected value to the lottery and a risk seeking decision maker should prefer to play the lottery. The second method is a *price based elicitation procedure* in which subjects are asked to state their maximum willingness to pay for the lottery if all outcomes are in gains and maximum willingness to pay to avoid playing the lottery if all outcomes are in the domain of losses. Harbaugh et al. (2002) and Harbaugh et al. (2010) use this method for a particularly simple class of lotteries of the form $L = (x, p)$. They find evidence for the fourfold classification of risk in the price based elicitation procedure but not in the choice-based task, where subjects appear to be risk neutral.

Fehr-Duda and Epper (2012) raise the following three objections about the findings in Harbaugh et al. (2010) that are more widely applicable.

1. They are skeptical about the findings of approximate risk neutrality in the choice-based tasks. They argue (p. 583) that these findings are not surprising: “Because it is so obvious that expected value equals outcome times probability, people may interpret the experiment as a test of their intelligence or numeracy rather than as a preference elicitation task.” For instance, an experimental subject, when asked for a choice between the lottery $(50, 0.10)$ and the expected value of the lottery, 5, may choose indifference between the two options.
2. A second issue they raise about the findings from lotteries of the form $L = (x, p)$ is that one cannot separately disentangle the parameters of the probability weighting function and the utility function, unless specific functional forms are used. They conclude (p. 583) with a cautionary note for experimental economists: “Therefore, we strongly recommend using multi-outcome prospects, with more than one nonzero outcome, over a wide range of probabilities to be able to generate a sufficiently rich database.”
3. Finally, the lotteries used in Harbaugh et al. (2010) deal with probabilities in the set $\{0.1, 0.4, 0.8\}$. We know from our discussion in Section 1.5 that violations of EU are more pronounced along the boundaries of the probability triangle. Indeed, in the interior of a probability triangle for, say, RDU, the indifference curves appear approximately linear. Figure 2.7 illustrates the situation for a three outcome lottery—the horizontal axis shows




Figure 2.7 Indifference curves under RDU over the full range of probabilities (left panel) and with the boundaries truncated (right panel).

Source: Reproduced with Permission of *Annual Review of Economics*, Volume 4, Helga Fehr-duda and Thomas Epper, “Probability and risk: foundations and economics implications of Probability-appendent risk preferences,” © 2012 by Annual Reviews <http://www.annualreviews.org>

the probability, p_3 , of the worst outcome and the vertical axis, the probability, p_1 , of the best outcome.²³ The indifference curves in the probability triangle are for RDU. In the left panel the indifference curves are drawn over the entire relevant space. The right panel shows the same indifference curves but the boundaries have been truncated; the indifference curves in the interior are now approximately linear.

Rieger et al. (2014) compare risk attitudes in 53 countries. They calculate the relative risk premium for gains and losses. There is strong support for risk aversion in gains and risk seeking in losses. The average relative risk premium in gains is 0.70 (risk aversion) and, in losses it is -0.46 (risk seeking). Individuals in richer countries are more risk averse in gains but also more risk seeking in losses. There are also gender effects—women are more risk averse in gains but they are also more risk seeking in losses.

We also consider the evidence on the fourfold classification of risk elsewhere in Part I; see, for instance, Section 2.9.3 and Figure 2.13 for confirmatory evidence.

2.4.3 A violation of EU and RDU that can be explained by PT

Recall Example 2.11 for which EU and RDU were both refuted. We now show that this example is consistent with PT because PT takes account of changes in the reference point, which EU and RDU cannot.

Example 2.14 (Support for PT against EU and RDU). Suppose that the decision maker has the utility function $v(y) = (y)^{0.6}$ and uses the Prelec probability weighting function $w(p) = e^{-(-\ln p)^{0.65}}$. Consider Example 2.11 that constituted a refutation of EU and RDU. Problem 11 in Example 2.11 involved the following choice

$$A = (x_0 + 1000, 0.5; x_0 + 2000, 0.5) \text{ versus } B = (x_0 + 1000, 0; x_0 + 1500, 1).$$

²³ Unlike our notational convention in Part I, here the three outcomes are ordered as $x_3 < x_2 < x_1$.

Let us take a status-quo interpretation of reference points. Thus, for Problem 11, the reference point is the endowment point, $x_0 + 1000$ and so $y = x - (x_0 + 1000)$. Thus, the value function for each of these lotteries is given, respectively, by

$$V(A) = 1000^{0.6} e^{-(\ln 0.5)^{0.65}} = 28.692.$$

$$V(B) = 500^{0.6} = 41.628.$$

PT predicts that lottery B is preferred over lottery A, which is true for most subjects. Problem 12 in Example 2.11 considered the following choice,

$$C = (x_0 + 1000, 0.5; x_0 + 2000, 0.5) \text{ versus } D = (x_0 + 1500, 1; x_0 + 2000, 0).$$

In this case, the reference point is $x_0 + 2000$ and so $y = x - (x_0 + 2000)$. Thus, the value function corresponding to each of these lotteries is

$$V(C) = -2.25 (1000)^{0.6} e^{-(\ln 0.5)^{0.65}} = -64.558.$$

$$V(D) = -2.25 (500)^{0.6} = -93.662.$$

PT predicts that C should be chosen over D, which is true for a majority of the subjects.

2.4.4 Some erroneous criticisms of PT

We now consider some of the common criticisms of PT that are either misplaced or overstated.

1. *PT is only an “as if” theory:* All the mainstream decision theories that we have looked at, EU/RDU/PT, are ‘as if’ theories. Each of these theories involves a relatively high degree of cognitive sophistication. There is no supposition that the mental processes of decision makers mirror closely any of these theories, and the theories make no assumptions about cognitive processes either. However, evidence from neuroeconomics confirms various aspects of PT, such as reference dependence, loss aversion, and non-linear probability weighting (see Part 9 for details).

So why does PT explain the evidence from many diverse problems so well, even if it is an “as if” theory? The answer probably lies in the individual components of PT that have strong empirical support. These components include reference dependence, loss aversion, diminished sensitivity to gains and losses, shape of the utility functions in the domain of gains and losses, and inverse S-shaped probability weighting functions. Thus, in many cases, the “as if” assumption actually turns out to be justified when one uses PT. By contrast, the empirically observed violations of EU, considered above, show that the “as if” supposition does not often hold in the case of EU.

2. *PT performs better only because it has additional parameters:* Another commonly heard criticism of PT is that “these theories have more parameters than EU and RDU, so it is ‘no surprise’ that they have greater explanatory power.”

Proponents of EU who make this criticism will find that EU is subjected to a similar criticism. EU has more parameters than expected value maximization (EV), so if the “no surprise” argument has any intrinsic merit, it would also apply to the better explanatory power of EU relative to EV (e.g., the St. Petersburg paradox). It is in the nature of scientific progress that one discovers new anomalies that existing theory cannot explain. One then needs to construct a new theory that explains everything that the old theory can explain, and also has the ability to explain the new anomalies. It is often unavoidable that the new theory has

a larger number of parameters. The intellectual victory of EU over EV when the St. Petersburg paradox arose is a good example. A trivial example can also be used to make the same point. Suppose that one observes a scatter of points that is unmistakably quadratic. Then fitting a straight line rather than a curve through the points may be parsimonious, yet hardly satisfactory.

The “no surprise” argument would have intellectual merit if the new theory explains less than the old theory particularly out-of-sample, is intractable, and lacks parsimony. This is not the case with PT. PT is elegant, tractable, based on robust psychological foundations, and relatively parsimonious. It explains everything that EU and RDU can explain; the converse is false.

A theory with more parameters may explain better the in-sample data but not necessarily the out-of-sample data. Indeed, one can use statistical tests that allow us to trade off greater explanatory power against the larger number of parameters (e.g., the Akaike information criterion and the Bayesian information criterion). This methodology is routinely used to test alternative models of learning (see Part 5 of the book). We have already discussed in the introductory chapter that PT does not predict better than the alternative theories, just because it has a larger number of parameters.

One can separately elicit the parameters of the utility and the weighting functions, parametrically or non-parametrically, under PT; Section 2.5 is devoted to this endeavor. The results strongly support the functional forms under PT. Mixture models that allow for the presence of both EU and PT decision makers find that the vast majority of subjects (80% in the study of Bruhin et al. (2010) reviewed in Section 2.9.4) follow PT. On these grounds, EU is simply a mis-specified theory for most decision makers and the extra parameters under PT are supported by the empirical evidence.

3. *PT has an arbitrary reference point, which makes it vacuous:* Not only is it difficult to specify a universal reference point, it might not be desirable to do so.
 - (a) PT is not unique among mainstream theories in not being able to strictly specify all components. For instance, Bayesian statistics does not specify a prior distribution but teaches us a method of updating a prior to a posterior distribution. This inability to specify a prior is not taken as a fatal shortcoming. Another example is that the predictions of neo-classical theory are often very sparse unless one is willing to specify explicitly, preferences and technology. A good example is general equilibrium theory; see, for instance, section 17E in Mas-Collel et al. (1995) titled “Anything goes . . .”. Indeed, the real bite in the predictions from many models comes from actually specifying preferences and technology. This is hardly considered to be fatal for neoclassical economics.
 - (b) What if human behavior is actually “sufficiently rich” in terms of the reference point? “Sufficiently rich” does not imply infinitely rich or random. Indeed, many candidates for a reference point seem very salient across many problems and contexts. Of these, the *status quo* or a fair/just entitlement (even if the fairness is subjective) is one of the most salient. In actual applications, this is often able to guide the researcher to a suitable reference point. In many contexts, a natural reference point arises on closer examination. For instance, in their explanation of the tax evasion puzzles (see Section 3.5 below), Dharmi and al-Nowaihi (2007, 2010b) assume the “legal after-tax liabilities” to be the reference point for a potential tax evader; this has the flavor of a status-quo point. Their calibrated results, based on this suggestion match closely the data from the problem.
 - (c) There has been progress in endogenous specifications of the reference point; see, for instance, the discussion in Section 2.8 on the Köszegi–Rabin framework. In this framework,

the reference point is determined within the model, in a manner that is consistent with rational expectations. The advantage is that this framework is not subject to the potential criticism of an arbitrary reference point. The disadvantage is that the required cognitive sophistication of the decision makers is very high.

A more serious criticism of PT comes from the evidence on human behavior for events of very low and very high probabilities. This is dealt with in Section 2.11.3 below. Indeed, newer theories that are able to handle events of very low and very high probabilities, such as *composite prospect theory* (see Section 2.11.3), rely heavily on the insights of PT because PT captures fundamentally important insights about human behavior. Other serious criticisms are discussed in Section 2.9. These have to do with the evidence from configural weights models, violations of gain–loss separability, the absence of “regret” in PT, and the interaction between probability weights and stake size.

2.5 Elicitation of utility and probability weighting functions in PT

In this section, we briefly consider some of the methods of elicitation under PT as well as some of the parameter estimates in the literature; for a more detailed discussion, see Wakker (2010).

2.5.1 Elicitation of the utility function under PT

In order to elicit the utility function under PT, let us suppose that the reference point is zero and proceed in the four steps suggested in Abdellaoui et al. (2007).

I. Determining probabilities p_l, p_g : $w^-(p_l) = \frac{1}{2}$ and $w^+(p_g) = \frac{1}{2}$, where w^- and w^+ are, respectively, probability weighting functions in the domain of losses and gains.

Suppose that, relative to the reference point, we are given (i) three outcomes in the loss domain, $l_0 < l < l^*$, and (ii) some level of probability p . Now we elicit two levels of outcomes in the loss domain, l_1 and l_2 , such that $l_2 < l_1 < l_0 < l < l^*$. We first find the utility midpoints using the method of Abdellaoui (2000) and Köbberling and Wakker (2003).

Suppose that the following indifferences hold

$$(l_1, p; l^*, 1 - p) \sim (l_0, p; l, 1 - p). \quad (2.28)$$

$$(l_2, p; l^*, 1 - p) \sim (l_1, p; l, 1 - p). \quad (2.29)$$

Under PT, for $i = 0, 1$, $(l_{i+1}, p; l^*, 1 - p) \sim (l_i, p; l, 1 - p)$ implies that

$$w^-(p)v(l_{i+1}) + (1 - w^-(p))v(l^*) = w^-(p)v(l_i) + (1 - w^-(p))v(l), \quad (2.30)$$

where v is the utility function under PT. Rewriting (2.30),

$$v(l_i) - v(l_{i+1}) = \frac{1 - w^-(p)}{w^-(p)} (v(l^*) - v(l)). \quad (2.31)$$

The right hand side is independent of the indices, $i, i + 1$, thus, for $i = 0, 1$, and using (2.28), (2.29), we have

$$v(l_0) - v(l_1) = v(l_1) - v(l_2). \quad (2.32)$$

Thus,

$$v(l_1) = \frac{1}{2} [v(l_0) + v(l_2)], \quad (2.33)$$

is the utility midpoint of $v(l_0)$ and $v(l_2)$. Once the values l_1 and l_2 are elicited, then we find the probability p_l such that $(l_1, 1) \sim (l_2, p_l; l_0, 1 - p_l)$. This indifference implies that

$$v(l_1) = w^-(p_l)v(l_2) + (1 - w^-(p_l))v(l_0). \quad (2.34)$$

Substituting (2.33) into (2.34) we get the required value

$$w^-(p_l) = \frac{1}{2}. \quad (2.35)$$

We now proceed to the domain of gains. We begin with a probability p and three outcomes in the domain of gains $g_0 > g > g^*$. Now we elicit two levels of gains g_1 and g_2 such that $g_2 > g_1 > g_0 > g > g^*$. Suppose that the following indifferences hold

$$(g^*, 1 - p; g_1, p) \sim (g, 1 - p; g_0, p). \quad (2.36)$$

$$(g^*, 1 - p; g_2, p) \sim (g, 1 - p; g_1, p). \quad (2.37)$$

Having elicited g_1 and g_2 we can now proceed analogously to the domain of losses to find a probability p_g such that

$$w^+(p_g) = \frac{1}{2}. \quad (2.38)$$

II. Eliciting utility in the loss domain.

Pick an outcome y_l in the domain of losses and normalize $v(y_l) = -1$. We now elicit the outcome $l_{0.5}$ in the domain of losses such that $(l_{0.5}, 1) \sim (y_l, p_l; 0, 1 - p_l)$, where p_l is found in (2.35). Under PT, this indifference implies that

$$v(l_{0.5}) = w^-(p_l)v(y_l).$$

Substituting (2.35) and $v(y_l) = -1$ we get

$$v(l_{0.5}) = -\frac{1}{2}.$$

We can now proceed to elicit the utility function in the interval $[y_l, 0]$. For instance, elicit an outcome $l_{0.25}$ such that $(l_{0.25}, 1) \sim (l_{0.5}, p_l; 0, 1 - p_l)$. This indifference implies that

$$v(l_{0.25}) = w^-(p_l)v(l_{0.5}) = -\frac{1}{4}.$$

Similarly we can find $l_{0.125}$ such that $(l_{0.125}, 1) \sim (l_{0.25}, p_l; 0, 1 - p_l)$, so $v(l_{0.125}) = w^-(p_l)v(l_{0.25}) = -\frac{1}{8}$, and so on.

III. Linking utility in the loss domain with utility in the gain domain.

We now relate utility in the loss domain from Step II with utility in the gain domain. This is an important step in the framework of Abdellaoui et al. (2007), which allows for the measurement of loss aversion.²⁴ We consider three indifferences.

(a) Consider one of the elicited outcomes in the interval $[y_l, 0]$ in Step II, say, l_s with the corresponding utility function $v(l_s) = -s$ (e.g., $v(l_{0.25}) = -0.25$). We now find an outcome l in the loss domain such that $(l_s, 1) \sim (l, 0.5; 0, 0.5)$. This implies that $v(l_s) = w^-(0.5)v(l)$, so

$$-s = w^-(0.5)v(l). \quad (2.39)$$

(b) Now elicit an outcome in the domain of gains, g , such that $(0, 1) \sim (l, 0.5; g, 0.5)$. This implies that $0 = w^-(0.5)v(l) + w^+(0.5)v(g)$. Using (2.39) we get

$$w^+(0.5)v(g) = s. \quad (2.40)$$

(c) Elicit the outcome g_s in the domain of gains so that $(g_s, 1) \sim (g, 0.5; 0, 0.5)$. This implies that $v(g_s) = w^+(0.5)v(g)$. Using (2.40), we get that $v(g_s) = s$. Thus, in terms of utility, the outcome g_s in the domain of gains is the mirror image of the outcome l_s in the domain of loss.

IV. Eliciting utility in the gain domain.

We now determine utility in the gain domain in the interval $[0, g_s]$, where g_s was determined in Step III. We elicit $g_{0.5} \in [0, g_s]$ such that $(g_{0.5}, 1) \sim (0, 1 - p_g, g_s, p_g)$. This gives rise to the condition $v(g_{0.5}) = w^+(p_g)v(g_s)$. Using (2.38) and $v(g_s) = s$ we get $v(g_{0.5}) = \frac{s}{2}$. Proceeding in an analogous manner, we can generate utility in the gains domain for other values in the interval $[0, g_s]$ such as $g_{0.25}, g_{0.125}$.

Figure 2.8 shows a non-parametric utility function fitted to the data from Abdellaoui et al. (2007) for EU and for PT. The non-parametric curve is concave in gains and convex in losses.²⁵ The degree of concavity and convexity is more pronounced relative to several other studies.²⁶ At the individual level, the most common pattern was an S-shaped utility function. Only one individual out of 48 behaved in a manner consistent with the typical assumption in EU of a concave utility function throughout. Similar results are obtained by Abdellaoui et al. (2013) when they elicit the preferences of private bankers and fund managers. This serves to contribute towards the external validity of lab experiments in such domains.

Figure 2.9 shows a fit of the power form of utility in (2.20) to the same data. The coefficient γ in the domain of gains is 0.75 and in the domain of loss it is 0.74. The curve has an extremely good fit and most points lie very close to the fitted curve.

One disadvantage of non-parametric estimation is that there might be substantial error propagation. Further, the procedure is not very efficient in the sense that too many iterations are needed, imposing greater cognitive requirements on the subjects. In response to these concerns, Abdellaoui et al. (2008) propose a parametric estimation method with the power form of utility. They

²⁴ In the trade-off method of Wakker and Deneffe (1996), one cannot measure simultaneously the utility in the domain of gains and losses, thus making it difficult to measure loss aversion.

²⁵ Notice that if we restricted attention only to EU, then the finding of risk seeking behavior in the domain of losses does not allow us to conclude that the utility function is convex. This is because risk seeking in the domain of losses might also arise due to the shape of the probability weighting function; see Abdellaoui et al. (2007).

²⁶ For instance, Tversky and Kahneman (1992), Fennema and van Assen (1998), Abdellaoui (2000), and Schunk and Betsch (2006).




Figure 2.8 Non-parametric estimates of utility for gains and losses under PT and EU based on median data.
Source: Abdellaoui et al. (2007). Reproduced with permission. Copyright, INFORMS, <http://www.informs.org>

confirm all the individual features of prospect theory: Concavity of utility in the domain of gains, convexity in the domain of losses, loss aversion, and inverse S-shaped probability weighting.

2.5.2 Elicitation of loss aversion under PT

There are many different definitions of loss aversion.²⁷ However, these definitions typically rely on the utility function and not on the probability weighting function. Hence, the elicitation of the utility function in Section 2.5.1 sets the stage for the elicitation of loss aversion.

The original definition of loss aversion, due to Kahneman and Tversky (1979), is given in Definition 2.31, i.e.,

$$\frac{-v(-y)}{v(y)} > 1; y > 0.$$

When the utility function takes the power form in (2.20), loss aversion is simply

$$\lambda = \frac{-v(-1)}{v(1)}.$$

²⁷ For some work on the behavioral foundations of loss aversion, see Peters (2012).




Figure 2.9 Parametric estimates of the power form of utility for gains and losses under PT based on median data.

Source: Abdellaoui et al. (2007). Reproduced with permission. Copyright, INFORMS, <http://www.informs.org>

Köbberling and Wakker (2005) consider the implications of an alternative definition of loss aversion. In this definition, suggested originally but informally by Benartzi and Thaler (1995), for a general utility function v ,

$$\lambda = \frac{v'_{\uparrow}(0)}{v'_{\downarrow}(0)}, \quad (2.41)$$

where $v'_{\uparrow}(0)$ and $v'_{\downarrow}(0)$ are, respectively, the left hand side and the right hand side limits around 0.

The Wakker and Tversky (1993) definition of loss aversion requires a greater slope of the utility function at any loss outcome relative to the slope at the absolute value of the loss, i.e.,

$$v'(-y) > v'(y), \forall y > 0.$$

Neilson (2002) proposed the following requirement for a loss averse decision maker

$$\frac{v(-x)}{-x} \geq \frac{v(y)}{y}, x > 0, y > 0.$$

Once the utility function has been estimated using the technique in Section 2.5.1, then any of these definitions can be applied to estimate loss aversion.

There are also some definitions of loss aversion that rely on the probability weighting function.

Definition 2.37 (Kahneman and Tversky, 1979): Let $\prec$ be a binary preference relation over lotteries. An individual is loss averse if $(y, 0.5; -y, 0.5) \prec (z, 0.5; -z, 0.5)$ where $y > z \geq 0$.

In OPT, Kahneman and Tversky (1979) defined loss aversion as the tendency of people to reject symmetric 50–50 bets (set $z = 0$ in Definition 2.37). Schmidt and Zank (2005) extend Definition 2.37 to the concept of *strong loss aversion*. As one might expect from Definition 2.37, in a comparison of the lotteries $(y, 0.5; -y, 0.5)$ and $(z, 0.5; -z, 0.5)$ under prospect theory, the weighting function plays a critical role. Thus, loss aversion, in this definition, involves a condition between changes in utility and the magnitude of probability weights.

Abdellaoui et al. (2007) summarize some of the estimates of loss aversion, based on aggregate data from previous studies; see, Table 2.3. The estimates from separate studies are sometimes not comparable because some report median values, while others report mean values.²⁸ Based on this small set of results, the Tversky and Kahneman (1992) estimate of 2.25 is quite representative. The estimates in Table 2.3 are for aggregate data. Abdellaoui et al. (2007) also report their own estimates from individual data, based on various definitions of loss aversion.²⁹ Using the Kahneman and Tversky (1979) definition, they detect the presence of loss aversion at the aggregate and the individual level; 39 out of 47 subjects are loss averse. Using the Wakker and Tversky (1993) definition, 30 out of 36 are loss averse. In the Köbberling and Wakker (2005) definition, 35 out of 47 subjects are loss averse. The mean estimates of loss aversion for these three studies are, respectively, 2.04, 1.71, and 8.27.

A large number of estimates of loss aversion have come from studies that examine exchange disparities. We consider these estimates separately in Section 3.2.1. Section 3.4 considers the concept of *myopic loss aversion*.

Table 2.3 Estimates of loss aversion in various studies. The two estimates in Bleichrodt et al. (2001) are based on two different datasets. The two estimates in Booij and van de Kuilen (2009) are based on high and low monetary amounts.

Study	Definition of loss aversion	Estimate
Fishburn and Kochenberger (1979)	$\frac{v'(-y)}{v'(y)}$	4.8
Tversky and Kahneman (1992)	$\frac{-v(-1)}{v(1)}$	2.25
Bleichrodt et al. (2001)	$\frac{v(-y)}{v(y)}$	2.17, 3.06
Schmidt and Traub (2002)	$\frac{v'(-y)}{v'(y)}$	1.43
Pennings and Smidts (2003)	$\frac{v'(-y)}{v'(y)}$	1.81
Booij and van de Kuilen (2009)	$\frac{v'_\uparrow(-y)}{v'_\downarrow(y)}$	1.79, 1.74

Source: Abdellaoui et al. (2007).

²⁸ In Table 2.3, median values are reported by Fishburn and Kochenberger (1979), Tversky and Kahneman (1992), and Bleichrodt et al. (2001).

²⁹ Estimates of loss aversion at the individual level are also provided in Schmidt and Traub (2002) and Bleichrodt and Pinto (2002) although these studies find a lower proportion of loss averse subjects.

2.5.3 Elicitation of the probability weighting function under PT

One can either estimate the probability weighting function parametrically or non-parametrically. Table 2.4 shows some parameter estimates reported in Bleichrodt and Pinto (2000) for three main probability weighting functions.

For several reasons, one might be interested in non-parametric estimation of the probability weighting function. Consider the trade-off method proposed by Wakker and Deneffe (1996). Suppose that we begin with three outcomes in the domain of gains, $R > r > x_0$ such that $(R, 1) \succeq (x_0, 1)$. The outcome x_1 is then elicited such that

$$(x_0, 1 - p; R, p) \sim (x_1, 1 - p; r, p). \quad (2.42)$$

Under PT, the indifference in (2.42) implies that

$$w(p)v(R) + (1 - w(p))v(x_0) = w(p)v(r) + (1 - w(p))v(x_1). \quad (2.43)$$

Now the outcome x_2 is elicited such that

$$(x_1, 1 - p; R, p) \sim (x_2, 1 - p; r, p)$$

which implies that

$$w(p)v(R) + (1 - w(p))v(x_1) = w(p)v(r) + (1 - w(p))v(x_2). \quad (2.44)$$

From (2.43) and (2.44)

$$(1 - w(p))(v(x_2) - v(x_1)) = w(p)(v(R) - v(r)). \quad (2.45)$$

Table 2.4 Estimates of the parameters of various probability weighting functions.

Probability weighting function	Parameter estimates
Tversky–Kahneman (1992)	Tversky–Kahneman (1992): $\tau = 0.61$ (gains), $\tau = 0.69$ (losses)
$w(p) = \frac{p^\tau}{[p^\tau + (1-p)^\tau]^{\frac{1}{\tau}}}$	Camerer and Ho (1994): $\tau = 0.56$ (gains)
	Wu and Gonzalez (1996): $\tau = 0.71$ (gains)
	Abdellaoui (2000): $\tau = 0.60$ (gains), $\tau = 0.70$ (losses)
	Wu and Gonzalez (1996): $\delta = 0.84$, $\eta = 0.68$ (gains)
Goldstein–Einhorn (1987)	Gonzalez and Wu (1999): $\delta = 0.77$, $\eta = 0.44$ (gains)
$w(p) = \frac{\delta p^\eta}{\delta p^\eta + (1-p)^\eta}$	Tversky and Fox (1995): $\delta = 0.77$, $\eta = 0.69$ (gains)
	Abdellaoui (2000): $\delta = 0.65$, $\eta = 0.60$ (gains)
	Abdellaoui (2000): $\delta = 0.84$, $\eta = 0.65$ (losses)
Prelec (1998)	
$w(p) = e^{-(-\ln p)^\alpha}$	Wu and Gonzalez (1996): $\alpha = 0.74$ (gains)

Source: Bleichrodt and Pinto (2000).

Notice that the RHS is independent of x_1 and x_2 . One can now continue these iterations of indifferences. For instance, the next stage would be to find an x_3 such that $(x_2, 1 - p; R, p) \sim (x_3, 1 - p; r, p)$. In general, for any x_i and x_{i-1} , the analogue of (2.45) is

$$(1 - w(p)) (v(x_i) - v(x_{i-1})) = w(p) (v(R) - v(r)). \quad (2.46)$$

Since the RHS is independent of x_i and x_{i-1} , iterating the process k times, we get

$$v(x_1) - v(x_0) = v(x_2) - v(x_1) = v(x_3) - v(x_2) = \dots = v(x_k) - v(x_{k-1}). \quad (2.47)$$

Normalize $v(x_0) = 0$ and $v(x_k) = 1$. From the first equality in (2.47), we get $v(x_1) = \frac{1}{2}v(x_2)$. From the second equality, we get $v(x_2) = \frac{2}{3}v(x_3)$. In general, for the j th equality, we get $v(x_j) = \frac{j}{j+1}v(x_{j+1})$, and from the last equality, we get $v(x_{k-1}) = \frac{k-1}{k}v(x_k) = \frac{k-1}{k}$. Rolling back the iterations, we get $v(x_{k-2}) = \frac{k-2}{k-1}v(x_{k-1}) = \frac{k-2}{k-1} \frac{k-1}{k} = \frac{k-2}{k}$. This establishes the pattern of recursion, so that we have

$$v(x_j) = \frac{j}{k}, j = 1, 2, \dots, k. \quad (2.48)$$

Once we obtain the utility function in this manner (or in the manner suggested in Section 2.5.1), we can follow Abdellaoui (2000). Choose any $x_j \in \{x_0, x_1, \dots, x_k\}$ (these are the first stage elicited values), and then elicit the probability q such that

$$(x_0, 1 - q; x_k, q) \sim (x_j, 1),$$

which implies that

$$w(q)v(x_k) + (1 - w(q))v(x_0) = v(x_j)$$

Using $v(x_k) = 1$, $v(x_0) = 0$, and $v(x_j) = j/k$, we get

$$w(q) = \frac{j}{k}.$$

Now we select another element from $\{x_0, x_1, \dots, x_k\}$ and repeat the procedure to get the probability weight associated with another probability value and so on for each point in $\{x_0, x_1, \dots, x_k\}$. Bleichrodt and Pinto (2000) apply this procedure to health outcomes (rather than monetary prizes). They find strong evidence of an inverse S-shaped probability weighting function at the individual and the aggregate levels. As to the shapes of the weighting function at the individual level, they find that 10.9% of the subjects had a concave weighting function, 0% had a linear weighting function, 4.4% had a convex weighting function, and 83.7% had an inverse S-shaped weighting function.

2.6 The axiomatic foundations of PT

We briefly consider the preference foundations of various elements of PT in this section. The interested reader can consult Wakker (2010) for a comprehensive treatment. Since the Prelec (1998) probability weighting function (see Definition 2.22) and the power form of utility (see Equation 2.20) are used extensively in Part 1, the reader is specifically given a flavor of their preference foundations.

2.6.1 Foundations of PT in terms of the trade-off consistency method

In this section we consider the preference foundations of PT based on the *trade-off consistency method* (Chateauneuf and Wakker, 1999; Köbberling and Wakker, 2003).³⁰ This method gives the preference foundations for EU, RDU, and PT under *risk* as well as under *uncertainty*. Our presentation here is based on Chateauneuf and Wakker (1999).³¹

For pedagogical simplicity, let $X = \mathbb{R}$ be the set of monetary outcomes.³² The reference point is normalized to 0. A typical prospect is given by $L = (y_1, p_1; y_2, p_2; \dots; y_n, p_n)$ (see Definition 2.30, with $x_0 = 0$), where $y_1 < y_2 < \dots < y_n$ are outcomes in X and $p_1, p_2, \dots, p_n$ are probabilities such that $p_i \geq 0$ and $\sum_i p_i = 1$. Let $\mathcal{L}_P$ be the set of all such prospects.

Let $\succeq$ be a preference relation over $\mathcal{L}_P$ and let $\succ$ and $\sim$ be the strict preference and indifference relations. Sure outcomes y_i are equivalent to the prospect $(y_i, 1)$. An outcome $y_i \succ 0$ is a gain, while $y_i < 0$ is a loss. Since we have assumed that X is the set of monetary outcomes, so for sure outcomes, $\succeq$ corresponds in a natural way with $\geq$.

If $\succeq$ is a *weak order* (i.e., complete and transitive), then we can find a preference representation such that for any two prospects $L_1, L_2 \in \mathcal{L}_P$ we have $L_1 \succeq L_2 \Leftrightarrow V(L_1) \geq V(L_2)$. Let $x = (x_1, x_2, \dots, x_n) \in X^n$ and fix the prospect $L_1 = (y_1, p_1; \dots; y_n, p_n)$. Preferences $\succeq$ are *continuous* if for all $x \in X^n$ and fixed probabilities $p_1, p_2, \dots, p_n$, the following two sets are closed

$$\{x : (x_1, p_1; \dots; x_n, p_n) \succeq L_1\}, \{x : (x_1, p_1; \dots; x_n, p_n) \leq L_1\}.$$

Let $L_1 = (y_1, p_1; \dots; y_n, p_n)$ and $L_2 = (x_1, p_1; \dots; x_n, p_n)$. Then, $\succeq$ satisfies *stochastic dominance* if $x_i \geq y_i$ for $i = 1, 2, \dots, n$ implies $L_2 \succeq L_1$ and if, in addition, $x_i \succ y_i$ for at least one i we have $L_2 \succ L_1$ (see Definition 1.9 and note that $x_i \geq y_i \Leftrightarrow x_i \succeq y_i$).

Suppose that we are given an arbitrary prospect $L = (y_1, p_1; \dots; y_n, p_n)$. Then, we define a new prospect $\alpha_j L$ which replaces the j th outcome in L , y_j , by α_j , i.e.,

³⁰ On the penultimate line of the second paragraph on p. 94 of his book, Wakker (2010) credits Pfanzagl (1968, Remark 9.4.5) with the recommendation of the trade-off tool. The interested reader can consult Wakker (2010) for a detailed statement of the relevant theorems and the references. For the preference foundations of PT in the continuous case, see Kothiyal et al. (2011).

³¹ For other axiomatizations of prospect theory, see Tversky and Kahneman (1992), Wakker and Tversky (1993), Chateauneuf and Wakker (1999), Zank (2001), Wakker and Zank (2002), Köbberling and Wakker (2003), Schmidt (2003), Schmidt and Zank (2007, 2008). All these axiomatizations assumed an exogenous reference point. For axiomatizations of the rank dependent model, see Quiggin (1981, 1982), Yaari (1987), Schmeidler (1989), Luce (1991), Abdellaoui (2002), Luce and Marley (2005), and Luce and Fishburn (1991). For axiomatic extensions of the RDU model see Green and Jullien (1988), Chew and Epstein (1989a), and Segal (1989, 1993). These sets of papers provide links between EU and RDU as well as showing the connection between RDU and the betweenness axiom.

³² The framework can be extended to more general spaces that do not restrict attention to monetary outcomes and higher dimensional spaces such as $\mathbb{R}^m$ where elements of X are m dimensional bundles of goods, etc.

$$\alpha_j L = (y_1, p_1; \dots; y_{j-1}, p_{j-1}; \alpha_j, p_j; y_{j+1}, p_{j+1}; \dots; y_n, p_n); y_{j-1} < \alpha_j < y_{j+1}.$$

Definition 2.38 Let there be four outcomes (all in gains or all on losses) a, b, c, d . Consider the preference relation $\succ^*$ such that $ab \succ^* cd$ (also written as $[a; b] \succ^* [c; d]$) if there exist prospects $L_1 = (y_1, p_1; y_2, p_2; \dots; y_n, p_n)$, $L_2 = (x_1, p_1; x_2, p_2; \dots; x_n, p_n)$ (with identical probabilities) and an index $j = 1, \dots, n$, such that

$$a_j L_1 \succeq b_j L_2 \text{ and } c_j L_1 \prec d_j L_2.$$

Example 2.15 Suppose that $L_1 = (80, 0.3; y_2, 0.7)$ and $L_2 = (90, 0.3; x_2, 0.7)$. Then, using Definition 2.38, $[1; 0] \succ^* [9; 8]$ if

$$(80, 0.3; 1, 0.7) \succeq (90, 0.3; 0, 0.7) \text{ and } (80, 0.3; 9, 0.7) \prec (90, 0.3; 8, 0.7).$$

Recall that v is the utility function under PT (Definition 2.31). The next lemma illustrates the usefulness of the preference relation $\succ^*$ for utility differences.

Lemma 2.1 Suppose that the decision maker follows PT. Then

$$[a; b] \succ^* [c; d] \Rightarrow v(a) - v(b) > v(c) - v(d).$$

Proof: Given Definition 2.38, the decision weights in each pair of prospects $(a_j L_1, b_j L_2)$ and $(c_j L_1, d_j L_2)$ only differ at the j th place; denote this decision weight by π_j . Then $a_j L_1 \succeq b_j L_2 \Leftrightarrow V(a_j L_1) \geq V(b_j L_2)$ or (using Definition 2.34)

$$\pi_j(v(a) - v(b)) \geq \sum_{i \neq j} \pi_i v(x_i) - \sum_{i \neq j} \pi_i v(y_i). \quad (2.49)$$

Analogously $c_j L_1 \prec d_j L_2$ implies that

$$\pi_j(v(c) - v(d)) < \sum_{i \neq j} \pi_i v(x_i) - \sum_{i \neq j} \pi_i v(y_i). \quad (2.50)$$

Using (2.49), (2.50) we get

$$v(a) - v(b) > v(c) - v(d).$$

The second inequality in the statement of the proposition can be derived analogously, so we omit it. ■

In the elicitation of the utility function under PT, see (2.48), notice that the elicited utility differences are constant. This is an implication of Lemma 2.1. To see this, the procedure elicits values $x_0, x_1, x_2, \dots$ such that $[x_{j+1}; x_j] \sim^* [x_1; x_0]$ so $[x_{j+1}; x_j] \succeq^* [x_1; x_0]$ and $[x_{j+1}; x_j] \preceq^* [x_1; x_0]$, which implies, from Lemma 2.1, that $v(x_{j+1}) - v(x_j) > v(x_1) - v(x_0)$, as claimed.

In order to ensure that this elicitation process does not run into contradictions, we impose the next condition.

Definition 2.39 (*Trade-off consistency*³³): Under trade-off consistency there cannot exist outcomes a, b, c, d such that $[a; b] \succ^* [c; d]$ and $[c; d] \succeq^* [a; b]$.

The main result is given in the next proposition.

Proposition 2.15 (*Chateauneuf and Wakker, 1999*): Assuming trade-off consistency, the following two statements are equivalent.

- (i) Cumulative prospect theory holds with a continuous utility function.
- (ii) The preference relation $\succeq$ satisfies the following conditions.
 - (a) Weak ordering.
 - (b) Continuity.
 - (c) Stochastic dominance.
 - (d) Trade-off consistency.

For an extension of this framework under risk to issues of uncertainty, a richer set of results, and a discussion of the related literature, see Köbberling and Wakker (2003), and Wakker (2010). The existing axiomatizations listed above, presuppose a reference point. Schmidt and Zank (2012) propose an alternative set of axioms to get a PT representation that does not impose a reference point. They rely on the idea of diminishing sensitivity around a reference point to endogenously identify a reference point.

2.6.2 Foundations of the Prelec weighting function

Under EU, one can use the reduction axiom to reduce compound lotteries to simple lotteries by using the product rule of probabilities. Under non-linear probability weighting, this method no longer works and we need something to replace the product rule of probabilities. Alternative preference foundations of the Prelec function propose different forms of the product rule.

Let X be the set of *outcomes*. Since outcomes are monetary losses or gains in all our applications, we may take X to be the set of real numbers, R .³⁴ We assume that the preferences of the decision maker are described by a real-valued function $u : X \rightarrow R$; $y \in X$ being preferred to $x \in X$ if, and only if, $u(y) \geq u(x)$. Three axiomatizations of the Prelec (1998) probability weighting function (see Definition 2.22) have been proposed. These are by Prelec (1998), Luce (2001), and al-Nowaihi and Dhami (2006). The main results are stated below to give readers a flavor of the literature.

We shall find it convenient to use the following notation: $(x; p_1, p_2)$ is the lottery (x, p_1) received with probability p_2 . It is straightforward to extend this notation, iteratively to $(x; p_1, p_2, p_3)$.

Definition 2.40 (*Reduction invariance; Luce, 2001*): The probability weighting function, w , satisfies reduction invariance if, for all $x \in X$, $p_1, p_2, q \in [0, 1]$, $\lambda > 0$, $(x; p_1, p_2) \sim (x; q) \Rightarrow (x; p_1^\lambda, p_2^\lambda) \sim (x; q^\lambda)$.

Theorem 2.1 (*Luce, 2001*): The probability weighting function, w , satisfies reduction invariance if, and only if, it is the Prelec function.

Definition 2.41 (*Power invariance, al-Nowaihi and Dhami, 2006*): The probability weighting function w satisfies power invariance if, for all $x \in X$, $p, q \in [0, 1]$, $\lambda > 0$,

³³ This condition is referred to as sign-comonotonic trade-off consistency in Wakker and Tversky (1993).

³⁴ If outcomes are bundles of goods, then X can be taken to be a subset of R^n . For more complex outcomes, X can be taken to be an appropriate topological vector space.

$$(x; p, p) \sim (x; q) \Rightarrow (x; p^\lambda, p^\lambda) \sim (x; q^\lambda) \text{ and } (x; p, p, p) \sim (x; q) \Rightarrow (x; p^\lambda, p^\lambda, p^\lambda) \sim (x; q^\lambda).$$

Comparing Definitions 2.40 and 2.41, we see that *power invariance* is simpler in that it requires two probabilities (p, q) instead of three (p_1, p_2, q) . On the other hand, *power invariance* requires two stages of compounding, instead of the single stage of *reduction invariance*. By contrast, *compound invariance* of Prelec (1998) involves no compounding but four probabilities and four outcomes.³⁵ We thus have a menu of testing options.

Theorem 2.2 (*al-Nowaihi and Dhami, 2006*): *The probability weighting function, w , satisfies power invariance if, and only if, it is the Prelec function.*

Corollary 2.2 (*al-Nowaihi and Dhami, 2006*): *Power invariance is equivalent to reduction invariance of Luce (2001) and compound invariance of Prelec (1998).*

2.6.3 Foundations of the power form of utility

Recall the power form of utility given in (2.20). We now discuss the preference foundations of this utility function due to al-Nowaihi et al. (2008).

Let $(x; p)$ stand for the *simple lottery* that pays $x \in R$ with probability $p \in [0, 1]$ and 0 otherwise. Let us use slightly different notation for lotteries in incremental form (Definition 2.30). Let

$$(\mathbf{x}; \mathbf{p}) = (x_{-m}, x_{-m+1}, \dots, x_0, x_1, \dots, x_n; p_{-m}, p_{-m+1}, \dots, p_0, p_1, \dots, p_n)$$

stand for the lottery that pays $x_i \in R$ with probability $p_i \in [0, 1]$, where $\sum_{i=-m}^n p_i = 1$ and $x_{-m} \leq x_{-m+1} \leq \dots \leq x_{-1} \leq x_0 = 0 \leq x_1 \leq x_2 \leq \dots \leq x_n$. If $\mathbf{x} = (x_1, x_2, \dots, x_n)$, then $-\mathbf{x} = (-x_1, -x_2, \dots, -x_n)$ and $\mathbf{x}^r = (x_n, x_{n-1}, \dots, x_1)$. Thus $(-\mathbf{x}^r; \mathbf{p}^r)$ stands for the lottery that pays $-x_i$ with probability p_i . If each $x_i \geq 0$ and, for some i , $p_i x_i > 0$, then we call $(\mathbf{x}; \mathbf{p})$ a *positive lottery*. For instance, if $(\mathbf{x}; \mathbf{p}) = (-7, 0, 4; 0.2, 0.5, 0.3)$, then $(-\mathbf{x}; \mathbf{p}) = (7, 0, -4; 0.2, 0.5, 0.3)$ and $(-\mathbf{x}^r; \mathbf{p}) = (-4, 0, 7; 0.3, 0.5, 0.2)$.

Definition 2.42 (*Tversky and Kahneman, 1992*): *The decision maker exhibits preference homogeneity if, for all lotteries $(\mathbf{x}; \mathbf{p})$, if c is the certainty equivalent of $(\mathbf{x}; \mathbf{p})$, then, for all $k \in R_+$, kc is the certainty equivalent of $(k\mathbf{x}; \mathbf{p})$.*

Definition 2.43 below, is an attempt to extend Tversky and Kahneman's (1992) concept of loss aversion to more general lotteries.

Definition 2.43 *Loss aversion holds for positive lotteries if, for some $\lambda > 1$, $\frac{|V(-\mathbf{x}^r; \mathbf{p}^r)|}{V(\mathbf{x}; \mathbf{p})} = \lambda$, for all positive lotteries, where V is the value function in prospect theory, given in (2.24). We call λ the coefficient of loss aversion.*

In the extreme case, where all outcomes are in gains, $V(-\mathbf{x}^r; \mathbf{p}^r)$ flips all outcomes to the domain of losses. The ratio $\frac{|V(-\mathbf{x}^r; \mathbf{p}^r)|}{V(\mathbf{x}; \mathbf{p})}$ captures the loss to gain utility under PT, which is a more general measure of loss aversion. In applications of PT and empirical implementations of PT, it is common to assume that in the power form of utility (see (2.20)), $\gamma^+ = \gamma^-$ and that the weighting

³⁵ Compound invariance of Prelec (1998) is as follows. For all outcomes $x, y, x', y' \in X$, probabilities $p, q, r, s \in [0, 1]$ and compounding integers $n \geq 1$: If $(x, p) \sim (y, q)$ and $(x, r) \sim (y, s)$; then $(x', p^n) \sim (y', q^n)$ implies $(x', r^n) \sim (y', s^n)$.

function is identical in the domain of gains and losses ($w^- = w^+$). The next theorem provides the necessary justification.

Theorem 2.3 (*al-Nowaihi et al., 2008*): (a) If preference homogeneity holds for simple lotteries, then the utility function for riskless outcomes, v , takes the form given in (2.20):

$$v(x) = x^{\gamma^+}, \text{ for } x \geq 0, v(x) = -\lambda (-x)^{\gamma^-}, \text{ for } x < 0, \text{ where } \gamma^+ > 0, \gamma^- > 0, \lambda > 0. \quad (2.51)$$

Conversely, if the value function for riskless outcomes takes the form (2.51), then preference homogeneity holds for all lotteries.

(b) If the utility function (2.51) for riskless outcomes exhibits loss aversion, then $\lambda > 1$ and $\gamma^+ = \gamma^-$.

(c) If preference homogeneity and loss aversion both hold for simple lotteries, then the utility function for riskless outcomes takes the form (2.51) with $\lambda > 1$, $\gamma^+ = \gamma^-$, and $w^- = w^+$. Conversely, if the utility function for riskless outcomes takes the form (2.51) with $\lambda > 1$, $\gamma^+ = \gamma^-$, and $w^- = w^+$, then preference homogeneity holds for all lotteries and loss aversion holds for all positive lotteries.

2.7 Third generation PT and stochastic reference points

Schmidt et al. (2008) propose an extension of PT to the case of stochastic reference points. They call their descriptive theory of decision making, *third generation prospect theory*, or PT³. Their framework addresses situations where individuals are endowed with a lottery and must choose between alternative actions that typically induce different probability distributions over outcomes. For instance, the decision maker who buys a house has an endowment in which his house may suffer harm due to fire or other natural hazards with some probability. His choice of action, for instance, to buy insurance induces another probability distribution over the relevant outcomes.

Consider a set of states of nature $S = \{s_1, s_2, \dots, s_m\}$. State $s_i \in S$ has an objective probability of occurrence, $0 \leq p_i \leq 1$ such that $\sum_{i=1}^m p_i = 1$. There is a set of state-contingent consequences from one's chosen action. The consequences, in any state, lie in the compact set $X = [a, b] \in \mathbb{R}$. To fix ideas, consider each of these consequences as terminal wealth. Suppose that F is the set of all *Savage acts*, where any $f \in F$ is a mapping $f : S \rightarrow X^m$; so $f(s_i) \in X$ is the outcome received by the decision maker when the state is s_i (see Section 1.3 for *Savage acts*).

Suppose that $r \in X^m$ is some *reference Savage act*; think of the reference act as the *state-wise entitlement* expected by the decision maker. Then, for $f, g \in F$, $f \succeq_r g$ denotes the preference of the decision maker for f over g when his reference act is r . In contrast to this *state-contingent reference act*, PT considers the special case when r is a *constant act*, i.e., $r(s_i) = r(s_j)$ for any $s_i, s_j \in S$.

Schmidt et al. (2008) propose the following preference representation. Suppose that a decision maker has the reference act r and wishes to compute his utility from the act f . Let

$$v(f(s_i), r(s_i)) = v(f(s_i) - r(s_i))$$

denote the utility in state s_i under the act f when the reference act is r . The function v is given in Definition 2.31, and a parametric form is given in (2.20). Thus, v is increasing and $v(r(s_i), r(s_i)) = 0$. Order the states in S in such a way that $f(s_i) - r(s_i)$ is increasing. Let S^+ be

the set of states for which $f(s_i) - r(s_i) \geq 0$ and S^- be the set of states for which $f(s_i) - r(s_i) < 0$. Then the utility of the decision maker from the act f when his reference act is r is given by

$$V(f, r) = \sum_{s_i \in S^-} \pi(s_i; f, r) v(f(s_i), r(s_i)) + \sum_{s_i \in S^+} \pi(s_i; f, r) v(f(s_i), r(s_i)), \quad (2.52)$$

where $\pi(s_i; f, r)$ is the decision weight corresponding to the state s_i . The decision weights are constructed as in Definition 2.33, separately cumulated in the domain of gains (S^+) and losses (S^-). The utility function V represents preferences in the sense that $f \succeq_r g \Leftrightarrow V(f, r) \geq V(g, r)$. As under PT, a decision maker who has these preferences does not choose stochastically dominated options. This implies that if for $f, g \in F$, $f(s_i) \geq g(s_i)$ for all $s_i \in S$, with strict inequality in at least one state, then $V(f, r) > V(g, r)$. EU, RDU, PT are all special cases of this framework. Sugden's (2003) *reference dependent subjective expected utility theory* (RDSEU) is also a special case when $\pi(s_i; f, r) = p_i$.

Schmidt et al. (2008) suggest the following parametric form. For the utility function, v , they suggest the parametric form in (2.20) and for the probability weighting function (that feeds into the calculations of decision weights) they propose the Tversky and Kahneman (1992) probability weighting function given in (2.5) with an identical weighting function in the domain of gains and losses. As parameter values they suggest the following. The parameter of loss aversion, $\lambda \in [1, 2.5]$, the utility function parameter in (2.20), $\gamma \in [0.5, 1]$, and the weighting function parameter in (2.5), $\tau \in [0.5, 1]$.

This framework can explain (i) the WTA/WTP disparity (see Section 3.2), provided the coefficient of loss aversion, $\lambda > 1$, and (ii) preference reversals (Section 1.5.6). Let us consider the WTA/WTP disparity. We first define *non-reversible preferences*. These preferences imply that if one prefers the act f over g when the reference act is g , one must continue to do so even if the reference act is altered to f .

Definition 2.44 *The preferences of a decision maker are said to be non-reversible if $f \succeq_g g \Rightarrow f \succeq_f g$. Preferences are said to be strictly non-reversible if $f \succeq_g g \Rightarrow f \succ_f g$, where $\succ$ is the strict preference relation.*

Proposition 2.16 (Schmidt et al., 2008): (a) *Strict non-reversibility implies the WTA/WTP disparity.*

(b) *The parameterized version of PT³ satisfies strict non-reversibility if the parameter of loss aversion $\lambda > 1$.*

Proof: In this case, the reference act is the constant act (e.g., possession or non-possession of a mug in all states). Let r be a constant act such that $r(s_i) = z$ for all $s_i \in S$ and let $r \sim_g g$. Clearly z is the WTA of the act g . Now suppose that r' is another constant act such that $r'(s_i) = z'$ for all $s_i \in S$ and let $r' \sim_{r'} g$. A moment's reflection reveals that z' is the WTP for the act g . Suppose preferences are strictly non-reversible. Then $r \sim_g g \Rightarrow r \succeq_g g$ and $g \succeq_g r$. But due to strict non-reversibility, $r \succeq_g g \Rightarrow r \succ_r g$. So $r \succ_r g \sim_{r'} r'$, thus, $\text{WTA} > \text{WTP}$.

(b) Left to the reader as an exercise. ■

2.8 Stochastic reference points in PT under rational expectations

Kőszegi and Rabin (2006) argue that existing evidence in support of the status quo as a reference point comes from experiments where the status quo is expected to remain unchanged in the

future. However, when future values of the relevant variables of interest, say, income or wealth, are uncertain, then expectations about the future may be critical in the formation of the reference point. Recent evidence supports the idea that expectations can influence reference points (Ericson and Fuster, 2011). In one set of experiments, they endowed subjects with a mug and then probabilistically offered them a chance of exchanging the mug with a second object, a pen. When the probability of exchanging the mug for the pen was low, they found that subjects were less willing to exchange their mugs when they had an opportunity to do so.

We now explain the basic framework in Köszegi and Rabin (2006). Suppose that $c \in \mathbb{R}$ is the level of consumption. Different values of consumption could correspond to, say, dated values of consumption or to the levels of consumption in different states of the world. Let $r \in \mathbb{R}$ be a reference point.

Suppose that the actions of a decision maker lead to stochastic draws of consumption c with some distribution function F . Define the expected utility of the decision maker, conditional on a fixed reference point r as

$$U(F | r) = E_c u = \int u(c | r) dF(c), \quad (2.53)$$

where u is a utility function of consumption, conditional on the reference point, and F is induced by a particular action of the decision maker.

The reference point r captures the beliefs of the decision maker about the possible outcomes, hence, a fixed reference point may not be reasonable. Since the consumption outcomes are stochastic, the reference point, r , may also be thought of as stochastic (as in Section 2.7) with a distribution function, G . The reference point reflects the decision maker's most recent beliefs about future outcomes. To quote from Köszegi and Rabin (2006): "Specifically, a person's reference point is her probabilistic beliefs about the impending outcome between the time she first focused on a decision and shortly before the time of consumption."

The expected utility of the decision maker when the reference point is stochastic is given by

$$U(F | G) = E_c E_r u = \int \left[\int u(c | r) dG(r) \right] dF(c). \quad (2.54)$$

This approach is similar to theories of disappointment aversion (Bell, 1985; Loomes and Sugden, 1986; Gul 1991; see Section 2.10.2). In these theories, one compares the consumption outcomes with a reference lottery that equals the certainty equivalent of the lottery. By contrast, in the Köszegi–Rabin framework, we compare each consumption outcome with each possible element of the reference lottery as in (2.54).

In PT, the utility function depends only on *gain-loss utility*, i.e., the utility of an outcome relative to the reference point. By contrast, in the Köszegi–Rabin framework, the conditional utility, u , depends on (i) the absolute level of *consumption utility*, and (ii) *gain-loss utility*, in the following manner.

$$u(c | r) = m(c) + l(c | r). \quad (2.55)$$

Remark 2.7 In a more general model, one may consider an N -dimensional vector of consumption, $c = (c_1, c_2, \dots, c_N) \in \mathbb{R}^N$ and an N -dimensional vector of reference points, $r = (r_1, r_2, \dots, r_N) \in \mathbb{R}^N$. The analogue of (2.55) in this case is given by:

$$m(c) = \sum_{j=1}^N m_j(c_j) \text{ and } l(c | r) = \sum_{j=1}^N l_j(c_j | r_j), \quad (2.56)$$

where m_j and l_j denote, respectively, consumption utility and gain-loss utility. For $j = 1, 2, \dots, N$, each m_j is a differentiable and strictly increasing function. Both components of utility, m and l , are additively separable across the N dimensions. The gain-loss utility across each dimension is assumed identical, and it has the following form

$$l_j(c_j | r_j) = \mu (m_j(c_j) - m_j(r_j)). \quad (2.57)$$

However, for the case of (2.55) we have

$$l(c | r) = \mu (m(c) - m(r)). \quad (2.58)$$

Example 2.16 John faces the consumption lottery $(c_1, p_1; c_2, p_2)$, where $c_1 < c_2$ and $p_1 + p_2 = 1$. Denote the probability distribution over these consumption levels by F . In theories of disappointment aversion, the reference point r is given by $\bar{r}$, the certainty equivalent, hence, the distribution G is degenerate at $\bar{r}$. The decision maker's utility in this case is given by

$$U = \sum_{i=1}^2 p_i v(c_i | \bar{r}),$$

where v is the standard utility function under PT (see Definition 2.31).

However, in the Köszegi–Rabin framework, uncertainty about the future induces a stochastic reference point. Suppose that the reference point r has the distribution G , such that $r = r_1$ with probability q_1 and $r = r_2$ with probability q_2 and $q_1 + q_2 = 1$. Then John's expected utility in (2.54) is given by

$$U(F | G) = p_1 \left[\sum_{j=1}^2 q_j u(c_1 | r_j) \right] + p_2 \left[\sum_{j=1}^2 q_j u(c_2 | r_j) \right]. \quad (2.59)$$

Substituting (2.55) and (2.58) in (2.59) we get

$$\begin{aligned} U(F | G) = & \sum_{i=1}^2 p_i m(c_i) + p_1 \left[\sum_{j=1}^2 q_j \mu (m(c_1) - m(r_j)) \right] \\ & + p_2 \left[\sum_{j=1}^2 q_j \mu (m(c_2) - m(r_j)) \right]. \end{aligned} \quad (2.60)$$

The first term in (2.60) is standard expected utility and the next two terms capture gain-loss utility for each realization of the consumption level, weighted by the probability of its occurrence. For instance, in the second term in (2.60), when the realized consumption level is c_1 , John compares his consumption with both possible realizations of the stochastic reference point; it is possible that he is elated in one of the comparisons and disappointed in the other. He performs a similar comparison in the third term when realized consumption is c_2 . Notice that in this approach, probability is weighted linearly, as under expected utility theory. This assumption is contrary to the strong evidence on non-linear probability weighting that we review extensively in Part 1 (Fehr-Duda and Epper, 2012).

The properties of gain-loss utility, μ , are as follows.

Definition 2.45 The gain-loss utility, $\mu(x)$, is continuous, differentiable for all $x \neq 0$, and strictly increasing. It has the following properties.

- (i) (Reference dependence) $\mu(0) = 0$.

- (ii) (Large stakes loss aversion) For $0 < x < y$, $\mu(y) - \mu(x) < \mu(-x) - \mu(-y)$.³⁶
- (iii) (Small stakes loss aversion) Let $\mu'_+(0) = \lim_{x \rightarrow 0^+} \mu'(x)$ and $\mu'_-(0) = \lim_{x \rightarrow 0^-} \mu'(x)$. Then $\frac{\mu'_-(0)}{\mu'_+(0)} = \lambda > 1$, where λ is the parameter of loss aversion.
- (iv) (Diminishing sensitivity) For $x > 0$, $\mu''(x) \leq 0$ and for $x < 0$, $\mu''(x) \geq 0$.

The properties listed in Definition 2.45 are similar to those in Definition 2.31, except for the difference between small and large stakes loss aversion. For alternative definitions of loss aversion, see Section 2.5.2. For small stakes gambles, it is assumed that μ is linear,

$$\mu(y) = \begin{cases} \eta y & \text{if } y \geq 0 \\ \eta \lambda y & \text{if } y < 0 \end{cases}, \quad \eta > 0, \lambda > 1. \quad (2.61)$$

In applying (2.54), one has to specify how the expectations about the reference point are formed. While the appropriate specification of a reference point is ultimately an empirical question, as a first pass, Köszegi and Rabin make the extreme assumption that the expectations of the decision maker are fully consistent with rational expectations in a manner that we specify below.

The next proposition formalizes two intuitive ideas that follow from Definition 2.45. First, if one's reference point is low, then most realized outcomes are in the domain of gains, hence, the decision maker's utility will be higher. Since the reference point is stochastic, a natural formalization of a "lower reference point" is to compare two distributions of reference points in which one distribution first order stochastically dominates the other. The dominated distribution gives a higher probability of lower reference points, hence, this should be the one that the decision maker should prefer. Second, individuals have a status-quo bias in the following sense. If a decision maker prefers the consumption bundle c to c' , when the reference point is c' , then the decision maker also prefers c to c' if the reference point is c . This is related to the idea of non-reversible preferences in Definition 2.44.

Proposition 2.17 (Köszegi and Rabin, 2006): Let μ satisfy the conditions in Definition 2.45.

- (i) Suppose that the distribution of consumption is given by F . Consider two distributions of the reference point, G and G' , such that G' first order stochastically dominates G . Then $U(F | G) \leq U(F | G')$, where U is defined in (2.54).
- (ii) Suppose that $c, c' \in \mathbb{R}$ and $c \neq c'$. Then

$$u(c | c') \geq u(c' | c') \Rightarrow u(c | c) \geq u(c' | c). \quad (2.62)$$

The next example illustrates Proposition 2.17.

Example 2.17 Suppose that in Example 2.16, $c_1 = 0$, $c_2 = 1$, $p_1 = p_2 = 0.5$. Let the reference point be $r_1 = 0$ with probability q_1 and $r_2 = 1$ with probability q_2 . Suppose that μ is linear as in (2.61) and $m(x) = x$, then, from (2.60), the utility function of the decision maker is

$$U(F | G) = \frac{1}{2} + \frac{1}{2} \eta (q_1 - q_2 \lambda).$$

Now consider another distribution of reference points, G' that dominates G . Under G' , the probability of reference point $r = 1$ is $q'_2 > q_2$ and $r = 0$ is $q'_1 < q_1$. Then $U(F | G) - U(F | G') = \frac{1}{2} \eta [(q_1 - q'_1) + \lambda (q'_2 - q_2)] > 0$, as predicted by Proposition 2.17(i).

³⁶ Suppose that μ is linear such that $\mu(x) = x$ for $x > 0$ and $\mu(x) = 2x$ for $x < 0$. Let $x = 5$ and $y = 10$ then $\mu(y) - \mu(x) = 5$ and $\mu(-x) - \mu(-y) = 10$.

The proof of Proposition 2.17(ii) is relatively straightforward. Let $u(c \mid c') \geq u(c' \mid c')$ and suppose that $u(c \mid c) < u(c' \mid c)$. Then, a few simple calculations show that substituting the values of $u(c \mid c)$, $u(c' \mid c)$ lead to a contradiction.

From (2.55), $u(c \mid r) = m(c) + l(c \mid r)$. Thus, the utility function in the Köszegi–Rabin framework differs from the standard PT utility function in having the first term, $m(c)$, that depends on the absolute level of consumption. When large stakes are involved, then changes in utility are driven more by the first term, i.e., changes in marginal utility that economists are typically used to thinking about. However, when small stakes are involved, m is likely to be linear, so one may wonder if the Köszegi–Rabin utility function $u(c \mid r)$ behaves in a manner that is similar to gain–loss utility in prospect theory. This is formalized in the next result.

Proposition 2.18 (Köszegi and Rabin, 2006): Suppose that stakes are small so that consumption utility, m , is linear and gain–loss utility μ satisfies the conditions in Definition 2.45. Then, there exists a utility function v that also satisfies the conditions in Definition 2.45 such that for all $c \in \mathbb{R}$ and $r \in \mathbb{R}$ we have³⁷

$$u(c \mid r) - u(r \mid r) = v(c - r).$$

However, for large stakes, when m is non-linear, Proposition 2.18 does not hold.

A choice set, D_k , $k = 1, 2, \dots$ is a set of probability distributions over $\mathbb{R}$ (possibly degenerate), so $D_k \in \Delta\mathbb{R}$. Suppose that the decision maker expects to choose $F_k^* \in D_k$. The decision maker's beliefs are represented by a distribution Q over the choice sets $\{D_k\}_{k \in \mathbb{N}}$ such that q_k is the probability that choice set D_k occurs and F_k^* is chosen.³⁸ Thus, the decision maker's expected choice, prior to the resolution of uncertainty is

$$\sum_{j \in \mathbb{N}} q_j F_j^*.$$

We now define three different kinds of equilibrium concepts, depending on how far in advance of the resolution of uncertainty, the decision maker makes his choice. But first it is worth going through an example that clarifies the various constructions involved.

Example 2.18 Consider the example shown in Figure 2.10. In this case, there are two choice sets, D_1 and D_2 . Each choice set is a degenerate probability distribution over meals.³⁹ Choice set D_1 gives a choice of eating out at one of two restaurants A or B. Choice set D_2 gives a choice of two takeaway meals to be eaten at home—C or D. In richer descriptions, each choice could be a probability distribution over a set of restaurants or a set of takeaway meals. Suppose that the decision maker knows that if he eats out (the choice set D_1), he prefers restaurant B (decision F_1^*), and if he eats at home (the choice set D_2), he prefers takeaway C (decision F_2^*). The decision maker's beliefs are represented by a distribution Q over the choice sets $\{D_1, D_2\}$, such that D_1 and D_2 are believed to occur with respective probabilities q and $1 - q$.

³⁷ When the consumption and the reference point are drawn from $\mathbb{R}^N$ then under the conditions of Proposition 2.18, there exist utility functions $(v_1, v_2, \dots, v_N) \in \mathbb{R}^N$ that also satisfy the conditions in Definition 2.45 such that for all $c = (c_1, c_2, \dots, c_N) \in \mathbb{R}^N$ and $r = (r_1, r_2, \dots, r_N) \in \mathbb{R}^N$ we have $u(c \mid r) - u(r \mid r) = \sum_{j=1}^N v_j(c_j - r_j)$.

³⁸ We consider here the case where the decision sets are sequences, $\{D_k\}_{k \in \mathbb{N}}$. For an extension to the continuous case, $\{D_k\}_{k \in \mathbb{R}}$, see Köszegi and Rabin (2006, 2007).

³⁹ We assume that the meals can be translated into some units on the real line.




Figure 2.10 The case of two choice sets.

Consider first, situations where decisions are made sufficiently in advance of the actual outcomes, such as in insurance decisions. In this case, the individual gets enough time to be *acclimated* to the decision. If $\{F_k^* \in D_k\}_{k \in \mathbb{N}}$ are the choices, then the individual regards $\sum_{j \in \mathbb{N}} q_j F_j^*$ as the reference point (or reference distribution). In this case, the relevant equilibrium is a *choice acclimating personal equilibrium* (CPE).

Definition 2.46 (Kőszegi and Rabin, 2007, p. 1058) Suppose that $\{D_k\}_{k \in \mathbb{N}}$ is a sequence of choice sets. Then $\{F_k^* \in D_k\}_{k \in \mathbb{N}}$ is choice acclimating personal equilibrium (CPE) if for each D_k ,

$$U\left(F_k^* \mid \sum_{j \in \mathbb{N}} q_j F_j^*\right) \geq U\left(F'_k \mid q_k F'_k + \sum_{j \in \mathbb{N}/k} q_j F_j^*\right) \text{ for all } F'_k \in D_k.$$

Thus, once the decision made by the decision maker becomes the reference point, then, in a CPE, for each choice set, no other decision in the choice set can make the decision maker better off. This is the sense in which rational expectations are used to close the model. In particular, the decision maker correctly forecasts the economic environment and his own reaction in each possible state of the environment and then maximizes utility conditional on these “correct” expectations. We give several applications of a CPE below, but first, let us continue with Example 2.18 for a simple illustration of Definition 2.46.

Example 2.19 Consider Example 2.18. The selection $\{F_1 = B, F_2 = C\}$ described in Example 2.18 is a CPE if the following two conditions are met.

$$U(B \mid qB + (1-q)C) \geq U(A \mid qA + (1-q)C). \quad (2.63)$$

$$U(C \mid qB + (1-q)C) \geq U(D \mid qB + (1-q)D). \quad (2.64)$$

Notice that in each of (2.63) and (2.64), if the decision maker deviates from one of the planned choices, then the new choice becomes a part of the stochastic reference point. This is a central feature of a CPE.

In contrast to a CPE, a second possibility is that the decision maker makes his decision too close to the actual resolution of uncertainty (e.g., buying insurance for a travel plan in order to take advantage of some last minute travel deals). In this case, the relevant equilibrium notion is one of an *unacclimating personal equilibrium* (UPE).

Definition 2.47 (Kőszegi and Rabin, 2007, p. 1056) Suppose that $\{D_k\}_{k \in \mathbb{N}}$ is a sequence of choice sets. Then $\{F_k^* \in D_k\}_{k \in \mathbb{N}}$ is an unacclimating personal equilibrium (UPE) if for each D_k ,

$$U\left(F_k^* \mid \sum_{j \in \mathbb{N}} q_j F_j^*\right) \geq U\left(F'_k \mid \sum_{j \in \mathbb{N}} q_j F_j^*\right) \text{ for all } F'_k \in D_k.$$

In Köszegi and Rabin (2006, p. 1143) a UPE is simply described as a “personal equilibrium” (PE).

In Definition 2.47, the decision maker expects to choose $\{F_k^* \in D_k\}_{k \in \mathbb{N}}$ and his beliefs about the distribution of the relevant choice sets are given by Q . Taking these expectations as the reference point, in a UPE, the decision maker indeed finds it optimal to choose $F_k^* \in D_k, k \in \mathbb{N}$.

Example 2.20 Consider Example 2.18. The selection $\{F_1 = B, F_2 = C\}$ described in Example 2.18 is a UPE, if the following two conditions are met.

$$U(B \mid qB + (1 - q)C) \geq U(A \mid qB + (1 - q)C). \quad (2.65)$$

$$U(C \mid qB + (1 - q)C) \geq U(D \mid qB + (1 - q)C). \quad (2.66)$$

The main difference from the conditions in (2.63) and (2.64) is that the reference point is now attuned only to the choices $F_1 = B$ and $F_2 = C$ that occur with respective probabilities q and $1 - q$.

What is the relevant equilibrium concept when a UPE leads to multiple equilibria? Unlike cases of multiple equilibria in strategic interaction, the setting of a game against nature is much simpler. The relevant refinement, embodied in a *preferred personal equilibrium* (PPE), is that restricting to the set of UPE, the decision maker picks the sequence $\{F_k\}_{k \in \mathbb{N}}$ that gives him the highest ex-ante expected utility.

Definition 2.48 (Köszegi and Rabin, 2006, p. 1144) Suppose that $\{D_k\}_{k \in \mathbb{N}}$ is a sequence of choice sets. Then $\{F_k \in D_k\}_{k \in \mathbb{N}}$ is a preferred personal equilibrium (PPE) if $\{F_k \in D_k\}_{k \in \mathbb{N}}$ is a UPE and for all UPE sequences $\{F'_k \in D_k\}_{k \in \mathbb{N}}$ we have that

$$U\left(\sum_{j \in \mathbb{N}} q_j F_j \mid \sum_{j \in \mathbb{N}} q_j F_j\right) \geq U\left(\sum_{j \in \mathbb{N}} q_j F'_j \mid \sum_{j \in \mathbb{N}} q_j F'_j\right).$$

Let us now consider a simple example that illustrates the difference between a CPE and a UPE.

Example 2.21 Suppose that a decision maker has current wealth w and he faces the lottery $L = (w - 100, 0.5; w, 0.5)$, thus, with probability 0.5, the decision maker faces a potentially large loss. The decision maker can pay an insurance premium of 55 and obtain, instead, the lottery $L' = (w - 55, 1)$. Denote the probability distributions over outcomes in the lotteries L and L' by respectively, F and F' . Should the decision maker buy insurance?

- (a) Consider first a CPE. In this case, there is significant time between the actual decision and the resolution of uncertainty. So, if the decision maker does not buy insurance, his reference point is the lottery L and if he buys insurance, the reference point is the lottery L' . In a CPE, it is profitable to buy insurance if

$$U(F \mid F) \leq U(F' \mid F'), \quad (2.67)$$

where U is defined in (2.54) to (2.58). Let consumption utility be linear, so $m(x) = x$ and let μ be given by (2.61). Then, (2.67) can be written as

$$\frac{1}{2} \left[(w - 100) + \frac{1}{2} \eta \lambda (-100) \right] + \frac{1}{2} \left[w + \frac{1}{2} \eta (100) \right] \leq w - 55, \quad (2.68)$$

where the LHS and RHS of (2.68) give, respectively, $U(F | F)$ and $U(F' | F')$. In computing $U(F' | F')$, note that the outcome under the lottery F' is a certain outcome so there is no gain-loss utility. However, gain-loss utility plays a critical role in the computation of $U(F | F)$. When the outcome is $w - 100$, the decision maker suffers a loss relative to the reference outcome w . Analogously, when the outcome is w , there is a feeling of gain relative to the reference outcome $w - 100$.

- (b) Consider now the case of a UPE. Buying insurance is a UPE, if having initially planned to buy the insurance, the individual finds it optimal to carry out the initial plan to buy rather than not buy insurance. Hence, in a UPE, it is profitable to buy insurance if

$$U(F | F') \leq U(F' | F'). \quad (2.69)$$

$U(F' | F')$ is already given by the RHS of (2.68). However, $U(F | F')$ is 2.22 given by

$$U(F | F') = \frac{1}{2} [w - 100 - 45\eta\lambda] + \frac{1}{2} [w + 55\eta]. \quad (2.70)$$

If we are now willing to specify the values of λ and η , then we can carry out the relevant comparisons for each of the two equilibria.

Example 2.22 (Stochastic dominance and CPE⁴⁰): One undesirable feature of a CPE is that the decision maker may choose stochastically dominated options. Let us use the gain-loss function of the form given in (2.61):

$$\mu(y) = \begin{cases} 0 & \text{if } y \geq 0 \\ \hat{\lambda}y & \text{if } y < 0 \end{cases}, \hat{\lambda} = \eta(\lambda - 1) \geq 0.$$

The decision maker is loss averse if $\hat{\lambda} > 0$ and loss neutral if $\hat{\lambda} = 0$. Consider the following two lotteries

$$A = (10, p; 20, 1 - p); B = (10, 1), 0 < p < 1.$$

Clearly, lottery A first order stochastically dominates B (the distribution function for lottery A lies everywhere below that for B). In a CPE, using a condition analogous to (2.67), the decision maker prefers lottery B to lottery A if

$$p10 + (1 - p)20 - p(1 - p)\eta(\lambda - 1)10 < 10$$

or, $B \succ A$ is equivalent to

$$\frac{1}{p} < \eta(\lambda - 1).$$

The LHS attains a minimum value of 1. Thus, if $\hat{\lambda} = \eta(\lambda - 1) < 1$, then the preference for the stochastically dominated option ($B \succ A$) can never arise for any p . For this reason, applications of this framework often make such an assumption. On the other hand, if

⁴⁰ I am grateful to Fabian Herweg for suggesting this example.

$\hat{\lambda} = \eta(\lambda - 1) > 1$, then for some values of p , a decision maker chooses the stochastically dominated option in a CPE.⁴¹

What is the relation between the Köszegi–Rabin framework and PT³? Schmidt et al. (2008, p. 207) argue that when the Köszegi–Rabin framework is recast in terms of Savage acts, it has an undesirable implication. It is possible that for two acts f and g , the probabilities of each state are the same but for at least some $s_i \in S$ it is the case that $f(s_i) \neq g(s_i)$. Since the probabilities are identical across all states, the Köszegi and Rabin framework would treat these two acts as identical, yet under PT³ these acts are not identical (see (2.52)).

There are at least two issues for applied economists in testing the Köszegi and Rabin framework. First, it is not clear how much time should elapse before the decision maker has become acclimatized to his decision, hence, the choice between a CPE and a UPE can be unclear. Since the predictions of a CPE and a UPE may be different, this is an important consideration in making the theory falsifiable. A second issue is the “strength” of empirical tests. Empirical tests that simply demonstrate “expectations driven reference points” are consistent with the Köszegi and Rabin framework, but not stringent enough to reject it. For instance, such tests may not be able to distinguish between the Köszegi–Rabin framework and theories of disappointment aversion. A stringent test must, therefore, incorporate the rational expectations built into the Köszegi and Rabin framework. It is worth bearing this in mind when you read Section 2.8.3.

In the remaining subsections we give examples of a CPE and a UPE as well as the empirical evidence for this class of models.

2.8.1 An example of a choice acclimating equilibrium (CPE): The newsvendor problem

We now consider an example of a CPE in a newsvendor problem, due to Herweg (2013). At the beginning of a *trading period* (say, a day) a newsvendor orders a quantity $q \geq 0$ of newspapers at a unit wholesale price, w ; this is the only decision made by the newsvendor. The vendor is uncertain about the demand for newspapers, x , which is known to be distributed over the interval $[x, \bar{x}]$ with density f and distribution F . The retail price of newspapers is $p \geq 0$ per unit. During the trading period, the ordered quantity q cannot be replenished. Any unsold newspapers, $q - x$, are sold at a scrap price of $s \geq 0$ per unit. We assume that

$$0 \leq s < w < p. \quad (2.71)$$

Thus, the vendor always makes a profit on any newspapers that are sold and a loss on any newspapers that are unsold. This allows gain–loss utility to play an important role in the vendor’s decisions. The ex-post state-dependent profits of the vendor are

$$\pi(q, x) = \begin{cases} \pi_L(q, x) = (p - w)x + (q - x)(s - w) & \text{if } x < q \\ \pi_H(q, x) = (p - w)q & \text{if } x \geq q \end{cases}. \quad (2.72)$$

Ex-ante, the newsvendor cannot know for certain what his profits will be. However, using the distribution of possible demand levels, F , and (2.72), the expected profit of the newsvendor is

⁴¹ This may create technical problems. For instance, non-concavities of the target function (Herweg, 2013) or non-existence of a contract that implements certain actions (Herweg et al., 2010).

$E\pi(q, x) = \int_{\underline{x}}^q \pi_L(q, x) dF(x) + \int_q^{\bar{x}} \pi_H(q, x) dF(x)$. From (2.72)

$$\pi_L = (p - s)x - (w - s)q, \quad (2.73)$$

hence, we can rewrite expected profits as

$$E\pi(q, x) = \left[-(w - s)qF(q) + (p - s) \int_{\underline{x}}^q x dF(x) \right] + (p - w)q(1 - F(q)). \quad (2.74)$$

The classical solution maximizes $E\pi(q, x)$ by a suitable choice of q . The first order condition is sufficient because the objective function is strictly concave. The solution, $q = q^*$, is given by

$$F(q^*) = \frac{p - w}{p - s} \in (0, 1). \quad (2.75)$$

Suppose now that the newsvendor has the Köszegi–Rabin utility, U , defined in (2.54) to (2.58). Thus, the ex-post utility of the newsvendor for a fixed q, x, r is given by

$$U = \pi(q, x) + \mu(\pi(q, x) - r), \quad (2.76)$$

where r is the reference point and the gain–loss utility is assumed to take the particularly simple form

$$\mu(y) = \begin{cases} 0 & \text{if } y \geq 0 \\ \lambda y & \text{if } y < 0 \end{cases}, \quad \lambda \in [0, 1]. \quad (2.77)$$

The restriction $\lambda \in [0, 1]$ ensures that the newsvendor's objective function is concave. The reference point itself is determined as in the Köszegi–Rabin framework. We assume that the choices are made sufficiently in advance of the outcome so that the concept of the choice acclimating personal equilibrium (CPE) in Definition 2.46 applies.

Suppose that the newsvendor has ordered some quantity, q , and a particular level of demand x has materialized. If the decision maker has acclimatized with the decision (CPE), the reference point, r , in (2.76) is $\pi(q, z)$ where z has the distribution function $F(z)$. Thus, conditional on given levels of q, x , the utility of the newsvendor is

$$U = \pi(q, x) + \int_{\underline{x}}^{\bar{x}} \mu(\pi(q, x) - \pi(q, z)) dF(z), \quad (2.78)$$

where μ is defined in (2.77). There are two possibilities.

I. ($x < q$): For any particular choice of q , followed by any realization of demand such that $x < q$, the material utility of the newsvendor is $\pi_L(q, x) = [(p - s)x - (w - s)q]$ defined in (2.73). The reference point is given by $\pi(q, z)$ and z has the distribution function $F(z)$. There are two subcases to consider: $z < q$ and $z \geq q$. When $z < q$, μ is defined on the term $\pi_L(q, x) - \pi_L(q, z) = (p - s)(x - z)$. For any $z \geq q$, μ is defined on the term $\pi_L(q, x) - \pi_H(q, z) = (p - s)(x - q) < 0$. Thus, the gain–loss utility (second term in (2.78)) is

$$\int_{\underline{x}}^x \mu((p-s)(x-z)) dF(z) + \int_x^q \mu((p-s)(x-z)) dF(z) + \int_q^{\bar{x}} \mu((p-s)(x-q)) dF(z). \quad (2.79)$$

Using the first row of (2.77), the first term in (2.79) is zero. Using the second row in (2.77), the second term equals $-\lambda(p-s) \int_x^q (z-x) dF(z)$ and the third equals $-\lambda(p-s)(q-x)(1-F(q))$. Hence, the Köszegi–Rabin utility of the newsvendor for a particular realization of x , when $x < q$ is given by

$$U^-(q, x) = [(p-s)x - (w-s)q] - \lambda(p-s) \int_x^q (z-x) dF(z) - \lambda(p-s)(q-x)(1-F(q)). \quad (2.80)$$

II. ($x \geq q$): In this case, the realized demand exceeds the ordered supply. From (2.77), gain–loss utility $\mu(y) = 0$ if $y \geq 0$. Hence, the vendor will be able to sell his entire stock, q at a price p per unit. Thus, his material utility is given by $\pi_H = (p-w)q$ (see (2.72)) and there is no gain–loss utility. Thus, his Köszegi–Rabin utility for a particular realization of x , when $x \geq q$ is given by

$$U^+(q, x) = (p-w)q. \quad (2.81)$$

The utility levels in (2.80), (2.81) are for fixed levels of x (either side of q). However, ex ante, the newsvendor does not know the realization of x . Thus, his expected utility is $EU(q) = \int_{\underline{x}}^q U^-(q, x) dF(x) + \int_q^{\bar{x}} U^+(q, x) dF(x)$. Substituting U^- and U^+ from (2.80), (2.81), we get

$$\begin{aligned} EU(q) &= \int_{\underline{x}}^q [(p-s)x - (w-s)q - \lambda(p-s) \int_x^q (z-x) dF(z) \\ &\quad - \lambda(p-s)(q-x)(1-F(q))] dF(x) + (p-w)q(1-F(q)). \end{aligned} \quad (2.82)$$

The vendor with Köszegi–Rabin preferences chooses q in order to maximize $EU(q)$ in (2.82).⁴² Since $0 \leq \lambda \leq 1$, $EU(q)$ is strictly concave. The reader can also check that corner solutions can be ruled out. This leaves only an interior solution $q = \hat{q}$, given by the solution to the first order condition.

$$EU'(\hat{q}) = 0 \Rightarrow (p-w) - (p-s)F(\hat{q}) - \lambda(p-s)F(\hat{q})(1-F(\hat{q})) = 0. \quad (2.83)$$

Evaluating the first order condition (2.83) at $q = q^*$, where q^* is given in (2.75), we get

$$EU'(q^*) = -\lambda(p-w) \left(\frac{w-s}{p-s} \right) < 0. \quad (2.84)$$

Since $EU(q)$ is strictly concave ($EU''(q) < 0$), we get from (2.84) that $\hat{q} < q^*$, i.e., a newsvendor with Köszegi–Rabin preferences finds it optimal to order a lower quantity. This reduction in the optimal quantity is driven by the loss aversion of the newsvendor. To see this, implicitly differentiate (2.83) with respect to λ to get

⁴² When $\lambda = 0$, the solution is given by (2.75).

$$\frac{\partial \hat{q}}{\partial \lambda} = \frac{(p-s)F(\hat{q})(1-F(\hat{q}))}{EU''(\hat{q})} < 0.$$

A loss averse newsvendor hedges his bet in the face of demand uncertainty by lowering the ordered amount, q , relative to the case where no loss aversion exists. If the actual demand turns out to be lower than expected, then with a lower q , a loss averse newsvendor suffers a smaller loss in utility on account of loss aversion.

2.8.2 An example of an unacclimating personal equilibrium (UPE): The taxi drivers' problem

We consider the phenomenon of the backward bending labor supply of New York cab drivers in Chapter 3. Here, we consider a stylized model of such a situation.⁴³ A taxi driver has the material utility

$$m(w, e) = we - \frac{1}{2}e^2, \quad (2.85)$$

where w is the market determined wage level, and $e \in [0, \bar{e}]$ is the taxi driver's effort level, measured in terms of number of hours worked. For any realization of the wage level, w , the taxi driver maximizes his material utility by a suitable choice of e . Since this is a strictly concave problem over a compact domain, the solution is given by

$$e^*(w) = w.$$

Thus, the labour supply curve is upward sloping.⁴⁴

Now suppose that the taxi driver has the Köszegi–Rabin utility in which the gain–loss utility applies to the income dimension as well as to the effort dimension. Thus, we need to use Remark 2.7 for the case $N = 2$. Suppose that in the optimal solution, the driver expects to put in effort $e_H \geq 0$ when the wage is w_H , and effort $e_L \geq 0$ when the wage is w_L . We also assume that the decision is made sufficiently close to the actual realization of the outcomes, so that the relevant solution concept is an unacclimating personal equilibrium, UPE (Definition 2.47).

Definition 2.49 *The taxi driver has an income target, $\bar{w}$, if $w_H e_H = w_L e_L = \bar{w}$.*

Let us assume that the taxi driver has an income target, as specified in Definition 2.49. Since $w_H > w_L$, one implication of having an income target is that $e_H < e_L$. So the driver works less when the wage is higher.

Since the driver plans to exert effort level e_j when wage is w_j , $j = H, L$, his stochastic reference point for income is

$$r_I = \begin{cases} w_H e_H & \text{with probability } p \\ w_L e_L & \text{with probability } 1 - p \end{cases},$$

⁴³ We present a simplified version of the taxi driver problem using the unpublished notes of Fabian Herweg. See also Köszegi and Rabin (2006) for a related application.

⁴⁴ In a more general model under EU, there are well-known opposing income and substitution effects such that the labor supply can be upward or downward sloping.

and the *reference point for cost of effort* is

$$r_C = \begin{cases} \frac{1}{2}e_H^2 & \text{with probability } p \\ \frac{1}{2}e_L^2 & \text{with probability } 1 - p \end{cases}.$$

Suppose that the realized wage is w , and the driver exerts the effort level e . Then, the Köszegi–Rabin utility of the driver is:⁴⁵

$$\begin{aligned} V(e \mid w, r_I, r_C) = & \left[we - \frac{1}{2}e^2 \right] + [p\mu(we - w_H e_H) + (1-p)\mu(we - w_L e_L)] \\ & + \left[p\mu\left(-\frac{1}{2}e^2 + \frac{1}{2}e_H^2\right) + (1-p)\mu\left(-\frac{1}{2}e^2 + \frac{1}{2}e_L^2\right) \right]. \end{aligned} \quad (2.86)$$

The first term in (2.86) is material utility $m(w, e)$. The second term is gain–loss utility from income. The actual income we is compared to the reference income $w_H e_H$ with probability p , and to the reference income $w_L e_L$ with probability $1 - p$. The third term is gain–loss utility from the effort level. The actual cost of effort $\frac{1}{2}e^2$ is compared to the reference cost $\frac{1}{2}e_H^2$ with probability p , and to the reference cost $\frac{1}{2}e_L^2$ with probability $1 - p$.

In a UPE, the planned choices of the driver are consistent with actual choices. Hence, we now check if the driver would like to carry out the planned choices in each state. Suppose that the driver exerts the effort level e . Then there are three cases to consider.

- C1. $e \leq e_H < e_L$. In terms of gain–loss utility, the driver is in the gain dimension with respect to effort in both states.
- C2. $e_H < e \leq e_L$. The driver is in the gain dimension with respect to the low wage state and in the loss dimension with respect to the high wage state.
- C3. $e_H < e_L \leq e$. The driver is in the loss dimension in both states.

The second dimension of gain–loss utility arises with respect to the income level, so we need to check actual income against reference income, i.e., $w_i e \gtrless \bar{w}$, $i = H, L$. We consider the various cases separately and assume that gain–loss utility, μ , takes the linear form, given in (2.61) with $\eta = 1$.

Case I. The high wage state, w_H : Suppose that the realized wage is w_H . We need to consider each of the subcases, C1, C2, and C3.

C1. $e \leq e_H < e_L$: In this case, since actual effort is lower than both levels of planned effort, so $w_i e < \bar{w}$ for $i = H, L$. Thus, the driver is in the loss dimension with respect to income and in the gain dimension with respect to effort. Using $w_H e_H = w_L e_L = \bar{w}$, the Köszegi–Rabin utility in (2.86) is

$$\begin{aligned} V(e \mid w_H, r_I, r_C) = & \left[w_H e - \frac{1}{2}e^2 \right] + [-\lambda p(\bar{w} - w_H e) - \lambda(1-p)(\bar{w} - w_H e)] \\ & + \left[\frac{p}{2}(e_H^2 - e^2) + \frac{(1-p)}{2}(e_L^2 - e^2) \right]. \end{aligned} \quad (2.87)$$

⁴⁵ It is more convenient in this example to replace the notation for the utility function U in the Köszegi–Rabin framework with the new notation V that keeps track of the effort level and the wage level in each case.

The first term is material utility, the second is gain–loss utility in the income domain, and the third is gain–loss utility in the effort domain. From (2.87), the first and second order derivatives with respect to e are: $V' = w_H(1 + \lambda) - 2e$ and $V'' = -2 < 0$. Since this is a high wage state, w_H , a UPE requires that the driver does wish to carry out his original plan of working at least e_H hours. Since $V'' < 0$, and $V' > 0$ for small values of e , a sufficient condition for ruling out the case $e < e_H$ is that $V' |_{e=e_H} \geq 0$ or $w_H(1 + \lambda) - 2e_H \geq 0$. Also $w_H e_H = \bar{w}$, so $e_H = \bar{w}/w_H$. In conjunction, these conditions imply that a sufficient condition for ruling out $e < e_H$ is

$$V' |_{e=e_H} \geq 0 \Leftrightarrow \bar{w} \leq \frac{1 + \lambda}{2} w_H^2. \quad (2.88)$$

Higher loss aversion, λ , is more conducive to the satisfaction of (2.88).

C2. $e_H < e \leq e_L$: In this case, $\bar{w} = w_H e_H < w_H e$ and $w_H e \leq w_H e_L = \bar{w}$. Thus, in the income dimension, the driver is in the domain of gains when wage is high and in the domain of losses when wage is low. In the effort dimension, the driver is in gains with respect to the low wage reference state and in losses with respect to the high wage reference state. Thus, using $w_H e_H = w_L e_L = \bar{w}$, the analogue of (2.87) is

$$\begin{aligned} V(e | w_H, r_I, r_C) = & \left[w_H e - \frac{1}{2} e^2 \right] + [p(w_H e - \bar{w}) - \lambda(1 - p)(\bar{w} - w_H e)] \\ & + \left[-\lambda \frac{p}{2} (e^2 - e_H^2) + \frac{1 - p}{2} (e_L^2 - e^2) \right]. \end{aligned} \quad (2.89)$$

In this case, one can check that $V' = w_H[1 + \lambda - p(\lambda - 1)] - e[2 + p(\lambda - 1)]$ and $V'' < 0$. In this case, $e_H < e$, so to rule it out we require $V' |_{e=e_H} \leq 0$. A simple calculation shows that if the following condition holds, then we cannot have $e_H < e$:

$$V' |_{e=e_H} \leq 0 \Leftrightarrow \bar{w} \geq \frac{w_H^2(1 + \lambda - p(\lambda - 1))}{2 + p(\lambda - 1)}. \quad (2.90)$$

C3. ($e_H < e_L \leq e$). In this case, the driver is in the domain of gains with respect to income and in the domain of losses with respect to cost of effort in both reference states. Thus, his Köszegi–Rabin utility is

$$\begin{aligned} V(e | w_H, r_I, r_C) = & \left[w_H e - \frac{1}{2} e^2 \right] + [p(w_H e - \bar{w}) + (1 - p)(w_H e - \bar{w})] \\ & - \left[\lambda \frac{p}{2} (e^2 - e_H^2) + \lambda \frac{1 - p}{2} (e^2 - e_L^2) \right]. \end{aligned} \quad (2.91)$$

A simple calculation shows that $V' = 2w_H - e(1 + \lambda)$, $V'_{e=0} > 0$, and $V'' < 0$. Using $e_H = \bar{w}/w_H$, and in order to keep to the planned effort levels (i.e., effort e_H when wage is w_H), we require that

$$V' |_{e=e_H} \leq 0 \Leftrightarrow \bar{w} \geq \frac{2w_H^2}{1 + \lambda}. \quad (2.92)$$

However, if (2.90) holds, then so too does (2.92) because with $\lambda > 1$, $\min \left\{ \frac{2}{1+\lambda}, \frac{(1+\lambda-p(\lambda-1))}{2+p(\lambda-1)} \right\} = \frac{2}{1+\lambda}$. Hence, we can rule out any departure from the planned effort level in the high wage state by using (2.88) and (2.90). The result can be summarized as follows.

Proposition 2.19 *There exists a reference income, $\bar{w}$, satisfying*

$$w_H^2 \frac{2}{1+\lambda} \leq \bar{w} \leq \frac{1+\lambda}{2} w_H^2 \quad (2.93)$$

such that the taxi driver finds it optimal to carry out his planned effort level in the high wage state.

Case II. The low wage state, w_L .

In the low wage state, w_L , one needs to check, analogously to Case I, all three subcases in C1, C2 and C3. The analogue of Proposition 2.19 in this case is given next.

Proposition 2.20 *There exists a reference income $\bar{w}$ and functions $a(w_L, \lambda)$ and $b(w_L, \lambda)$ satisfying*

$$a(w_L, \lambda) \leq \bar{w} \leq b(w_L, \lambda), \quad (2.94)$$

such that the taxi driver finds it optimal to carry out his planned effort level in the low wage state.

We leave the proof of Proposition 2.20 to the reader. Putting Propositions 2.19 and 2.20 together we get the final result.

Proposition 2.21 *There exists a UPE with an income target, $\bar{w}$, if the target jointly satisfies (2.93) and (2.94). The state-contingent effort levels are $e_i = \bar{w}/w_i$, $i = L, H$.*

We have derived in this example the conditions under which a wage target is justified. But we have not derived the optimal wage target from the set of all wage targets. In order to compute the optimal wage target, one may proceed as suggested in the definition of a PPE (Definition 2.48). For each wage target, compute the utility from the wage target, and then choose the target that maximizes this utility.⁴⁶

2.8.3 Some empirical evidence for Köszegi–Rabin preferences

In Abeler et al. (2011), subjects perform a repetitive and tedious task that has no intrinsic value to the experimenter. They count the number of zeros in a table of 150 randomly ordered digits that are either 1 or 0. In the first stage, subjects were paid a piece rate of 10 cents for each correct table; the payment was received with probability 1. In the second task, they could decide how long they wanted to work, up to a maximum allowed time of 60 minutes.

Subjects in the second task received their accumulated earnings at the piece rate of 20 cents per correct table. However, the payment was received with probability 0.5, and with probability 0.5, they received a fixed payment, $f > 0$. In the LO treatment, $f = f_{LO} = 3$ euros and in the HI treatment, $f = f_{HI} = 7$ euros. Let w be the piece rate per correct table, e the number of correct

⁴⁶ In personal communication, Fabian Herweg points out that there exist also other UPEs that are not income target effort plans. It may well be that the PPE is one of these non-target UPEs.

tables (proxy for effort), and $\phi(e)$ the cost of effort function, such that $\phi' > 0$ and $\phi'' > 0$. Consider the predictions of three different classes of model.

1. Absence of a reference point: Consider the canonical textbook model and the following additively separable utility function

$$U(e, f, w) = [0.5u(we) + 0.5u(f)] - \phi(e),$$

where u is a continuous, differentiable, strictly increasing, and concave utility function. Individuals choose e to maximize U , given w, f . The first order condition $wu'(we) = 2\phi'(e)$ is independent of f and it can be solved for the optimal effort level, e^* , which is predicted to be independent of f . In particular, when $u(x) = x$, the solution is given by

$$\phi'(e^*) = \frac{w}{2}, \quad (2.95)$$

where the LHS is the marginal cost of effort and the RHS is the marginal benefit.

2. Exogenous reference point: Suppose that there is no reference point for effort,⁴⁷ linear weighting of probabilities for simplicity, and the reference wage is exogenously given at w_r . Then, under prospect theory,

$$V(e, f, w) = [0.5v(we - w_re) + 0.5v(f - w_r)] - \phi(e).$$

The first order condition is $(w - w_r)v' = 2\phi'(e)$. The optimal effort level is predicted to be independent of f .

3. Endogenous reference: Suppose that the reference point is given as in the Köszegi–Rabin framework, so that the utility function is defined in (2.54) to (2.58). In particular, let us restrict to the linear case, so $m(c) = c$, and gain–loss utility is also linear as in (2.61). We assume that the relevant equilibrium concept is a CPE (Definition 2.46). We also assume that there is no reference point for effort but there is a reference point for wage income. There are two cases:
 - 3.1 The decision maker chooses e such that $we < f$. Prior to the resolution of uncertainty about the state, the consumption utility is $m(we) = we$ with probability 0.5 and $m(f) = f$ with probability 0.5. Each of these possible levels of consumption with the associated probabilities also serve as the reference points.

$$r = \begin{cases} we & \text{with probability 0.5} \\ f & \text{with probability 0.5} \end{cases}.$$

Each consumption level must now be compared with this stochastic reference point, and then we must aggregate over all consumption levels (see (2.54)). Denoting the Köszegi–Rabin utility in this case by U^- , we have

$$U^-(e | r) = \frac{1}{2} \left[we - \phi(e) + \frac{1}{2}\eta\lambda (we - f) \right] + \frac{1}{2} \left[f - \phi(e) + \frac{1}{2}\eta(f - we) \right]. \quad (2.96)$$

⁴⁷ This result is unchanged even if there is a positive but exogenous reference point for effort.

In the first term in (2.96), the realization of consumption utility is we . Relative to the reference point we , the decision maker experiences no gain or loss. However, relative to the reference point f , the decision maker experiences a loss. In the second term, the realization of consumption utility is f . In this case, the decision maker feels a gain relative to the potential consumption level we , and no gain or loss relative to the consumption level, f . Let $\hat{e}_-$ be the optimal consumption level when $we < f$. The first order condition, which is sufficient, is given by

$$\phi'(\hat{e}_-) = \frac{w}{2} + \frac{w}{4}\eta(\lambda - 1) = \phi'(e^*) + \frac{w}{4}\eta(\lambda - 1), \quad (2.97)$$

where e^* is the optimal effort in the absence of a reference point, see (2.95). From (2.97), if there were no loss aversion, $\lambda = 1$, then $\phi'(\hat{e}_-) = \phi'(e^*)$. Recall from (2.61) that $\lambda > 1$, so in the presence of loss aversion, $\phi'(\hat{e}_-) > \phi'(e^*)$, thus, given $\phi'' > 0$, we have $\hat{e}_- > e^*$. By implication, $we^* < f$.

- 3.2 The decision maker chooses e such that $we > f$. Denoting the Köszegi–Rabin utility in this case by U^+ , the expression analogous to (2.96) is given by

$$U^+(e | r) = \frac{1}{2} \left[we - \phi(e) + \frac{1}{2}\eta(we - f) \right] + \frac{1}{2} \left[f - \phi(e) + \frac{1}{2}\eta\lambda(f - we) \right].$$

The optimal effort level, in this case, $\hat{e}_+$, is given by the solution to the following first order condition

$$\phi'(\hat{e}_+) = \frac{w}{2} - \frac{w}{4}\eta(\lambda - 1) = \phi'(e^*) - \frac{w}{4}\eta(\lambda - 1). \quad (2.98)$$

In this case, the marginal benefit of effort at the optimal point is lower relative to the canonical model; $\phi'(\hat{e}_+) < \phi'(e^*)$, so $\hat{e}_+ < e^*$. Thus, if e^* is high enough, so that $we^* > f$, then loss aversion serves to reduce effort below e^* . By implication, $we^* > f$.

Thus, once we allow for an endogenous reference point, the fixed payment f has an effect on effort, while in the other cases it has no effect. As f increases to a high enough level as in case 3.1, so $we^* < f$, we should, on average, observe an increase in effort under endogenous reference point. An increase in f increases the loss-aversion adjusted cost of falling behind ($we < f$), hence, individuals exert more effort. This can be stated as a testable hypothesis:

H1: Average effort is higher when $f = f_H$, relative to the case $f = f_L$.

A second implication is as follows. The presence of gain–loss utility increases effort by increasing the marginal benefit of effort if $we < f$. The magnitude of this increase depends on the strength of loss aversion, λ . Gain–loss utility also reduces the marginal utility of effort when $we > f$ and, so, reduces optimal effort. In sum, loss aversion pushes effort in the direction of $we = f$. In the limit, as $\lambda \rightarrow \infty$, the first order conditions give $w\hat{e}_- = w\hat{e}_+ = f$; thus it is always optimal for the decision maker to stop exerting effort as soon as wage income equals f . Thus, f has the most powerful effect on effort in this case. This gives rise to the second testable hypothesis.

- H2:** The probability of stopping at $we = f_{LO}$ is higher when $f = f_{LO}$ relative to the case $f = f_{HI}$. Similarly, the probability of stopping at $we = f_{HI}$ is higher when $f = f_{HI}$ relative to the case $f = f_{LO}$. These effects are enhanced by an increase in loss aversion.




Figure 2.11 Histogram of accumulated earnings (in euros) at which a subject stopped.

Source: Abeler et al. (2011), with permission from the American Economic Association.

Figure 2.11 shows the histogram of earnings at which individuals stopped for two different treatments, $f_{LO} = 3$ and $f_{HI} = 7$. Individuals make fairly dispersed choices. However, there are sharp spikes at $f_{LO} = 3$ in the LO treatment, and at $f_{HI} = 7$ in the HI treatment; 15% and 16.7% of the subjects make these respective choices. These are also the modal choices in each treatment and are statistically different. By contrast, in the LO treatment, only 3.3% stop at an earnings of 7 euros and in the HI treatment, only 1.7% stop at an earnings of 3 euros.

The results are even sharper when we consider stopping values in some range that is close to the predicted value. For instance, 30% of the subjects stop in the earnings range of 2–4 euros in the LO treatment and 38.3% stop in the earnings range of 6–8 euros in the HI treatment. Average effort levels are higher under the HI treatment as compared to the LO treatment. Using a proxy for loss aversion, based on binary lottery choices, subjects who exhibit higher loss aversion are more likely to stop when their earnings are closer to f . The authors take their results as support for Köszegi–Rabin preferences and for the importance of expectations in the formation of reference points.

In the experiments of Abeler et al. (2011), subjects were not told that they were randomly assigned to one of two cases, $f = 3$ or $f = 7$. Hence, they might have inferred the value of their likely effort from the fixed payment, f , giving rise to treatment effect. Ericson and Fuster (2011) correct for this feature and perform new experiments to test the Köszegi–Rabin framework. Their first experiment is described in Figure 2.12. Experimental subjects are endowed with a mug. They then get an opportunity (Op) to exchange the mug for a pen with probability p . With probability $1 - p$, there is no opportunity to exchange (Nop) and they end the experiment with a mug. If there is an opportunity to exchange, then they can exchange with a pen or elect to keep the mug. The decision maker has to commit ex-ante (prior to getting the opportunity to exchange) if he would like to exchange the mug for a pen.




Figure 2.12 Decision tree for experiment 1 in Ericson and Fuster (2011).

Source: Keith M. Marzill Ericson and Andreas Fuster, "Expectations as endowments: evidence on reference-dependent preferences from exchange and valuation experiments," *The Quarterly Journal of Economics* (2011) 126(4): 1879–907, by permission of Oxford University Press.

Let u_P be the utility from owning a pen and u_M be the utility from owning a mug; we assume $u_P > u_M$. Under expected utility, committing ex-ante to exchange the mug for the pen gives the expected utility $pu_P + (1 - p)u_M$ and committing not to exchange gives u_M . Thus, exchange is better than no-exchange if $u_P > u_M$, which is independent of p .

Under expected utility and status-quo reference points, we can show that p does not have an effect on the decision to exchange. Here, the reference point is u_M because the decision maker is endowed with a mug. If the decision maker commits to exchange, his ex-ante utility is $p(u_P - u_M) + (1 - p)(u_M - u_M)$, and if he does not, it is $u_M - u_M = 0$. Thus, exchange is preferred if $u_P > u_M$, which is independent of p .

Under Köszegi–Rabin utility, we show below that p plays a role in the exchange decision. Hence, one can use experiments to compare the relative predictions of these theories. For pedagogical simplicity, we choose the CPE concept.

I. Suppose that the decision maker plans, ex-ante, to exchange the mug for the pen.

Let the realized utility of the decision maker be u_i , $i = P, M$ at any of the terminal nodes in Figure 2.12 (at the upper node in Figure 2.12, this is u_P , and at the lower node it is u_M). Then the Köszegi–Rabin utility, at that node, taking into account the stochastic reference point, r (u_P with probability p and u_M with probability $1 - p$), is

$$U(u_i | r) = u_i + [p\mu(u_i - u_P) + (1 - p)\mu(u_i - u_M)], \quad (2.99)$$

where μ is given in its linear form in (2.61). For ease of exposition, we shall assume $\eta(\lambda - 1) > 1$ (this can be easily implemented in experiments).

The ex-ante Köszegi–Rabin utility at the initial node in Figure 2.12, given the decision to “exchange,” is $U^{Ex} = pU(u_P | r) + (1 - p)U(u_M | r)$, where U is defined in (2.99). Thus, using μ , given in (2.61), and $u_P > u_M$, we get

$$U^{Ex} = p[u_P + (1 - p)\eta(u_P - u_M)] + (1 - p)[u_M - p\lambda\eta(u_P - u_M)],$$

or,

$$U^{Ex} = u_M + p(u_P - u_M)(1 - (1 - p)\eta(\lambda - 1)). \quad (2.100)$$

II. Suppose the ex-ante plan is not to exchange the mug.

In this case, the decision maker is always left with a mug in the end. The CPE equilibrium concept requires that the reference point in this case is u_M in all states. Thus, gain-loss utility is zero. The ex-ante Köszegi–Rabin utility in the case of “no-exchange” is simply $U^{Nex} = u_M$. Thus, the ex-ante plan to exchange the mug dominates if $U^{Ex} > U^{Nex}$, or

$$u_M + p(u_P - u_M)(1 - (1 - p)\eta(\lambda - 1)) > u_M.$$

or,

$$p > p^* = 1 - \frac{1}{\eta(\lambda - 1)}. \quad (2.101)$$

By assumption, $\eta(\lambda - 1) > 1$, hence, $0 < p^* < 1$. From (2.101), as the probability of exchanging the mug, p , increases, the decision maker is more likely to choose to exchange. Thus, in equilibrium, a higher expectation of exchange does lead to greater exchange. By contrast in the other models, p should not influence the exchange decision.

Köszegi and Rabin (2007) show that a decision maker is more risk averse when he anticipates a risk and has the possibility to insure against that risk, relative to the case where he does not anticipate the risk. A related result can be found in Köszegi and Rabin (2009) in a life-cycle consumption model. In this model, an increase in consumption today reduces the expectation of consumption tomorrow, hence, it reduces the reference point tomorrow. This is in contrast to models of habit formation where an increase in consumption today increases the reference point for consumption tomorrow by increasing the habitual consumption level.

In the Ericson and Fuster (2011) experiments, two treatments were considered: Treatment T^L where $p = 0.1$ and treatment T^H where $p = 0.9$. The findings are as follows. In treatment T^L , 22.% of the subjects chose to exchange the mug, while in treatment T^H , this figure is 55.6%; these differences are statistically significant.⁴⁸ Probit regressions that include treatment and gender indicators also confirm this result. In Treatment T^L , where $p = 0.1$, the decision maker who wishes to exchange his mug experiences a greater risk so the typical decision maker appears more risk averse by choosing not to exchange; this exemplifies the remark on the Köszegi and Rabin (2007) result alluded to in the last paragraph.

Heffetz and List (2014) modify several aspects of the experimental design in Ericson and Fuster (2011). They find that the expectations effects predicted by the Köszegi–Rabin framework are quite weak.

2.9 Limitations of PT

In this section, we consider evidence that contradicts PT in certain contexts. This evidence also typically contradicts all other mainstream decision theories in economics, such as EU and RDU.

⁴⁸ See also Smith (2008) for experimental evidence that supports the prediction of an endowment effect based on the Köszegi–Rabin framework.

2.9.1 The evidence and critique from configural weights models

We first consider the work of Michael Birnbaum and his associates who have conducted experiments that contradict the predictions of OPT (due to Kahneman and Tversky, 1979). Many of these experiments exploit the heuristics used in the editing phase of OPT (see Section 2.11.1), which are not a part of PT (due to Tversky and Kahneman, 1992). The following terminology should be helpful in understanding the experimental results.

1. (Branches) For any lottery, such as $L = (x_1, p_1; x_2, p_2; \dots; x_n, p_n)$, the i th *branch* is the ordered pair (x_i, p_i) . In a two branch lottery $L = (x_1, p_1; x_2, p_2)$, $x_1 < x_2$, (x_2, p_2) is sometimes known as the *upper branch* and (x_1, p_1) is known as the *lower branch*.
2. (Coalescing) Under EU, RDU, and PT if the same outcome occurs more than once in a lottery, we can simply add up the corresponding probabilities and write a lottery in *coalesced* form. For instance, the lottery $(50, 0.2; 50, 0.3; 75, 0.5)$, also known as a *split lottery*, can be written in the coalesced form $(50, 0.5; 75, 0.5)$. Under EU, RDU, PT, the decision maker should be indifferent among the split and the coalesced forms. Depending on which branch of a lottery is split up we get *upper coalescing* (branch with the highest outcome is split) and *lower coalescing* (branch with the lowest outcome is split).
3. (Consequence monotonicity) If one outcome in a gamble is improved, then this should be preferred by the decision maker. Thus, if more is preferred to less, then the lottery $(200, 0.5; 0, 0.5)$ should be preferred to $(100, 0.5; 0, 0.5)$.
4. (Branch independence) Suppose that there are two lotteries L_1, L_2 such that (1) $L_1 \succeq L_2$, (2) there is a common branch in the two lotteries, say, (x_i, p_i) for some i , and (3) the outcomes in each lottery are all distinct. Branch independence implies that if, in both lotteries, we replace branch i with another branch (y_i, p_i) , $x_i \neq y_i$ it is still the case that $L_1 \succeq L_2$. Consider the following additional property: (4) All corresponding probabilities are identical, i.e., for any i if (x_i, p_i) is the i th branch of L_1 and (y_i, q_i) is the i th branch of L_2 , then $p_i = q_i$. When all four properties hold, we have the case of *restricted branch independence* (RBI).

Experiments indicate that coalescing lotteries can alter choices among lotteries.⁴⁹ Consider, for instance, the choice between the following two pairs of lotteries.

$$A = (50, 0.05; 50, 0.10; 100, 0.85) \text{ versus } B = (7, 0.05; 100, 0.10; 100, 0.85) \\ A' = (50, 0.15; 100, 0.85) \text{ versus } B' = (7, 0.05; 100, 0.95)$$

In this pair of choices, if coalescing holds, then lotteries A, A' are identical and lotteries B, B' are also identical. Under EU, RDU, or PT, $A \succeq B$ implies $A' \succeq B'$. In contrast, for 200 subjects, Birnbaum (2004a) found that for 63% of the subjects, $B \succeq A$ and for 80% of the subjects, $A' \succeq B'$, which violates EU, RDU, and PT for a majority of the subjects.

Now consider a choice between the following two lotteries, C, D, where C first order stochastically dominates D,

$$C = (12, 0.05; 14, 0.05; 96, 0.90) \text{ versus } D = (12, 0.10; 90, 0.05; 96, 0.85).$$

⁴⁹ See, for instance, Birnbaum (1999, 2004a, 2007), Humphrey (1995, 1998, 2000), Starmer and Sugden (1993).

Birnbaum and Navarrete (1998) found that 70% of the undergraduate students in their sample chose D over C.⁵⁰ Birnbaum (2004b) found that these results are not due to subject error in reporting their true preferences. Further, he showed that splitting the branches of a tree in a particular manner can reduce violations of stochastic dominance. Consider, for instance, C' , D' that, in coalesced form, are identical to C , D respectively (so C' also first order stochastically dominates D'),

$$C' = (12, 0.05; 14, 0.05; 96, 0.85; 96, 0.05) \text{ versus } D' = (12, 0.05; 12, 0.05; 90, 0.05; 96, 0.85).$$

In C' , relative to C , the top branch has been split, while in D' , relative to D , the bottom branch is split. The splitting is done in such a way that the corresponding probabilities of all branches, when outcomes are ordered monotonically, are identical. This is known as the *canonical split form*. Birnbaum (2004b) found that 71% of the subjects violated stochastic dominance in a choice between C , D but only 5.6% did so when the choice was between C' , D' .

Birnbaum (2008) reports the following two pairs of lotteries from Birnbaum (2001).⁵¹

$$E = (0, 0.50; 68, 0.07; 92, 0.43) \text{ versus } F = (0, 0.52; 92, 0.48).$$

$$E' = (0, 0.50; 68, 0.07; 97, 0.43) \text{ versus } F' = (0, 0.52; 92, 0.05; 97, 0.43).$$

RDU and PT imply that $F \succeq E$ if and only if $E' \succeq F'$ (the reader should check this). In contrast, out of a sample of 1438 respondents, for 66%, $E \succeq F$, while for 62%, $E' \succeq F'$, which contradicts RDU/PT.⁵²

Birnbaum and Chavez (1997) find violations of restricted branch independence (RBI) but the direction of these violations is opposite to that predicted by PT.

2.9.2 Violations of gain-loss separability

Consider a lottery, as in (2.21), that is expressed in incremental form,

$$A = (y_{-m}, p_{-m}; y_{-m+1}, p_{-m+1}; \dots; y_{-1}, p_{-1}; 0, p_0; y_1, p_1; y_2, p_2; \dots; y_n, p_n),$$

where the reference outcome relative to the reference point is zero. We shall make use of the observation that in prospect theory, $v(0) = 0$, so we can assign any probability to an outcome of zero. Let the sum of probabilities in the domain of losses and gains be denoted, respectively, by $p_- = \sum_{i=-1}^{-m} p_{-i}$ and $p_+ = \sum_{j=0}^n p_j$. Then, denote by A^- and A^+ , respectively, the negative and positive components of A , having adjusted for the probability of a zero outcome in each case such that the respective probabilities in each component add upto 1.

$$A^- = (y_{-m}, p_{-m}; y_{-m+1}, p_{-m+1}; \dots; y_{-1}, p_{-1}; 0, p_+) , A^+ = (0, p_-; y_1, p_1; \dots; y_n, p_n).$$

The probabilities in each of the lotteries in incremental form, A^- and A^+ , add up to one. Using Definition 2.34, we can write the value function for lottery A as

⁵⁰ Birnbaum et al. (1999) found similar results on a new dataset.

⁵¹ These lotteries were designed to test the property of *upper tail independence* in Wu (1994).

⁵² See also Birnbaum (2005) for newer experiments that also arrive at similar findings.

$$V(A) = V(A^-) + V(A^+), \quad (2.102)$$

where $V(A^-)$ and $V(A^+)$ give the value function under PT corresponding to the two lotteries A^- and A^+ . Consider another lottery in incremental form, B , whose positive and negative components are respectively B^- and B^+ . The decision maker is given a choice between the two pairs of prospects (i) A^+ and B^+ , and (ii) A^- and B^- . Suppose that the following preference pattern is observed,

$$A^+ \succeq B^+, A^- \succeq B^-. \quad (2.103)$$

(2.103) implies that $V(A^+) \geq V(B^+)$ and $V(A^-) \geq V(B^-)$. Adding these two inequalities and using (2.102) we get

$$V(A) \geq V(B) \Leftrightarrow A \succeq B. \quad (2.104)$$

The pattern of preferences in (2.103), (2.104) constitute the property of *gain-loss separability* and it is satisfied by EU/RDU/OPT/PT. Wu and Markle (2008) give experimental evidence to reject gain-loss separability. They took the following lotteries

$$\begin{aligned} A^+ &= (0, 0.5; 4200, 0.5), \quad B^+ = (0, 0.25; 3000, 0.75) \\ A^- &= (0, 0.5; -3000, 0.5), \quad B^- = (0, 0.75; -4500, 0.25) \end{aligned}$$

In a within-subjects experimental design, they found that for a clear majority of subjects $B^+ \succeq A^+$, $B^- \succeq A^-$, yet for a slight majority of subjects, it was found that $A \succeq B$.

2.9.3 Interaction between probability weights and stake sizes

The *stake size effect* refers to the observed change in relative risk aversion as the stake size in the domain of gains goes up; the idea goes back at least to Markowitz (1952). Empirical evidence suggests that the stake size effect arises from probability weighting rather than from the utility function. This calls into question the assumed separability between probabilities and outcomes that is assumed in EU, RDU, and PT

The stake size effect has been verified in many studies but the stakes in these studies have typically not been very high or the high stakes have been hypothetical.⁵³ However, in the context of decision theory, doubts have been expressed about the use of hypothetical stakes relative to incentive compatible elicitation mechanisms.⁵⁴

Not surprisingly, some of the high stakes experiments have come from subjects in developing countries; these studies report a significant stake effect in the domain of gains (Binswanger, 1981; Kachelmeier and Shehata, 1992). Post et al. (2008) consider data with very high stakes from the TV show *Deal or no Deal* and they report changes in risk aversion depending on the initial expectations and changes in the resolution of uncertainty. They offer support for PT over EU in explaining the observed facts.

⁵³ See, for instance, Hogarth and Einhorn (1990), Bosch-Domenech and Silvestre (1999), and Weber and Chapman (2005).

⁵⁴ See, for instance, Camerer and Hogarth (1999), Beattie and Loomes (1997), and Holt and Laury (2002, 2005).

What is the source of the stake size effect? Most empirical studies have attempted to identify the source as the utility function. However, some empirical studies have identified the role of probability weighting.⁵⁵ An important advance was made by Fehr-Duda et al. (2010) who dispensed with hypothetical rewards and considered lotteries in the domain of gains and losses. We now consider their work.

Experiments were conducted in Beijing in 2005 with 56 two-outcome lotteries—28 in the domain of gains and 28 in the domain of losses. One of the lotteries each in the domain of gains and losses was then randomly selected and played for real. In order to allow for lotteries in the loss domain, subjects were given initial endowments that covered the maximum possible loss. Outcomes and probabilities varied widely in the experiment and high stake lotteries were constructed from the low stake lotteries by scaling up the amounts by integer values.

It was assumed that the subjects used RDU (which nests EU as a special case). The functional forms used were as follows. The utility is of the power form as in (2.20), and the probability weighting function is the two-parameter Goldstein–Einhorn function given in (2.7). A finite mixture model was used to detect the presence of decision makers with heterogeneous preferences.⁵⁶

Substantial heterogeneity in choices was found. About 27% of the subjects were expected utility types (labeled EUT types in Figure 2.13); they exhibited linear probability weighting functions and linear utility functions, so they are essentially expected value maximizers. The remaining 73% were non-expected utility types (labeled Non EUT types in Figure 2.13) and their probability weighting functions departed significantly from linearity. The experimental results for the two types of decision makers, separated by the gain and loss domains, and by stakes (high and low), are shown in Figure 2.13. On the vertical axis we have the relative risk premium (see 2.1). The main results arising from Figure 2.13 are as follows.

1. Fourfold classification of risk: The fourfold classification of risk that is explained above in Section 2.4.2 can be seen to hold. For each type of decision maker, in the domain of gains, for high and low stakes gains that are of low probability, we have $RRP < 0$ (see (2.1) for the definition of RRP), hence, the decision maker is risk seeking. However, for higher probabilities in the domain of gains, for both EU and non-EU types, $RRP > 0$, so the decision maker is risk averse. The reverse picture emerges in the loss domain. Here, for both high and low stakes, $RRP > 0$ for small probabilities (risk aversion) and $RRP < 0$ for high probabilities (risk seeking).
2. The stake size effect: For any probability, p , in Figure 2.13, the stake size effect can be seen to hold for both types of decision makers (EU and non-EU) in the domain of gains. In the domain of gains, RRP is higher at each level of probability (except $p = 0.95$) when stakes are high as compared to when the stakes are low. However, no clear pattern on the stake size effect emerges in the domain of losses.

What causes the stake size effect to be prominent for gains but not for losses? The curvature of the utility function captured by the parameter γ does not vary with the stake size for both types of decision makers. Statistically, one cannot reject the hypothesis that the curvature of the utility function is stake-invariant. Since the only other component of the objective function is probability weights, it seems likely that this causes the stake size effect. This intuition is confirmed statistically.

⁵⁵ See, for instance, Kachelmeier and Shehata (1992), Tversky and Kahneman (1992), Camerer (1992), and Etchart-Vincent (2004).

⁵⁶ The model is not designed to pick out prospect theory preferences because there are no mixed gambles, so loss aversion cannot be detected.




Figure 2.13 Median relative risk premium (RRP) for different types of decision makers as the probability, p , is varied.

Source: *Journal of Risk and Uncertainty*, 40(2): 147–80 (2010) “Rationality on the rise: why relative risk aversion increases with stake size,” Helga Fehr-Duda, Adrian Bruhin, Thomas Epper, and Renate Schubert, Figure 2, © Springer Science + Business Media, LLC 2010, with kind permission from Springer Science and Business Media.

The Goldstein–Einhorn function is a two-parameter family of probability weighting functions, hence, the inference about the contribution of probability weighting to the stake size effect must be based on the joint contribution of the two parameters, δ and η (see equation (2.7)). This is achieved by constructing 95% confidence bands based on the percentile bootstrap method, and using 2000 replications.

The probability weighting function for the EU types is approximately linear, so in Figure 2.14 we report only the results in the gain and loss domains for the non EU types. The probability weighting functions for high and low stakes in the domain of losses are quite similar, but there is a more marked difference in the domain of gains. In terms of the estimated parameters in the domain of gains, δ , which captures the elevation of the Goldstein–Einhorn function, decreases from 1.304 to 0.913 as stakes increase. This corresponds to a lower degree of optimism. At the same time, η goes up with stakes, which mitigates slightly the inverse S-shape of the weighting function.

In the lower panel of Figure 2.14, if the confidence band includes the zero line, then we cannot reject the null hypothesis of stake-independent probability weights; this is the case in the




Figure 2.14 Stake dependent probability weights in the domain of losses and gains (upper panel) and 95% confidence bands based on the percentile bootstrap method (lower panel).

Source: *Journal of Risk and Uncertainty*, 40(2): 147–80 (2010) “Rationality on the rise: why relative risk aversion increases with stake size,” Helga Fehr-Duda, Adrian Bruhin, Thomas Epper, and Renate Schubert, Figure 3, © Springer Science + Business Media, LLC 2010, with kind permission from Springer Science and Business Media.

domain of losses. There is also no noticeable difference between probability weights by stake in the domain of losses. However, in the domain of gains we can reject the hypothesis that the probability weighting function is stake independent (lower panel of Figure 2.14).

Not much is known about the determinants of probability weights but there have been several suggestions in the literature. These include differences in probability weights in the domain of gains and losses,⁵⁷ dependence on the timing of the resolution of uncertainty,⁵⁸ role of emotions,⁵⁹ and the *affect heuristic*.⁶⁰ Fehr-Duda et al., (2011) consider the effect of *incidental*

⁵⁷ See, for instance, Abdellaoui (2000) and Bruhin et al. (2010).

⁵⁸ See, for instance, Abdellaoui, Diecidue, and Öncüler (2011), and Noussair and Wu (2006).

⁵⁹ See, for instance, Walther (2003) who showed that anticipation of elation and disappointment lead to departures from linear probability weighting. The literature on regret aversion and disappointment aversion also identifies the effect of anticipated emotions on the shape of the probability weighting function; see, for instance, Bell (1982), Loomes and Sugden (1986), and Gul (1991). For other empirical evidence, see Fehr-Duda et al. (2011).

⁶⁰ See, for instance, Rottenstreich and Hsee (2001).

moods on probability weighting by estimating a probability weighting function of the Goldstein–Einhorn form. Their subjects are asked if their day has been better or worse than normal on a scale of 0 (worse) to 5 (promising) and their probability weights for lotteries are elicited. Women who were in a better mood tended to weight probabilities more optimistically. No effect was found for men, most of whom behaved in a manner consistent with expected utility maximization.

2.9.4 Heterogeneity in preferences

Given a choice between two lotteries, non-stochastic versions of EU and most behavioral decision theories predict a preference for either of the two lotteries (unless one is indifferent between the two). In contrast, in most experiments, we observe heterogeneity in choices in the sense that there is a fraction of subjects who prefer each lottery. Experimental economists have tried to accommodate this stylized fact by allowing for subjects in experiments to make noisy decisions. The interested reader can consult Hey (2014) for a detailed and readable account of the various issues that we only briefly touch on here.

Suppose that subjects in experiments are asked to make a choice between two lotteries, L_1 and L_2 . If the utility function V represents the decision maker's preferences in some decision theory, then we should observe L_1 to be chosen over L_2 if $V(L_1) > V(L_2)$. In this case, one typically conducts a test of differences in proportions to determine if the predictions of the relevant theory hold (Kahneman and Tversky, 1979). However, for whatever reasons, subjects in experiments may not state their preferences accurately, so they may exhibit a preference of L_1 over L_2 if $V(L_1) > V(L_2) + \varepsilon$, where ε is some random variable. This error term may be assumed to be common across all subjects in experiments (Camerer, 1989), or might vary by subjects in experiments (Carbone and Hey, 1994), or might even vary by the questions asked in experiments.

By assuming a specific error structure, one can engage in more sophisticated testing such as maximum likelihood rather than a simple test of proportions. For instance, Hey and Orme (1994) assumed that ε is normally distributed, while Carbone and Hey (1994) assume that ε has an extreme value distribution. Hey's (2014, p. 821) reading of the evidence based on experiments, which allow for noisy choices is this: "SEU is not bad, and where it is, then it is Rank Dependent that is better." But in other studies, there is strong support for prospect theory (Kothiyal et al., 2014; Abdellaoui, 2000; Erev et al., 2010).

One complication in models that are based on a uniform error structure across individuals is that the propensity to make decision errors varies with observed characteristics (von Gaudecker et al., 2011). Andersen et al. (2012) show that ignoring such heterogeneity may result in spurious correlation between background variables and risk attitudes. In particular, they show that cognitive ability is related to noisy decision making and that this can cause both positive and negative relationships between risk attitudes and cognitive ability, depending on the construction of the elicitation task.

An even more satisfactory method of testing is to allow for a mixture of types in the experimental population who follow different behavioral decision theories. In addition, subjects may choose with noise. Identifying heterogeneity in types, based on a *mixture model approach* has been an active area of research in recent years. In Fehr-Duda et al. (2010) that we considered above (Section 2.9.3), the evidence shows a mixture of types in the population. About 23% are expected value maximizers, while the behavior of the remaining 73% is consistent with non-linear probability weighting.

Several recent papers ask the question: Is there a mixture of individuals in the population who follow different decision theories?⁶¹ Harrison and Rutström (2009) calculated the probability that choices made by decision makers in their experiments are consistent with EU or with PT. They found that each of these two decision theories is equally likely to account for the observed choices. However, the authors recognize that these proportions are likely to vary with the task, context, and demographics of the subjects in experiments.⁶²

Bruhin et al. (2010) use a finite mixture model and run three sets of experiments—two in Zurich and one in China. Across the three experiments, they find that about 20% follow expected value maximization, or EU maximization with risk neutrality. The remaining 80% exhibit non-linear probability weighting; 30% exhibit pronounced non-linear probability weighting and 50% somewhat less pronounced weighting. These proportions are roughly similar across the three experiments. Interestingly, these proportions (20% EU and 80% PT) are very similar to those found by Conte et al. (2011) who also use finite mixture models for British experimental subjects in experiments, although there are several differences in the procedures followed by the two studies.

Figure 2.15 shows the estimated probability weighting functions for the student and the representative population in the study by Fehr-Duda and Epper (2012) that we have described above. For the student population (left panel of Figure 2.15), 80% exhibit a pronounced departure from linearity, of the inverse S-shaped form (Type I), while the remaining appear to be expected value maximizers (Type I). For the representative sample (right panel of Figure 2.15), 67% exhibit the inverse S-shaped form and have a more elevated probability weighting function (Type I), so they




Figure 2.15 Two types of decision makers for the student and the representative population.

Source: Reproduced with permission of *Annual Review of Economics*, Volume 4, Helga Fehr-Duda and Thomas Epper, "Probability and risk: foundations and economic implications of Probability-dependent risk preferences," © 2012 by Annual Reviews <http://www.annualreviews.org>

⁶¹ See Hey and Orme (1994) for an earlier paper that revealed substantial heterogeneity between individuals.

⁶² For evidence from diverse geographical regions, see Harrison et al. (2010).

are more optimistic. The remaining 33% also exhibit a pronounced departure from linearity but they underweight probabilities significantly over a large range of higher probabilities (Type II). There are no expected value types in the representative sample. These results make a strong case for moving away from models of decision theory with linear probability weighting.

Mixture models would seem to have a better chance of explaining the dispersion in experimental results relative to any of the theories on their own. Heterogeneity in risk preferences can also arise in other dimensions. For gender based heterogeneity in risk attitudes, see Eckel and Grossman (2008), and Croson and Gneezy (2009); women are often more risk averse than men.⁶³ Dohmen et al. (2012) consider the intergenerational transmission of risk attitudes. Dohmen et al. (2011) uncover substantial heterogeneity of risk attitudes that depend on gender, age, parental background, and even height. Dohmen et al. (2010) find that lower cognitive ability is correlated with higher risk aversion as well as greater time impatience.

2.10 A selection of other behavioral theories

In this section, we consider some of the other important theories of behavioral decision making.

2.10.1 Regret theory

The ideas of *regret*, and its counterpart, *rejoice*, are not directly incorporated into EU/RDU/OPT/PT. In a recent book, Kahneman (2012, p. 288) lists the absence of the emotions of regret and disappointment as one of the blind spots of prospect theory. *Regret theory* tries to plug this gap.⁶⁴

Suppose that there are n states of the world, $s_1, s_2, \dots, s_n$, that can arise, respectively, with the objective probabilities $p_1, p_2, \dots, p_n$. Suppose that a decision maker can take two possible actions, A_1 and A_2 , each of which induces an outcome for each possible state. Let action A_1 induce outcomes $x_1, x_2, \dots, x_n$ and action A_2 induce the outcomes $y_1, y_2, \dots, y_n$. Suppose that outcomes are real numbers, say, levels of wealth.

Each state s_i is associated with a regret-rejoice function $R(x_i, y_i)$, where x_i, y_i are the respective outcomes under actions A_1 and A_2 . When $R > 0$, the decision maker rejoices, and when $R < 0$, the decision maker regrets. If action A_1 is taken, the function $R(x_i, y_i)$ has the following properties:⁶⁵

- P1. $R(x_i, y_i)$ is positive if $x_i > y_i$, negative if $y_i > x_i$ and zero if $x_i = y_i$.
- P2. $R(x_i, y_i)$ is increasing in the first argument and decreasing in the second.
- P3. R is skew symmetric (or antisymmetric), i.e., $R(x_i, y_i) = -R(y_i, x_i)$.
- P4. (Regret aversion) If $x > y > z > 0$ (where “ $>$ ” means greater than), then

$$R(x, y) + R(y, z) < R(x, z). \quad (2.105)$$

⁶³ For a review of several studies of gender differences in risk seeking, which support the finding that men are more risk taking than women, see Charness and Gneezy (2012).

⁶⁴ Important early ideas that formed the basis of subsequent developments in this area are Bell (1982), Loomes and Sugden (1982), and Fishburn (1982a).

⁶⁵ Not all of these properties appeared in the early versions of regret theory. For instance, P2 and P4 were introduced by Loomes and Sugden (1987).

From P1, in state s_i , the decision maker rejoices from taking action A_1 when it gives a relatively higher outcome ($x_i > y_i$). If $x_i > y_i$, rejoicing increases as the gap between x_i and y_i widens (property P2). P3 imposes strict symmetry into the model. The early formulations of regret theory allowed only for pairwise comparisons among tuples of the form x_i, y_i . Hence, transitivity could be easily shown to be violated. Loomes and Sugden (1987) then introduced property P4 that enabled comparisons between outcome triples.

In regret theory, action A_1 is preferred to action A_2 if, and only if,

$$\sum_{i=1}^n p_i R(x_i, y_i) > 0. \quad (2.106)$$

Regret theory is equivalent to EU in the special case where $R(x_i, y_i) = u(x_i) - u(y_i)$.

Regret aversion leads to fanning-out of the indifference curves in the probability triangle; see Sugden (1986). Hence, it is able to explain the Allais paradox. The early versions of regret theory can also easily account for intransitivity. Some early evidence indicated that the pattern of observed intransitivity was consistent with the predictions of regret theory (Loomes et al., 1989, 1991). However, subsequent evidence on observed patterns of intransitivity has been inconsistent with the predictions of regret theory (Starmer and Sugden, 1998). Furthermore, preferences under regret theory do not require that the decision maker should be indifferent between stochastically equivalent prospects (one of the exercises gives an example of such prospects). And so, regret theory is also able to account for violations of monotonicity. Some of these issues are further developed in exercises.

Bleichrodt et al. (2010) use the trade-off method to devise two separate tests for regret theory. Their two methods give similar results. In each case, they find support for an implication of regret aversion—namely that the regret function is convex. For a recent evaluation of regret theory, its ability to explain several facts that are puzzling for other decision theories, and a review of the relevant empirical evidence, see Bleichrodt and Wakker (2014).

2.10.2 The theory of disappointment aversion

Gul (1991) proposed a model of *disappointment aversion* in which decision makers distinguish between outcomes in a lottery that cause ex-post *elation* and *disappointment*. The reference point relative to which outcomes cause elation and disappointment is the *certainty equivalent* of the lottery. More formally, consider the lottery $L = (x_1, p_1; x_2, p_2; \dots; x_n, p_n)$ such that $x_1 < x_2 < \dots < x_n$. Denote the certainty equivalent of this lottery by C_L . Suppose that

$$x_1 < \dots < x_{k-1} \leq C_L < x_k < \dots < x_n.$$

If the ex-post outcome is $x \leq x_{k-1}$ the decision maker feels disappointed and if it is $x > x_{k-1}$, then the decision maker feels elated. Define $q = \sum_{i=k}^n p_i$ as the sum of elation probabilities and $1 - q$ as the sum of disappointment probabilities. A *disappointment-averse* decision maker then evaluates the lottery L as follows (Gul, 1991, Theorem 1):

$$U(L) = \sum_{j=1}^{k-1} \pi_j^D u(x_j) + \sum_{i=k}^n \pi_i^E u(x_i), \quad (2.107)$$

where u is a utility function, π_j^D is the decision weight accorded to the j th disappointing outcome, and π_i^E to the i th elating outcome. For $\theta > -1$, we define:

$$\pi_i^E = \left(\frac{q}{1 + (1-q)\theta} \right) \frac{p_i}{q} = \frac{p_i}{1 + (1-q)\theta}, k \leq i \leq n. \quad (2.108)$$

$$\pi_j^D = \left(1 - \frac{q}{1 + (1-q)\theta} \right) \frac{p_j}{1-q} = (1+\theta) \frac{p_j}{1 + (1-q)\theta}, 1 \leq j \leq k-1. \quad (2.109)$$

The parameter θ measures the sensitivity of the decision maker towards disappointment. To see this, suppose that an elating outcome, $i \geq k$, and a disappointing outcome, $j \leq k-1$, have identical probabilities $p_i = p_j = p$, then $\pi_j^D = (1+\theta) \pi_i^E$. When $\theta > 0$, the decision weight assigned to a disappointing outcome is magnified relative to the case if it were actually an elating outcome. In the special case of $\theta = 0$, we get the formulation under EU, i.e., $\pi_i^E = p_i$ and $\pi_j^D = p_j$. The greater is θ , the more risk averse the decision maker appears to be to an outside observer. The construction of the model has similarities with prospect theory because the certainty equivalent of the lottery plays the role of an endogenous reference point.⁶⁶

Let us compare the model of disappointment aversion with RDU for the case of binary lotteries, $L_1 = (x_1, 1-p; x_2, p)$, where $x_1 < x_2$. In this case, the certainty equivalent of L_1 lies between x_1 and x_2 . Thus, x_1 is the disappointing outcome and x_2 is the elating outcome. Using (2.107), (2.108), (2.109), and defining $\tilde{w}(p) = \frac{p}{1+(1-p)\theta}$ and $\tilde{w}(1-p) = \frac{1-p}{1+(1-p)\theta}$ we get⁶⁷

$$U(L_1) = (1+\theta) \tilde{w}(1-p)u(x_1) + \tilde{w}(p)u(x_2). \quad (2.110)$$

Define a new probability function $w(p) = \frac{p}{1+(1-p)\theta}$ that is convex and subproportional (see Definition 2.24) for $\theta > 0$. Then (2.110) can be written as

$$U(L_1) = [1 - w(p)]u(x_1) + w(p)u(x_2). \quad (2.111)$$

(2.111) is observationally equivalent to the utility function under RDU. But there are other differences between the two theories. Consider a three-outcome lottery $L_1 = (x_1, p_1; x_2, p_2; x_3, p_3)$, $x_1 < x_2 < x_3$ and $p_1 + p_2 + p_3 = 1$. Fehr-Duda and Epper (2012) show that in a probability triangle with p_1 on the horizontal axis and p_3 on the vertical axis, the indifference curves under disappointment aversion are linear and they fan-out in the bottom half and fan-in in the top half. This contrasts with the indifference curves under RDU which are non-linear. The reason is that under disappointment aversion, the betweenness axiom is satisfied but it does not hold under RDU. Furthermore, consider a concave utility function. Then under RDU, the decision maker can be risk averse or risk seeking, depending on the shape of the probability weighting function. However, under disappointment aversion and $\theta > 0$, the decision maker can only be globally risk averse with a concave utility function; see Fehr-Duda and Epper (2012).

Abdellaoui and Bleichrodt (2007) find that Gul's model of disappointment aversion is too parsimonious and data suggests that disappointment aversion is not constant but variable.⁶⁸ Gill and Prowse (2012) find evidence for the general class of models of disappointment aversion, which they classify to include, Gul's disappointment aversion model, regret aversion, and the Köszegi–Rabin endogenous reference point model. They conduct a novel experiment in which

⁶⁶ See Routledge and Zin (2010) for a two-parameter generalization of the model of disappointment aversion.

⁶⁷ Note that the reference point, the certainty equivalent of the lottery, C_L , is endogenous, and given by $u(C_L) = U(L|C_L)$, so in the general case, there is a fixed point to solve for (though not in the case with just two outcomes).

⁶⁸ Camerer and Ho (1994) and Hey and Orme (1994) also report findings that are negative for models of disappointment aversion.

two players sequentially choose effort levels in a tournament setting. Given their assumptions, the second mover's choice of effort is independent of the first mover's effort choice. Some of the results appear to be sensitive to the specification of the cost of effort function. Empirically, they find that the second mover's choice is influenced by the first mover's choice, which is consistent with models of disappointment aversion.

2.10.3 Saliency and context in decision theory

We now consider an alternative model of decision making under risk that highlights the *saliency* outcomes (Bordalo et al., 2012). Let $S = \{s_1, s_2, \dots, s_n\}$ be the set of possible states. State $s_i \in S$ occurs with probability $p_i \geq 0$, and $\sum_{i=1}^n p_i = 1$. Consider two lotteries or choices $\{L_1, L_2\}$, and let x_s^j be the payoff in state $s \in S$ when lottery L_j , $j = 1, 2$ is chosen. Let $x_s = (x_s^1, x_s^2)$ be the payoffs in the two choices when the state is s ; $x_s^{\min}$ and $x_s^{\max}$ give the minimum and maximum payoffs in state s . Sometimes we shall denote x_s^{-i} as the payoff of lottery L_j in state s .

In most alternatives to EU, whether a particular outcome is overweighted (more salient) or underweighted (less salient) depends on the probability associated with the outcomes. For instance, in the Prelec (1998) probability weighting function, low probabilities are overweighted and high probabilities are underweighted. However, Bordalo et al. (2012) propose that the notion of saliency is a feature of the relative monetary payoffs of the two lotteries in each state.

Definition 2.50 *The saliency of state s for lottery L_i , $i = 1, 2$, is given by a continuous and bounded saliency function, $\sigma(x_s^i, x_s^{-i})$, that has the following properties.*

- (a) *Ordering: Let $s, \tilde{s} \in S$. If $[x_s^{\min}, x_s^{\max}] \subset [x_{\tilde{s}}^{\min}, x_{\tilde{s}}^{\max}]$, then $\sigma(x_s^i, x_s^{-i}) < \sigma(x_{\tilde{s}}^i, x_{\tilde{s}}^{-i})$.*
- (b) *Diminishing sensitivity: Suppose that in some state s , the outcomes in both lotteries are strictly positive, i.e., $x_s^i > 0$, $i = 1, 2$. Then, for any $\epsilon > 0$,*

$$\sigma(x_s^i + \epsilon, x_s^{-i} + \epsilon) < \sigma(x_s^i, x_s^{-i}).$$

- (c) *Reflection: Let $s, \tilde{s} \in S$ and $x_s^j > 0$, $x_{\tilde{s}}^j > 0$, then*

$$\sigma(x_s^i, x_s^{-i}) < \sigma(x_{\tilde{s}}^i, x_{\tilde{s}}^{-i}) \Leftrightarrow \sigma(-x_s^i, -x_s^{-i}) < \sigma(-x_{\tilde{s}}^i, -x_{\tilde{s}}^{-i}).$$

Ordering implies that a state is more salient if the distance between the two payoffs increases. Diminishing sensitivity implies that a state is less salient if, preserving the distance between the two outcomes, the outcomes are further away from zero. So zero plays the role of a reference point in this analysis. Reflection implies that saliency is preserved under sign changes of outcomes. The last two properties also capture the notion of *context* in the model.

A decision maker in neoclassical theory, such as EU, treats all states as equally salient (in the sense defined above). However, decision makers may assign higher probability weights to more salient states relative to less salient ones. Unlike decision makers in neoclassical theory, following Gennaioli and Shleifer (2010), we may term such decision makers as *local thinkers*.

Before we see how local thinkers make decisions, consider an example of a saliency function that satisfies all the conditions in Definition 2.50.

$$\sigma(x_s^i, x_s^{-i}) = \frac{|x_s^i - x_s^{-i}|}{|x_s^i| + |x_s^{-i}| + \theta}; \theta > 0. \quad (2.112)$$

In (2.112), ordering is captured by the term in the numerator, and diminishing sensitivity by the denominator. Ordering and diminishing sensitivity are the two most important properties used in this model. Yet another property, symmetry, is intuitive when there are only two lotteries. Symmetry implies that $\sigma(x_s^i, x_s^{-i}) = \sigma(x_s^{-i}, x_s^i)$ and is satisfied by (2.112).

Example 2.23 (*Saliency of states in the Allais paradox*): Consider the Kahneman and Tversky (1979) version of the Allais paradox in its common consequence form that we have considered above. The decision maker chooses between the two lotteries

$$L_1(z) = (0, 0.01; z, 0.66; 2500, 0.33) \text{ and } L_2(z) = (z, 1), z > 0.$$

When $z = 2400$, most decision makers choose the safe choice, $L_2(2400)$, and when $z = 0$ decision makers make the risk seeking choice, $L_1(0)$. This violates EU. Let $z = 2400$, denote the three states by s_1, s_2, s_3 , and rewrite $L_2(2400) = (2400, 0.01; 2400, 0.66; 2400, 0.33)$. Then, the three state-dependent outcome tuples in the two lotteries are $x_1 = (0, 2400)$, $x_2 = (2400, 2400)$, $x_3 = (2500, 2400)$ with associated probabilities $p_1 = 0.01$, $p_2 = 0.66$, $p_3 = 0.33$. We now consider the relative saliency of these states. Using diminishing sensitivity,

$$\sigma(0 + 100, 2400 + 100) < \sigma(0, 2400). \quad (2.113)$$

Using the ordering property,

$$\sigma(2400, 2500) < \sigma(100, 2500). \quad (2.114)$$

From (2.113), (2.114),

$$\sigma(2400, 2500) < \sigma(0, 2400). \quad (2.115)$$

Also, by the ordering property,

$$\sigma(2400, 2400) < \sigma(2400, 2500). \quad (2.116)$$

From (2.115), (2.116), the relative saliency of the three states is given by

$$\sigma(2400, 2400) < \sigma(2400, 2500) < \sigma(0, 2400). \quad (2.117)$$

If we wish to invoke symmetry, then we can also write the middle term as $\sigma(2500, 2400)$. We leave it to the reader to describe the saliency of the various states when $z = 0$.

Definition 2.51 Let $s, \tilde{s} \in S$. Then state s is relatively more salient for lottery L_i if $\sigma(x_s^i, x_s^{-i}) > \sigma(x_{\tilde{s}}^i, x_{\tilde{s}}^{-i})$. Using this method we can rank the n states in the order of saliency for lottery L_i . Let this ranking be $k_s^i \in \{1, 2, \dots, n\}$ where a lower number denotes higher saliency. Then, given the saliency rankings of states, for lottery L_i , a “local thinker” transforms the relative probabilities of states s and $\tilde{s}$, $\frac{p_s}{p_{\tilde{s}}}$, as follows

$$\frac{p_s^i}{p_{\tilde{s}}^i} = \delta^{k_s^i - k_{\tilde{s}}^i} \left(\frac{p_s}{p_{\tilde{s}}} \right), \delta \in (0, 1], \quad (2.118)$$

where the ratio of transformed probabilities is denoted by $\frac{p_s^i}{p_{\tilde{s}}^i}$. We normalize $\sum_{s \in S} p_s^i = 1$. The decision weight associated with the state s for lottery L_i is given by

$$p_s^i = p_s \omega_s^i, \text{ where } \omega_s^i = \frac{\delta^{k_s^i}}{\sum_{r \in S} \delta^{k_r^i} p_r}. \quad (2.119)$$

By definition,

$$\begin{aligned} \sum_{s \in S} p_s^i &= \sum_{s \in S} \omega_s^i p_s = \sum_{s \in S} p_s \left(\frac{\delta^{k_s^i}}{\sum_{r \in S} \delta^{k_r^i} p_r} \right) \\ &= \frac{1}{\sum_{r \in S} \delta^{k_r^i} p_r} \sum_{s \in S} p_s \delta^{k_s^i} \\ &= 1. \end{aligned}$$

Thus, the average distortion in probabilities must be zero (because $\sum_{s \in S} p_s = 1$). From (2.118), if state s is relatively more salient than state $\tilde{s}$, then $k_s^i < k_{\tilde{s}}^i$, so $\delta^{k_s^i - k_{\tilde{s}}^i} > 1$, thus, the transformed probabilities of state s are also relatively higher. Conversely, if s is relatively less salient, then the transformed probability of that state is relatively lower. The size of δ is a measure of how much the probabilities are transformed. If $\delta = 1$, then $\delta^{k_s^i - k_{\tilde{s}}^i} = 1$, so there is no distortion of probabilities and we have the case of EU; by contrast when $\delta < 1$, the decision maker is a local thinker. From (2.119), if $\delta^{k_s^i}$ is higher than the average $\sum_{r \in S} \delta^{k_r^i} p_r$, i.e., $\omega_s^i > 1$, then state s receives a higher decision weight; in other words it is relatively more salient. Thus, unlike some other behavioral theories, e.g., RDU, OPT, PT, the distortion in probabilities does not depend on the size of the probabilities, but rather on the salience of the outcomes in different states. Thus, low probability states are only overweighted if they are salient, otherwise not.

Definition 2.52 The objective function of a “local thinker” for the lottery L^i , denoted by $V^{LT}(L^i)$ is given by

$$V^{LT}(L^i) = \sum_{s \in S} p_s^i u(x_s),$$

where $u(x_s)$ is the utility of outcome x_s .

Let us revisit Example 2.23. Empirical evidence shows that most decision makers prefer $L_2(2400)$ to $L_1(2400)$. The explanation based on the framework of a local thinker is along the following lines. From (2.117), the most salient state for lottery $L_1(2400)$ is the one in which the outcome is 0, despite the fact that the outcome 0 occurs with a very small probability of 0.01. This leads to a substantial overweighting of this outcome (as is the case under RDU/OPT/PT, but for an entirely different reason). Thus, the decision maker prefers the lottery $L_2(2400)$ to $L_1(2400)$.

The model can explain the Allais paradox, preference reversal, and certain context dependent choices. Bordalo et al. (2013a) apply this model to explain four puzzles in finance. For instance, they can explain why the prices of assets with skewed returns are so high; the reason is that skewed returns are relatively more salient. For an alternative account, based on inverse S-shaped probability weighting functions, under RDU/PT, see Section 3.9. They can also explain the equity premium puzzle because downside risk is more salient than upside risk, an implication that does the job of loss aversion under PT. Bordalo et al. (2013b) apply the theory to consumer choice. They consider various attributes of consumer goods such as quality, price, and context. They then determine the consumer’s willingness to pay for the good, depending on which attributes are more salient.

There is strong empirical support for the separate components of prospect theory that is also able to explain most of the phenomena that this theory of salience can explain. Hence, what is needed are empirical tests that stringently pit the predictions of the two theories against each other in domains where these predictions differ.

2.10.4 Case-based decision theory

One criticism of the Bayesian approach is that it is agnostic about the source of beliefs. One possible response to this criticism is the *case base decision theory* approach developed in Gilboa and Schmeidler (1995, 1996, 1997, 2001).⁶⁹ A central construct in this theory is a *similarity function*. When evaluating a new choice, individuals use a similarity function to compare it with all other cases in their database. The degree of similarity with other cases as well as the historical outcomes associated with choices made in the other cases then guide individuals in forming beliefs about their new choice.

The essence of the idea can be explained as follows. Suppose that an individual has some action set $A = \{a_1, a_2, \dots, a_n\}$. There is also some universal set of *cases*, or possible situations, $C = \{c_1, c_2, \dots, c_m\}$. Suppose that, at some instant in time, an individual has experience with the subset of cases $\bar{C} = \{c_1, c_2, \dots, c_n\}$, $n \leq m$. Denote the utility arising from taking action a_i (in the past) in case c_j by u_{ij} . Then, for all i, j , the memory of the individual is given by $\Omega = \{A, \bar{C}, C, \{u_{ij}\}\}$. Suppose now that some new case arises, say, $c_k \in C$, $c_k \notin \bar{C}$. What action should the decision maker take?

Define an exogenously given similarity function $s(c_j, c_k)$, $c_j \in \bar{C}$. This captures the similarity between the new case and any case $c_j \in \bar{C}$. Suppose that the decision maker wishes to evaluate the utility from taking action a_i under the new case, c_k . Let M_i be the subset of cases in which action a_i was taken in the past. Then the utility from taking action a_i in case c_k , given M_i , is denoted by $U(a_i | c_k, M_i)$; this is sometimes called a *similarity weighted utility*:

$$U_i = U(a_i | c_k, M_i) = \sum_j s(c_j, c_k) u_{ij}; \quad c_j \in M_i. \quad (2.120)$$

In (2.120), all cases in which action a_i was not taken in the past are irrelevant for the calculation and assigned zero utility, which is taken to be the *default aspiration level* of the decision maker. The chosen action a^* then satisfies

$$a^* \in \arg \max_{a_i \in A} U_i. \quad (2.121)$$

One implication of this framework is that if an action was never taken in the past, then it might never be taken in the future. However, because the aspiration utility level is zero, if all actions taken in the past yield negative utility, then the decision maker randomly chooses an action from among the set of actions not tried in the past. In this case, an action that was never chosen in the past can get chosen.

Gilboa and Schmeidler (1997) ask: What does the behavior of a case-based decision maker converge to in the long run? To answer this question, they postulate dynamics for the aspiration level as well as a specific form of the similarity function. In particular, this similarity function rules out the possibility that the outcome of an act reveals any information that is relevant for the

⁶⁹ For axiomatizations, see Billot et al. (2005) and Gilboa et al. (2006).

evaluation of another act.⁷⁰ Based on these assumptions they find that the long-run behavior of the decision maker is optimal in the sense that it converges to one that would be chosen by an expected utility maximizer.

Guerdjikova (2008) shows that the optimality result fails if the specific similarity function used by Gilboa and Schmeidler (1997) is altered. She goes on to show how a modification of the similarity function and the dynamics of the aspiration levels can ensure the achievement of optimality. Extensive empirical evidence is required to sharpen and test these models. Using agent-based computational methods, Pape and Kurtz (2013) find evidence in support of case-based decision theory. Bleichrodt et al. (2013) also report evidence that is supportive of the theory.

Two further issues that need to be addressed are as follows. First, often individuals construct *hypothetical cases* and case studies in their minds before making current choices, yet such cases are not a part of the existing theory. These hypothetical actions might be completely unrelated to past actions or payoffs, yet, be crucial in making current choices. Indeed, the ability to construct hypothetical cases is often cited as an ability that is unique to humans. Second, convergence of the model to an expected utility framework needs to deal with the issue of the empirical rejection of the expected utility model.

2.10.5 Configural weights models

We now briefly outline some *configural weight models* that have been proposed to deal with the violations of PT outlined in this section, followed by an assessment.⁷¹ Like other non-EU alternatives, these models also weight probabilities in a non-linear manner. In these models, one first writes down the decision tree corresponding to the problem.

Definition 2.53 (*Lotteries and ranks*): Consider some lottery, $L = (x_1, p_1; x_2, p_2; \dots; x_n, p_n)$ where $x_1 \geq x_2 \geq \dots \geq x_n$.⁷² Outcome x_j is then said to have rank j , $j = 1, 2, \dots, n$.

THE RANK AFFECTED MULTIPLICATIVE WEIGHTS MODEL (RAM)

In the *rank affected multiplicative weights model* (RAM), the utility from lottery, L , is:⁷³

$$RAM(L) = \frac{\sum_{i=1}^n a(i, n, s) t(p_i) u(x_i)}{\sum_{j=1}^n a(j, n, s) t(p_j)}. \quad (2.122)$$

Indices i, j denote the *rank* of the outcome (see Definition 2.53). s is the *augmented sign* of the branch's consequence (positive, zero, or negative). The weighting function under RAM, $t(p)$, is an increasing function of p that takes the form,

$$t(p) = p^\nu; 0 < \nu < 1. \quad (2.123)$$

⁷⁰ One can immediately think of many examples to convince oneself that this is quite a strong assumption. For instance, a car accident while being drunk may convince an individual to avoid drinking while operating other forms of machinery.

⁷¹ The interested reader can consult Birnbaum (2008) for an exhaustive set of references and a detailed survey of these models.

⁷² Note that configural weights models write the successive outcomes in a lottery (reading from left to right) in decreasing order. This convention is used only in this section in the book.

⁷³ The reader can consult Birnbaum (1997) for the details.

The utility function $u(x)$ typically takes the power form,

$$u(x) = x^\beta; 0 < \beta < 1. \quad (2.124)$$

$a(i, n, s)$ is the *rank and sign augmented branch weight* corresponding to the outcome that has rank i . The sign s of an outcome can be positive, negative, or zero. For two branch gambles, e.g., the lottery $L = (100, 0.5; 0, 0.5)$, the branch corresponding to the lower outcome (\$0) is associated with $a = 2$ and the higher outcome is associated with $a = 1$. For m branch gambles, such that $m \geq 3$, the weights (the a s) are simply the ranks of the outcomes for that branch. For example, for the set of outcomes, $x_m \leq x_{m-1} \leq \dots \leq x_1$ the weights are, respectively, $m, m-1, m-2, \dots, 1$. Thus, lower outcomes are given higher weights, which is one of the main features of the model. The presence of the sign, s , of the outcome in $a(i, n, s)$ allows one to alter the branch weight, depending on the sign. Finally, using (2.123), we get

$$\lim_{p \rightarrow 0} \frac{t(p)}{p^\nu} = 1. \quad (2.125)$$

TRANSFER OF ATTENTION IN EXCHANGE MODELS (TAX)

Transfer of attention in exchange models (TAX) assume that the decision maker has a fixed and limited amount of attention. Hence, it is assumed, as in RAM models, that the decision maker transfers weight from higher ranked to lower ranked outcomes. Consider a *three-branch* lottery, L , given by $L = (x_1, p_1; x_2, p_2; x_3, p_3)$, $x_1 \geq x_2 \geq x_3$. Then, the utility of lottery L to the decision maker under TAX is given by

$$TAX(L) = \frac{w_1 u(x_1) + w_2 u(x_2) + w_3 u(x_3)}{w_1 + w_2 + w_3}, \quad (2.126)$$

where,

$$w_1 = t(p_1) - \frac{2}{4} \delta t(p_1); w_2 = t(p_2) - \frac{1}{4} \delta t(p_2) + \frac{1}{4} \delta t(p_1); w_3 = t(p_3) + \frac{1}{4} \delta t(p_1) + \frac{1}{4} \delta t(p_2); \quad (2.127)$$

$t(p)$ is given in (2.123) and $\delta > 0$ is some constant. In this version of the TAX model, also known as the *special TAX model*, due to Birnbaum and Stegner (1979), one transfers a fixed amount of probability weight from a higher outcome, proportionately, to lower outcomes. Thus, $w_1 + w_2 + w_3 = t(p_1) + t(p_2) + t(p_3)$. From w_1 given in (2.127), it follows that the weight on the highest outcome, x_1 , is reduced by $\frac{2}{4} \delta t(p_1)$, and each of the lower two outcomes gains one half of this weight. Similarly, the middle ranked outcome, x_2 , sheds weight $\frac{1}{4} \delta t(p_2)$, which is gained by the worse outcome, x_3 .

THE GAINS DECOMPOSITION UTILITY MODEL (GDU)

In the GDU model, multi-branch lotteries are reduced to simple, two-branch, lotteries.⁷⁴ For instance, consider the binary lottery $L_2 = (x_1, p; x_2, 1-p)$ such that $x_1 \geq x_2$. Then, under the *gains decomposition utility model* (GDU), the utility of this binary lottery is:

⁷⁴ For the GDU model, see Luce (2000), and Marley and Luce (2001, 2005).

$$GDU(L_2) = w(p)u(x_1) + (1 - w(p))u(x_2). \quad (2.128)$$

Notice that the decision weights are formed exactly as under RDU. Three-branch lotteries such as $L_3 = (x_1, p_1; x_2, p_2; x_3, 1 - p_1 - p_2)$, $x_1 \geq x_2 \geq x_3$, are then evaluated as follows. First, define the lottery

$$L_4 = \left(x_1, \frac{p_1}{p_1 + p_2}; x_2, 1 - \frac{p_1}{p_1 + p_2} \right),$$

which offers a binary choice between the two best outcomes with the probabilities normalized to add up to one. Next, consider the lottery

$$L_5 = (L_4, p_1 + p_2; x_3, 1 - p_1 - p_2),$$

which offers a binary choice between the worst outcome x_3 in lottery L_3 and the lottery L_4 . The worth of the lottery L_3 to the decision maker under the GDU model is

$$GDU(L_3) = w(p_1 + p_2)GDU(L_4) + [1 - w(p_1 + p_2)]u(x_3).$$

The weighting functions used in this theory are the standard weighting functions. Thus, $\lim_{p \rightarrow 0} \frac{w(p)}{p^\gamma} = \infty$ (see Definition 2.41).

AN EVALUATION OF THE CONFIGURAL WEIGHTS MODELS

Section 2.9.1 described several anomalies of EU/RDU/OPT/PT. These included coalescing, stochastic dominance, pattern of RBI, and upper tail independence. The configural weights models are able to explain these anomalies (Birnbbaum, 2008). So what are economists to make of these models, which seem to be the only ones that can explain the relevant evidence?

1. The main difference in predictions between PT and the configural weights models arises in the case of *event splitting* (see Section 2.9.1); see Birnbbaum (2008). This occurs when there are two identical outcomes in a lottery. This does not apply when all outcomes in a lottery are distinct, as for instance is the case with most of our applications (see Chapter 3).
2. For lotteries of the form $(x, p; y, 1 - p)$, where $x < 0 < y$, Birnbbaum (2008) reports that PT and TAX give close results. In many cases in economics, one is interested in binary decisions, often associated with binary outcomes, such as to pay or not pay taxes, to insure or not insure, to quit or not to quit one's job. So, for TAX to offer different predictions, we need lotteries with at least three outcomes.
3. More generally, it is difficult to distinguish between PT and TAX within the probability triangle. Hence, where PT is successful (e.g., the experiments reported by Kahneman and Tversky (1979), Tversky and Kahneman (1992) and, more generally, events that can be represented in the probability triangle), TAX is also successful. But for the "new" paradoxes, TAX succeeds where PT fails.
4. In many applications that involve use of hypothesis H2 (i.e., behavior for near-zero probabilities), configural weights models fail (see, e.g., (2.125)) as do EU, RDU, and PT, while *composite prospect theory* (CPT) succeeds in these cases (see Section 2.11). In these cases, the configural weights models can benefit by incorporating hypothesis H2 and the insights generated by composite prospect theory (described in Section 2.11).

5. To many economists, the practice of transferring an arbitrary amount of weight from a higher outcome to a lower outcome (see, e.g., (2.127)) will appear unattractive and mechanical. Proponents of configural weights models will argue that the weight to be transferred is calculated to fit the existing data; see Birnbaum (2008). But this begs the question of uniqueness; is there a unique way to transfer weight from the upper to the lower branches that leads to a successful prediction? This is the sense in which these models may appear arbitrary to some. Although this is not a fatal flaw, it is also not clear what restrictions on preferences will lead to the RAM and TAX models or even if these restrictions are likely to be reasonable.

Overall, economists are likely to be persuaded more by the evidence presented by the proponents of the configural weights models and less by the configural weights models that have been proposed to explain the evidence. Explaining the evidence gathered in this literature remains a challenge for behavioral decision theory.

2.11 Human behavior for extreme probability events

Many behavioral economists believe that for situations of risk, uncertainty, and ambiguity, RDU, OPT, and PT are the most satisfactory alternatives to EU (Wakker, 2010). Recall hypotheses H1 and H2 from Section 2.2. RDU and PT only take account of hypothesis H1, however, no decision theory accounts for the more general hypothesis H2 in a satisfactory manner. al-Nowaihi and Dhami (2010a) propose a modification of PT, *composite prospect theory* (CPT), that is able to account for H2. They also show how RDU can be modified to produce *composite rank dependent utility theory* (CRDU) that can also explain H2, but it is less satisfactory as compared to CPT. For the benefit of the readers we reproduce H2 here:

H2: There are two kinds of decision makers:

(H2a) A fraction $\mu \in [0, 1]$ of decision makers (i) ignore events of extremely low probability and, (ii) treat extremely high probability events as certain.⁷⁵ For all other (non-extreme) probabilities they behave as in H1, i.e., they overweight small probabilities and underweight large ones.

(H2b) A fraction $1 - \mu$ of the decision makers behave as in H1 for all levels of probabilities.

We may also view H2 as a *bimodal perception of risk*. Namely, that some individuals ignore very low probabilities altogether. Others do not ignore very low probabilities, particularly in the domain of losses; see for instance, McClelland et al. (1993) in an insurance context. Similar arguments can be constructed from the work of Viscusi (1998). Kunreuther et al. (1978) argue that the bimodal response to low probability events is pervasive in field studies. In their OPT, Kahneman and Tversky (1979) were acutely aware of H2.

The bimodal perception of risks framework and OPT (see Section 2.11.1 below) do not specify the size of μ . It is likely that μ depends on the context and the type of problem and that it can be influenced by the media, family, friends, and public policy. The size of μ is also likely to be influenced by emotions, experience, time available to make a decision, bounded rationality, framing,

⁷⁵ In the context of the take-up of insurance for low probability natural hazards, the results from one set of experiments by Kunreuther et al. (1978) are consistent with $\mu = 0.8$.

incentive effects, and so on.⁷⁶ For the purposes of our exposition we shall simply take μ as given. The endogeneity of μ is likely to be a fruitful avenue for future work.

2.11.1 Original prospect theory (OPT)

The only decision theory that incorporates hypothesis H2 is Kahneman and Tversky's (1979) *original prospect theory* (OPT). OPT makes a distinction between an *editing* and an *evaluation/decision* phase. From our perspective, the most important aspect of OPT is that decision makers decide *which improbable events to treat as impossible and which probable events to treat as certain*. Kahneman and Tversky (1979) employ an objective function similar to (2.24) in Definition 2.34, with one main difference: they use point transformations of probabilities, $\pi_i = w(p_i)$. From Example 2.7, we know that decision makers who use such a transformation can choose stochastically dominated options (Fishburn, 1978). Based on $\pi_i = w(p_i)$, and taking account of H2, they drew their decision weights function (or probability weighting function), $\pi(p)$, as in Figure 2.16; for our purposes the exact shape is less important than the gaps at the endpoints.

Kahneman and Tversky (1979, pp. 282–3) summarize the evidence for H2, as follows.

The sharp drops or apparent discontinuities of $\pi(p)$ at the end-points are consistent with the notion that there is a limit to how small a decision weight can be attached to an event, if it is given any weight at all. A similar quantum of doubt could impose an upper limit on any decision weight that is less than unity . . . the simplification of prospects can lead the decision maker to discard events of extremely low probability and to treat events of extremely high probability as if they were certain. Because people are limited in their ability to comprehend and evaluate extreme probabilities, highly unlikely events are either ignored or overweighted, and the difference between high probability and certainty is either neglected or exaggerated. Consequently $\pi(p)$ is not well-behaved near the end-points.




Figure 2.16 A hypothetical probability weighting function.

Source: Kahneman and Tversky (1979, p. 282) with permission of The Econometric Society.

⁷⁶ See chapter 4 (Anomalies on the demand side) in Kunreuther and Pauly (2005).

In this quote, Kahneman and Tversky are very explicit about human behavior near $p = 0$. Such events, in their words, are either *ignored* by some fraction μ of individuals (H2a) or *overweighted* by the complementary fraction $1 - \mu$ (H2b). This is what creates the gaps at the end points in Figure 2.16. However, the informal and heuristic approach of OPT is not always practical in dealing with H2.

We briefly describe some of the other simplifications or heuristics undertaken in the editing phase, although these do not necessarily deal with extreme probabilities. *Coalescing* or *combination* is the property that if an outcome appears more than once in a lottery, its probabilities can be added. For instance, the lottery $(30, 0.2; 30, 0.3; 75, 0.5)$, also known as a *split lottery*, is equivalent to the lottery written in the coalesced form as $(30, 0.5; 75, 0.5)$. The *cancellation heuristic* is used to eliminate common components among the lotteries being compared. The *dominance heuristic* is used to recognize and eliminate lotteries that are stochastically dominated.

In OPT, following the *editing phase*, the decision maker is left with a set of lotteries, which are converted into incremental form (Definition 2.30). In the *evaluation phase*, the decision maker then proceeds to evaluate these lotteries as in Definition 2.34 above.

OPT was heavily criticized for the feature that decision makers could choose stochastically dominated options. In response to these difficulties, Tversky and Kahneman (1992) modified OPT to *cumulative prospect theory* (PT) by incorporating the insight of *cumulative transformation of probabilities* from RDU. Hence, they needed a probability weighting function, $w(p)$, that, unlike the one in Figure 2.16, was *defined* for the entire probability domain $[0, 1]$ (otherwise Definition 2.33 could not be applied). Modifying OPT to PT meant eliminating the psychologically rich *editing phase*, which may have been a difficult choice for its authors. However, since PT lacks the relevant editing phase, it cannot account for H2a.

2.11.2 The evidence on behavior under low probability events

In this section we present a brief overview of the evidence for hypothesis H2. The interested reader can consult al-Nowaihi and Dhami (2010a) for the details. In all the low probability anomalies in this section, the most natural explanation is that a fraction of the individuals behave as in H2a. Thus, they ignore the very low probability of the potentially high levels of harm that might be caused by their actions (actions with low levels of harm are not as interesting). We note that under EU, but particularly under RDU and PT, decision makers place too great a salience on low probability events because for most probability weighting functions, $\lim_{p \rightarrow 0} \frac{w(p)}{p} = \infty$ (see Proposition 2.12).

INSURANCE FOR LOW PROBABILITY EVENTS

Kunreuther et al. (1978) provides striking evidence that decision makers buy inadequate non-mandatory insurance against low probability natural hazards, e.g., earthquakes, floods, and hurricane damage in areas prone to these hazards.⁷⁷ EU predicts that a risk averse decision maker facing an actuarially fair premium will, in the absence of transactions costs, buy full insurance for all probabilities, however small. Kunreuther et al. (1978, chapter 7) presented subjects with varying potential losses with various probabilities, and fixed expected loss. Subjects faced actuarially

⁷⁷ In the foreword, Arrow (Kunreuther et al., 1978, p. vii) writes: "The following study is path breaking in opening up a new field of inquiry, the large scale field study of risk-taking behavior." The reader can also consult similar but more recent evidence in Kunreuther and Pauly (2004, 2005).

fair, unfair, or subsidized premiums. In each case, they found that as the probability of the loss decreases, there is a point below which the take-up of insurance robustly drops dramatically.⁷⁸

The lack of interest in buying insurance for low probability natural hazards arose despite active government attempts to (i) provide subsidy to overcome transaction costs, (ii) reduce premiums below their actuarially fair rates, (iii) provide reinsurance for firms, and (iv) provide the relevant information. Furthermore, insurees were aware of the losses due to natural hazards (many overestimated them). Moral hazard issues (e.g., expectation of federal aid in the event of disaster) were not found to be important.

Arrow's own reading of the evidence in Kunreuther et al. (1978) is that the problem is on the demand side rather than on the supply side. Arrow writes (Kunreuther et al., 1978, p. viii) "Clearly, a good part of the obstacle [to buying insurance] was the lack of interest on the part of purchasers." Kunreuther et al. (1978, p. 238) write: "Based on these results, we hypothesize that most homeowners in hazard-prone areas have not even considered how they would recover should they suffer flood or earthquake damage. Rather they treat such events as having a probability of occurrence sufficiently low to permit them to ignore the consequences." This behavior is in close conformity to H2a.

THE BECKER (1968) PARADOX

A celebrated result, the *Becker (1968) proposition*, states that the most efficient way to deter crime is to impose the "severest possible penalty, F , with the lowest possible probability of detection and conviction, p ." By reducing p , society can economize on the costs of enforcement, such as policing and trial costs. But by increasing F , which is relatively less costly, the deterrence effect of the punishment is maintained. In an appropriately specified model, one can prove the following proposition

Proposition 2.22 (*Dhami and al-Nowaihi, 2013*): *Under EU, risk-neutrality or risk-aversion, potentially infinitely severe punishments,⁷⁹ and a utility that is unbounded below, the Becker proposition implies that crime would be deterred completely, however small the probability of detection and conviction.*

Kolm (1973) memorably phrased this proposition as: *it is efficient to hang offenders with probability zero*. Empirical evidence strongly suggests that the behavior of a significant number of people does not conform to the Becker proposition. This is known as the *Becker paradox*. For instance, Levitt (2004) shows that the estimated contribution of capital punishment in deterring crime in the US, over the period 1973–1991, was zero.⁸⁰

Under RDU and PT, because decision makers heavily overweight the small probability of a punishment (see Proposition 2.12), they can be shown to be deterred by Becker type punishments even more as compared to EU and so the Becker paradox remains; see Dhami and al-Nowaihi (2013) for a formal proof. There are several competing explanations for the Becker paradox but none of them are ultimately satisfactory; see al-Nowaihi and Dhami (2010a) for the details.

A natural explanation for the Becker paradox is that a fraction μ of individuals ignore (or heavily underweight) the very low probability of getting caught (as in H2a), so they commit a crime despite high levels of punishment.

⁷⁸ These results were shown to be robust to changes in subject population, changes in experimental format, order of presentation, presenting the risks separately or simultaneously, bundling the risks, compounding over time, and introducing "no claims bonuses."

⁷⁹ For instance, ruinous fines, slavery, torture, extraction of body parts (all of which have been historically important), and modern capital punishment.

⁸⁰ See also Radelet and Ackers (1996) and Polinsky and Shavell (2007).

EVIDENCE FROM JUMPING RED TRAFFIC LIGHTS

Consider the act of running a red traffic light. There is at least a *small probability* of an accident. al-Nowaihi and Dhami (2010a) show that the nature of this problem is very similar to the Becker proposition, except that the consequences are *self-inflicted* and potentially have *infinite costs*, such as loss of one's life. Indeed, under EU, we should not expect anyone to run red traffic lights. Decision makers using RDU and PT should be even more reluctant to run red traffic lights because of the feature $\lim_{p \rightarrow 0} \frac{w(p)}{p} = \infty$ (see Proposition 2.12) in most probability weighting functions.

Bar-Ilan and Sacerdote (2001, 2004) estimate that there are 260,000 accidents per year in the US caused by red-light running with implied costs of car repair alone of the order of \$520 million per year. Clearly, this is an activity of economic significance and it is implausible to assume that running red traffic lights are simply "mistakes." Using Israeli data, Bar-Ilan (2000) calculated that the expected gain from jumping one red traffic light is, at most, one minute (the length of a typical light cycle). Given the known probabilities, they find that if a slight injury causes a loss greater or equal to 0.9 days, a risk-neutral person who follows EU will be deterred by that risk alone. However, the corresponding numbers for the additional risks of serious and fatal injuries are 13.9 days and 69.4 days respectively, which should ensure complete compliance with the law. However, evidence is to the contrary.

Clearly, EU combined with risk aversion struggles to explain this evidence. So too do RDU and PT. A plausible explanation for observed red traffic light running is that a fraction μ of individuals run red traffic lights because they ignore or seriously underweight the very low probability of an accident (as under H2a).

OTHER EXAMPLES OF HYPOTHESIS H2

Consider the usage of hand-held mobile phones in moving vehicles. A user of mobile phones faces potentially *infinite punishment* (e.g., loss of one's and/or the family's life) with *low probability*, in the event of an accident.⁸¹ Evidence from the UK indicates that up to 40% of individuals drive and talk on mobile phones (The Royal Society for the Prevention of Accidents, 2005). Pöystia et al. (2005) report that two thirds of Finnish drivers and 85% of American drivers use their phone while driving, which increases the risk of an accident by two to six fold. Hands-free equipment, although now obligatory in many countries, seems not to offer essential safety advantages.

People were initially reluctant to use non-mandatory seat belts in cars despite publicly available evidence that seat belts save lives. Prior to 1985, in the US, only 10–20% of motorists wore seat belts voluntarily (Williams and Lund, 1986). Car accidents may be perceived by individuals to be *very low probability events*, particularly if individuals are overconfident of their driving abilities.⁸² Although accidents are *potentially fatal*, people seem to ignore the low probability of facing such events.

Even as evidence accumulated on the dangers of breast cancer (*low probability, potentially fatal, event*⁸³) women only sparingly took up the offer of breast cancer examination. In the US, this changed only after the greatly publicized events of the mastectomies of Betty Ford and Happy

⁸¹ Extensive evidence suggests that the perceived probability of an accident might be even lower than the actual probability because drivers are *overconfident* of their driving abilities. Taylor and Brown (1988) suggest that upto 90% of car accidents might be caused by overconfidence.

⁸² People assign confidence intervals to their estimates that are too narrow and 90% of those surveyed report that they have above-average levels of intelligence and emotional ability (Weinstein, 1980).

⁸³ We now know that the *conditional* probability of breast cancer if there is such a problem in close relatives is not low. However, we refer here to data from a time when such a link was less well understood.

Rockefeller; see Kunreuther et al. (1978, p. xiii, 13–14). Hence, the fraction of individuals to whom stylized fact H2(a) applies, μ , is not static.

CONCLUSION FROM THESE DISPARATE CONTEXTS

Two main conclusions arise from these examples. First, human behavior for low probability events cannot be easily explained by most decision theories. EU and the associated auxiliary assumptions are unable to explain the stylized facts. RDU and PT make the problem even worse because $\lim_{p \rightarrow 0} \frac{w(p)}{p} = \infty$ (see Proposition 2.12). Second, a natural explanation for these phenomena is that many individuals (the fraction, μ) simply ignore or seriously underweight very low probability events, even when the consequence is a huge loss.

2.11.3 Composite Prospect Theory

The theoretical solution offered by al-Nowaihi and Dhami (2010a) to incorporate hypothesis H2 is a proposal for a new probability weighting function—the *composite Prelec probability weighting function* (CPF). The CPF is generated by *smooth pasting* three standard Prelec functions; see Figure 2.17 for the general shape of a CPF. In Figure 2.17, decision makers heavily underweight very low probabilities in the range $[0, p_1]$ in the sense that is made clear in Remark 2.8.

Remark 2.8 *The CPF zero-underweights infinitesimal probabilities, i.e., $\lim_{p \rightarrow 0} \frac{w(p)}{p} = 0$, and zero-overweights near-one probabilities, i.e., $\lim_{p \rightarrow 1} \frac{1-w(p)}{1-p} = 0$. This contrasts with the standard Prelec function where $\lim_{p \rightarrow 0} \frac{w(p)}{p} = \infty$ and $\lim_{p \rightarrow 1} \frac{1-w(p)}{1-p} = \infty$ (see Proposition 2.12).*

Decision makers who use the CPF ignore very low probability events by assigning low subjective probability weight, w , to them, hence, conforming with H2a. By contrast, decision makers who use the standard Prelec function heavily overweight very low probabilities (see Proposition 2.12), and so they conform to H2b. Events in the interval $[p_3, 1]$ are overweighted in the CPF, as suggested by Kahneman and Tversky (1979, pp. 282–3). In the middle segment, $p \in [\underline{p}, \bar{p}]$, the CPF is inverse S-shaped, identical to a standard Prelec function. al-Nowaihi and Dhami (2010a)




Figure 2.17 The composite Prelec weighting function (CPF).

show formally that CPF is *flexible* and has preference foundations. The standard Prelec function is a special case of the CPF.

A BRIEF, MORE FORMAL TREATMENT OF THE CPF

The CPF takes three segments from three different Prelec functions (see Definition 2.22) and joins them up in a continuous manner. The two points, $\underline{p}, \bar{p}$, in Figure 2.17 are the cutoff points between the three Prelec functions. Define,

$$\underline{p} = e^{-\left(\frac{\beta}{\beta_0}\right)^{\frac{1}{\alpha_0-1}}}, \bar{p} = e^{-\left(\frac{\beta}{\beta_1}\right)^{\frac{1}{\alpha_1-1}}}. \quad (2.129)$$

Definition 2.54 (*al-Nowaihi and Dhami, 2010a*): By the composite Prelec weighting function we mean the probability weighting function $w : [0, 1] \rightarrow [0, 1]$ given by

$$w(p) = \begin{cases} 0 & \text{if } p = 0 \\ e^{-\beta_0(-\ln p)^{\alpha_0}} & \text{if } 0 < p \leq \underline{p} \\ e^{-\beta(-\ln p)^{\alpha}} & \text{if } \underline{p} < p \leq \bar{p} \\ e^{-\beta_1(-\ln p)^{\alpha_1}} & \text{if } \bar{p} < p \leq 1 \end{cases} \quad (2.130)$$

where $\underline{p}$ and $\bar{p}$ are given by (2.129) and

$$0 < \alpha < 1, \beta > 0; \alpha_0 > 1, \beta_0 > 0; \alpha_1 > 1, \beta_1 > 0, \beta_0 < 1/\beta^{\frac{\alpha_0-1}{1-\alpha}}, \beta_1 > 1/\beta^{\frac{\alpha_1-1}{1-\alpha}}. \quad (2.131)$$

The restrictions $\alpha > 0, \beta > 0, \beta_0 > 0$, and $\beta_1 > 0$, in (2.131), are required by the axiomatic derivations of the Prelec function; see *al-Nowaihi and Dhami (2006)*. The restriction $\beta_0 < 1/\beta^{\frac{\alpha_0-1}{1-\alpha}}$ and $\beta_1 > 1/\beta^{\frac{\alpha_1-1}{1-\alpha}}$ ensure that the interval $(\underline{p}, \bar{p})$ is not empty. The interval limits are chosen so that the CPF in (2.130) is continuous across them.

Proposition 2.23 (*al-Nowaihi and Dhami, 2010a*): The composite Prelec function (Definition 2.55) is a probability weighting function in the sense of Definition 2.19.

Define p_1, p_2, p_3 (see Figure 2.17) as follows:

$$p_1 = e^{-\left(\frac{1}{\beta_0}\right)^{\frac{1}{\alpha_0-1}}}, p_2 = e^{-\left(\frac{1}{\beta}\right)^{\frac{1}{\alpha-1}}}, p_3 = e^{-\left(\frac{1}{\beta_1}\right)^{\frac{1}{\alpha_1-1}}}. \quad (2.132)$$

Proposition 2.24 (*al-Nowaihi and Dhami, 2010a*): (a) $p_1 < \underline{p} < p_2 < \bar{p} < p_3$. (b) $p \in (0, p_1) \Rightarrow w(p) < p$. (c) $p \in (p_1, p_2) \Rightarrow w(p) > p$. (d) $p \in (\bar{p}, p_3) \Rightarrow w(p) < p$. (e) $p \in (p_3, 1) \Rightarrow w(p) > p$.

By Proposition 2.24, a CPF overweights low probabilities in the range (p_1, p_2) and underweights high probabilities in the range (p_2, p_3) , thus, it is inverse S-shaped in the middle. But, in addition, and unlike all the standard probability weighting functions, it underweights near-zero probabilities in the range $(0, p_1)$, and overweights near-one probabilities in the range $(p_3, 1)$ as required in H2a.

The restrictions $\alpha_0 > 1, \alpha_1 > 1$ in (2.131) ensure that a CPF has the following properties, that help explain human behavior for extremely low probability events.

Proposition 2.25 For the CPF (Definition 2.54), $\forall \gamma > 0$, the CPF (i) zero-underweights infinitesimal probabilities, i.e., $\lim_{p \rightarrow 0} \frac{w(p)}{p^\gamma} = 0$, and (ii) zero-overweights near-one probabilities, i.e., $\lim_{p \rightarrow 1} \frac{1-w(p)}{1-p} = 0$.

The CPF has not yet been empirically estimated, so its ability to explain H2 is still an open question. The larger number of parameters in the CPF relative to any of the standard weighting functions can pose difficulties in estimation. Indeed, as Peter Wakker points out in personal communication, one needs to conduct proper statistical tests such as Akaike's criterion for weighing the pros and the cons of additional parameters. We welcome such tests. However, most economic situations of interest where the CPF has real bite are low probability events. Over this range, a CPF is simply a Prelec function with $\alpha > 1$. Hence, at a very minimum, one might wish to estimate if for low probabilities we observe a separation of decision makers into two groups: One with a Prelec parameter $\alpha > 1$ and another with $\alpha < 1$. Clearly, these low probabilities must be lower relative to those that are used in the standard experiments, say, one in 10,000 or one in 20,000. These could be the approximate chances of a house being burgled in some area, or the probability of suffering an accident while driving and talking on a mobile phone.

DEFINITIONS OF CPT AND CRDU

We now give definitions of *composite prospect theory* (CPT) and *composite rank dependent utility* (CRDU).

Definition 2.55 (Composite prospect theory (CPT), al-Nowaihi and Dhami, 2010a): Under CPT, a fraction $1 - \mu$ of the population uses cumulative prospect theory, PT, with any of its standard probability weighting functions, such as the standard Prelec function, so $\lim_{p \rightarrow 0} \frac{w(p)}{p} = \infty$. This fraction conforms to hypothesis H2b. The remaining fraction, μ , uses PT but replaces the standard probability weighting function with the composite Prelec probability weighting function, CPF, for which $\lim_{p \rightarrow 0} \frac{w(p)}{p} = 0$ (see Propositions 2.12 and 2.25). This fraction conforms to H2a.

From Definition 2.55, for the respective fractions μ and $1 - \mu$, the decision weights in Definition 2.33 are calculated by using, respectively, the CPF and the standard probability weighting function. Once these probability weights are calculated, the objective function in Definition 2.34 is used to compute the utility from any lottery in incremental form.

Definition 2.56 (Composite rank dependent utility (CRDU); al-Nowaihi and Dhami, 2010a) In CRDU, RDU replaces PT in Definition 2.55.

From Definition 2.56, we use the appropriate probability functions (CPF for the fraction μ and any standard function for the fraction $1 - \mu$) to calculate the decision weights in Definition 2.25, followed by the calculation of the value of the objective function in Definition 2.26.

Other than the probability weighting functions, CPT and CRDU share all other elements with PT and RDU, respectively. In this sense, both are conservative extensions of existing and established theories. However, the extension enables the explanation of hypothesis H2, which is more general than H1. Thus, we have the nested structure: CPT can explain everything that PT can, which in turn can explain everything that RDU can, which in turn can explain everything that

EU can, which in turn explains everything that expected value maximization (EV) can. However, the converse is false in each case.⁸⁴

CRDU could be thought of as a special case of CPT in the absence of a reference point.⁸⁵ However, this rules out the psychologically powerful and robust notions of loss aversion and reference dependence, which have strong explanatory power in a variety of contexts and are too important to be ignored.⁸⁶ This is likely to weaken the explanatory power of CRDU relative to CPT.

CPT has potential applications to a wide range of economically important problems such as insurance, and crime and punishment. Some of these applications are considered in the exercises.

2.12 Risk preferences and time preferences

In economics, it is typical to give a separate treatment of risk preferences and time preferences and assume that these two kinds of preferences are unrelated. Yet, the outcomes of most human decisions are uncertain and take time to materialize, which is suggestive of a potential link between the two. An emerging literature has tried to grapple with the complex interaction between these two entities. The contributions have been both theoretical and empirical.⁸⁷ Issues of time discounting are considered more fully in Part 3 where we also consider the effect of uncertainty on time preferences. In this section, we consider some of the theoretical and empirical contributions that explore the links between these two areas but with greater focus on risk preferences.

2.12.1 Can non-linear probability weighting explain temporal behavior under risk?

We begin by considering the contribution by Andreoni and Sprenger (2012b). In order to separate the role of risk and time preferences, they use a test method known as *convex time budgets*. This method and the basic research question here is adapted from their earlier work in Andreoni and Sprenger (2012a). At date 0, subjects are allocated 100 tokens. They can allocate integer amounts between an earlier date, t , or a later date, $t + k$. A generic allocation is denoted by $(n, 100 - n)$, where $0 \leq n \leq 100$. These tokens can then be exchanged for cash at dates t and $t + k$, respectively, at exchange rates of a_t and a_{t+k} per token. Two values of k are considered: $k = 28$ and $k = 56$.

In the experiments, $a_{t+k} = \$0.2$ per token and $a_t = \frac{a_{t+k}}{1+r}$, where r is the gross interest rate over k days.⁸⁸ The interest rate was varied from 0% to 2116.6% in discrete steps. Suppose that we refer to the income realized at any date as the consumption at that date. The future is uncertain in the sense that the payments at t and $t + k$, respectively, are honored with probabilities p_1 and p_2 that

⁸⁴ See al-Nowaihi and Dhami (2010a) for these claims. CPT can also incorporate the insight from *third generation prospect theory* of Schmidt et al. (2008). One may also refer to CPT as *fourth generation prospect theory* if one prefers this kind of terminology.

⁸⁵ This requires, of course, that we collapse lotteries in incremental form in CPT back to ordinary lotteries first (i.e., add the reference point to each outcome), and then apply CRDU.

⁸⁶ See, for instance, the applications discussed in Camerer (2000), Kahneman and Tversky (2000), and Barberis and Thaler (2003).

⁸⁷ For the theoretical contributions, see, for instance, Halevy (2008), Walther (2010), Baucells and HeuKamp (2012), Epper and Fehr-Duda (2014b). For the empirical contributions, see for instance, Andreoni and Sprenger (2012b) and Epper and Fehr-Duda (2014a).

⁸⁸ The standardized daily net interest is, therefore, $(1 + r)^{1/k} - 1$.

are independent of each other. If the payments at both dates go through, then the consumption stream is given by

$$(c_t, c_{t+k}) = (na_t, (100 - n)a_{t+k}). \quad (2.133)$$

Using $a_t = \frac{a_{t+k}}{1+r}$, the budget constraint of the consumer, expressed in date $t+k$ units, is given by⁸⁹

$$(1+r)c_t + c_{t+k} = 100a_{t+k}. \quad (2.134)$$

Given the uncertainty in the payments, the decision maker faces the following prospect

$$L(p_1, p_2) = ((c_t, c_{t+k}), p_1 p_2; (c_t, 0), p_1(1 - p_2); (0, c_{t+k}), (1 - p_1)p_2; (0, 0), (1 - p_1)(1 - p_2)). \quad (2.135)$$

In order to evaluate the utility of the consumer, we follow the *portfolio approach* in Chew and Epstein (1990). There are at least two other possibilities—the *separable case* and the *recursive case*; see Appendix A1 in Epper and Fehr-Duda (2014a). Andreoni and Sprenger (2012b) use the separable case. The separable and recursive cases do not correctly predict the behavior in the left panel of Figure 2.19, while the portfolio case correctly predicts it, as we show below. It is worth noting how the three cases reduce lotteries when both time and risk are present; see the self-explanatory Figure 2.18.

In this section, we shall consider only the portfolio case. Denote the *intertemporal discounted utility* from the consumption profile (c_t, c_{t+k}) by $D(t)v(c_t) + D(t+k)v(c_{t+k})$, where D is a discount function that is differentiable throughout, is decreasing, and has continuous partial derivatives up to order two. v is a utility function such that

$$v(0) = 0, v' > 0 \text{ and } v'' < 0. \quad (2.136)$$




Figure 2.18 Reduction of lotteries when both time and risk are present, in three different cases.

Source: Epper and Fehr-Duda (2014a) with permission from the American Economic Association.

⁸⁹ We could alternatively have expressed the budget constraint in date t units as $c_t + \frac{c_{t+k}}{(1+r)} = 100a_t$.

In a baseline condition, (1, 1), the prospect is $L(1, 1)$ and it is compared to a scaled down version (λ, λ) , with the corresponding prospect $L(\lambda, \lambda)$, where $0 < \lambda < 1$, in order to check for the common ratio effect. Two different risk conditions are considered, (1, q) and (q , 1). In each case, their respective scaled down conditions $(\lambda, \lambda q)$ and $(\lambda q, \lambda)$ are also considered to check for the common ratio effect under risk.

Proposition 2.26 *In any delay-discounting model (exponential discounting is a special case), where the decision maker follows expected utility, suppose that the optimal number of tokens purchased by the decision maker is n^* .*

- (i) *The optimal number of tokens, n^* , and the optimal consumption levels, c_t^* and c_{t+k}^* , do not change when p_1 and p_2 are scaled up or down by a common factor.*
- (ii) *The optimal number of tokens, n^* , and optimal consumption c_t^* , decrease as the interest rate, r , goes up.*

Proof: (i) In any delay-discounting model, the utility of the prospect in (2.135), given $v(0) = 0$, is given by

$$V_D = p_1 p_2 [D(t)v(c_t) + D(t+k)v(c_{t+k})] + p_1(1-p_2)D(t)v(c_t) + (1-p_1)p_2 D(t+k)v(c_{t+k}). \quad (2.137)$$

The decision maker optimally chooses c_t and c_{t+k} , subject to the constraint in (2.134). Substituting c_{t+k} from (2.134) in (2.137), we get the following unconstrained problem in which the decision maker chooses c_t ; the number of tokens can then be recovered from (2.133). The unconstrained problem is

$$c_t^* \in \arg \max_{0 \leq c_t \leq 100a_t} V_D = p_1 D(t)v(c_t) + p_2 D(t+k)v(100a_{t+k} - (1+r)c_t).$$

Using (2.133), the optimal number of tokens is $n^* = \frac{c_t^*}{a_t}$. V_D is a strictly concave function on a compact set, so the condition for a unique interior maximum value, c_t^* , is given by

$$\frac{v'(c_t)}{v'(c_{t+k})} = (1+r) \frac{D(t+k)}{D(t)} \frac{p_2}{p_1}. \quad (2.138)$$

(2.138) is the key condition that unlocks the main comparative statics in the model. A scaling of p_1 and p_2 by a common factor λ does not alter the ratio $\frac{p_2}{p_1}$. Thus, both sides of (2.138) are unaffected, so c_t^* (and n^*) is unchanged.

(ii) Given that v is strictly concave, any increase in the LHS must relatively increase $v'(c_t)$, or equivalently, reduce c_t . As r increases, the RHS of (2.138) increases so c_t^* falls and $n^* = \frac{c_t^*}{a_t}$ also falls. ■

Based on Proposition 2.26 we expect the following pairs of conditions to give identical results in terms of the optimal consumption levels: (i) (1, 1) and (λ, λ) . (ii) (1, q) and $(\lambda, \lambda q)$. (iii) (q , 1) and $(\lambda q, \lambda)$. A representative set of results is shown in Figure 2.19, which reports the average consumption level of decision makers at date t , $\bar{c}_t$, as a function of $(1+r)\frac{p_2}{p_1}$.⁹⁰

⁹⁰ Andreoni and Sprenger (2012b) give the full set of results, i.e., first for $k = 26$ they give a comparison of the three pairs of conditions listed above, for $\lambda = 0.5$ and $q = 0.8$. Then they repeat the same comparison for $k = 52$.




Figure 2.19 Average consumption, $\bar{c}_t$, in various conditions. In the left panel, the delay is $k = 28$ days, and we compare the two conditions $(1, 1)$ and $(0.5, 0.5)$. In the right panel, we consider a delay of $k = 56$ days, and we compare the conditions $(1, 0.8)$ versus $(0.5, 0.4)$ and $(0.8, 1)$ and $(0.4, 0.5)$.

Source: Epper and Fehr-Duda (2014a) with permission from the American Economic Association.

Figure 2.19 shows that Proposition 2.26 is violated. For instance, in the left panel of Figure 2.19, the average optimal level of consumption is predicted to be identical in the two conditions $(1, 1)$ and $(0.5, 0.5)$, yet average consumption is initially higher in the $(1, 1)$ condition and then drops below the level for the $(0.5, 0.5)$ condition, i.e., the curves cross over. This result has been replicated by Cheung (2014) and Miao and Zhong (2012). In the $(1, 1)$ condition, where there is complete certainty of payment, about 80% of the allocations are at one corner ($n^* = 100$ or $n^* = 0$). But in the risky condition, $(0.5, 0.5)$, subjects try to diversify and only 26.1% are at a corner solution.

Consider the right panel of Figure 2.19 where $q = 0.8$ and $\lambda = 0.5$. First consider the conditions $(1, q)$ versus $(\lambda, \lambda q)$. The average consumption in the $(0.5, 0.4)$ condition is always lower than the $(1, 0.8)$ condition. The converse is true in a comparison of the conditions $(q, 1)$ and $(\lambda q, \lambda)$ (i.e., the conditions $(0.8, 1)$ and $(0.4, 0.5)$).

Andreoni and Sprenger (2012b) argue that models of non-linear probability weighting, such as RDU, cannot explain these results either. Their preferred explanation is models where the common ratio property (see (2.138)) does not hold; models in this category include those by Kreps and Porteus (1978), Chew and Epstein (1989b), and Epstein and Zin (1989). However, reworking through the same set of results, Epper and Fehr-Duda (2014a) show that the results can be explained by RDU if the decision maker cares for portfolio risk rather than employ the separable case that is used by Andreoni and Sprenger (2012b); see Figure 2.18.

Let us now consider a decision maker who follows RDU and faces the lottery in (2.135). There are four possible outcomes, but each outcome is a pair of consumption values. How are these pairs to be ranked? Consider the binary relation $<$ such that $(c_t, 0) < (0, c_{t+k})$ means that c_{t+k} received at date $t + k$ is “ranked higher” than c_t received at date t . We assume that

By contrast, we report in Figure 2.19, a comparison between the first pair for $k = 26$ and a comparison for the remaining two pairs for $k = 52$.

$$\begin{cases} (0, c_{t+k}) \prec (c_t, 0) & \text{if } D(t+k)v(c_{t+k}) < D(t)v(c_t) \\ (c_t, 0) \prec (0, c_{t+k}) & \text{if } D(t)v(c_t) < D(t+k)v(c_{t+k}) \end{cases} \quad (2.139)$$

With this definition, it is always the case that $(0, 0) \prec (c_t, 0) \prec (c_t, c_{t+k})$ and $(0, 0) \prec (0, c_{t+k}) \prec (c_t, c_{t+k})$. Depending on which row of (2.139) applies, we get two possibilities, I and II.

$$\begin{cases} I: & (0, 0) \prec (0, c_{t+k}) \prec (c_t, 0) \prec (c_t, c_{t+k}) \\ II: & (0, 0) \prec (c_t, 0) \prec (0, c_{t+k}) \prec (c_t, c_{t+k}) \end{cases} \quad (2.140)$$

The lottery in (2.135) gives the associated probability of each of the outcomes in cases I and II in (2.140). Let w be a probability weighting function. The decision weights (see Definition 2.25) of the lowest and the highest outcomes in cases I and II are identically given by

$$\pi(0, 0) = 1 - w(p_1 + p_2 - p_1 p_2), \pi(c_t, c_{t+k}) = w(p_1 p_2). \quad (2.141)$$

The decision weights for the two intermediate outcomes in cases I and II are:

$$\begin{cases} I: & \pi_I(0, c_{t+k}) = w(p_1 + p_2 - p_1 p_2) - w(p_1), \pi_I(c_t, 0) = w(p_1) - w(p_1 p_2) \\ II: & \pi_{II}(c_t, 0) = w(p_1 + p_2 - p_1 p_2) - w(p_2), \pi_{II}(0, c_{t+k}) = w(p_2) - w(p_1 p_2) \end{cases} \quad (2.142)$$

Let the utility function of an outcome continue to be given by v , with the conditions in (2.136). Given the decision weights in (2.141) and (2.142), we can now define the objective function of the decision maker who follows RDU in the two cases in (2.139), V_R^i , $i = I, II$.

$$\begin{aligned} V_R^i = & \pi_i(0, c_{t+k})D(t+k)v(c_{t+k}) + \pi_i(c_t, 0)D(t)v(c_t) \\ & + \pi(c_t, c_{t+k})[D(t)v(c_t) + D(t+k)v(c_{t+k})], i = I, II. \end{aligned}$$

A decision maker who follows RDU, maximizes his objective function V_R^i , $i = I, II$, subject to the constraint in (2.134). The first order condition, the analogue of (2.138), is given by

$$\frac{v'(c_t)}{v'(c_{t+k})} = (1+r) \frac{D(t+k)}{D(t)} \varphi(p_1, p_2), \quad (2.143)$$

where,

$$\varphi(p_1, p_2) = \begin{cases} \frac{w(p_2)}{w(1)} & \text{if } p_1 = 1 \\ \frac{w(1)}{w(p_1)} & \text{if } p_2 = 1 \\ \frac{\varphi_I(p_1, p_2) = \frac{w(p_1 p_2) + w(p_1 + p_2 - p_1 p_2) - w(p_1)}{w(p_1)}}{\frac{w(p_2)}{w(p_1)}} & \text{if } p_1 \neq 1, p_2 \neq 1, (0, c_{t+k}) \prec (c_t, 0) \\ \frac{\varphi_{II}(p_1, p_2) = \frac{w(p_1 p_2) + w(p_1 + p_2 - p_1 p_2) - w(p_2)}{w(p_2)}}{\frac{w(p_1)}{w(p_2)}} & \text{if } p_1 \neq 1, p_2 \neq 1, (c_t, 0) \prec (0, c_{t+k}) \end{cases} \quad (2.144)$$

In order to compare the average level of consumption under delay-discounting models and RDU we compare (2.138) and (2.143). The relevant comparison is between the terms $\frac{p_2}{p_1}$ and $\varphi(p_1, p_2)$.

We now consider the first order condition (2.143) for the three cases: (i) $(1, 1)$ versus (λ, λ) , (ii) $(1, q)$ versus $(\lambda, \lambda q)$, and (iii) $(q, 1)$ versus $(\lambda q, \lambda)$. Some illustrative results for the case of the

Table 2.5 Various computations to explain the relevant empirical findings in different conditions.

Conditions (1,1) and (0.5,0.5)	Conditions (1,0.8) and (0.5,0.4)	Conditions (0.8,1) and (0.4,0.5)
$\frac{p_1}{p_2} = 1$	$\frac{p_1}{p_2} = 1.25$	$\frac{p_1}{p_2} = 0.8$
$\varphi(1, 1) = 1$	$\varphi(1, 0.8) = 0.624$	$\varphi(0.8, 1) = 1.604$
$\varphi_I(0.5, 0.5) = 1.053$	$\varphi_I(0.5, 0.4) = 0.912$	$\varphi_I(0.4, 0.5) = 1.166$
$\varphi_{II}(0.5, 0.5) = 0.950$	$\varphi_{II}(0.5, 0.4) = 0.858$	$\varphi_{II}(0.4, 0.5) = 1.097$

Source: Epner and Fehr-Duda (2014a).

Prelec function (Definition 2.22) with $\alpha = 0.5$ and $\beta = 1$, and for $\lambda = 0.5$ and $q = 0.8$ are summarized in Table 2.5.

We are now interested in the empirical findings encapsulated in Figure 2.19.

I. Baseline condition: Consider first the comparison between the conditions (1, 1) and (λ, λ) . The illustrative example in Table 2.5 (first column) is consistent with the crossover of the two curves in the left panel of Figure 2.19 ($\varphi_{II}(0.5, 0.5) < \varphi(1, 1) < \varphi_I(0.5, 0.5)$). Now consider the more general case. For most probability weighting functions, including the Prelec function, that are inverse S-shaped, there is an initial concave segment, followed by a convex segment. The point of inflection that separates these segments is typically well below the probability of 0.5, which is used in the experiments in Andreoni and Sprenger (2012b), hence, we are in the convex region of the weighting function. Assuming a convex probability weighting function, we get

$$w(ap_1 + (1-a)p_2) < aw(p_1) + (1-a)w(p_2), 0 < a < 1. \quad (2.145)$$

From (2.144)

$$\varphi_I(\lambda, \lambda) = \frac{w(\lambda^2) + w(2\lambda - \lambda^2) - w(\lambda)}{w(\lambda)}, \varphi_{II}(\lambda, \lambda) = \frac{w(\lambda)}{w(\lambda^2) + w(2\lambda - \lambda^2) - w(\lambda)}. \quad (2.146)$$

From (2.145), (2.146) we get

$$\varphi_I(\lambda, \lambda) > \varphi(1, 1) = 1 > \varphi_{II}(\lambda, \lambda). \quad (2.147)$$

When $(0, c_{t+k}) < (c_t, 0)$ we have $\varphi_I(\lambda, \lambda) > 1$, thus, the RHS of (2.143) is greater under the condition (λ, λ) relative to the condition (1, 1). Since v is concave, it follows that c_t is relatively lower under the condition (λ, λ) . On the other hand, in the condition $(c_t, 0) < (0, c_{t+k})$ we have that $1 > \varphi_{II}(\lambda, \lambda)$ so c_t is greater under the condition (λ, λ) relative to the condition (1, 1). This implies the crossing over of the two curves that we observe in the left panel of Figure 2.19.

II. Differential risk conditions: Consider first the comparison $(1, q)$ versus $(\lambda, \lambda q)$ (the other comparison is $(q, 1)$ versus $(\lambda q, \lambda)$). In this case, from (2.144)

$$\begin{cases} \varphi_I(\lambda, \lambda q) = \frac{w(\lambda^2 q) + w(\lambda + \lambda q - \lambda^2 q) - w(\lambda)}{w(\lambda)} \\ \varphi_{II}(\lambda, \lambda q) = \frac{w(\lambda q)}{w(\lambda^2 q) + w(\lambda + \lambda q - \lambda^2 q) - w(\lambda q)} \end{cases}. \quad (2.148)$$

In the experiments of Andreoni and Sprenger (2012b), $\lambda = 0.5$ and $q = 0.8$. Most probability weighting functions are quite flat over this range. Thus, let us use a linear approximation

$w(p) \approx \widehat{w}(p) = a + bp$. Suppose that we apply this approximation to $\varphi_I(\lambda, \lambda q)$ in (2.148), then we get $\frac{a+b\lambda q}{a+b\lambda} = \frac{\widehat{w}(\lambda q)}{\widehat{w}(\lambda)}$. The same approximation applies to $\varphi_{II}(\lambda, \lambda q)$, thus,

$$\varphi_I(\lambda, \lambda q) \approx \frac{\widehat{w}(\lambda q)}{\widehat{w}(\lambda)} \approx \varphi_{II}(\lambda, \lambda q). \quad (2.149)$$

Thus, $\frac{\varphi_I(\lambda, \lambda q)}{\varphi_{II}(\lambda, \lambda q)} \approx 1$. From the RHS of (2.143), optimal consumption c_t is, thus, almost unchanged as the rank of outcomes change between $(0, c_{t+k}) < (c_t, 0)$ and $(c_t, 0) < (0, c_{t+k})$. From the subproportionality condition in Definition 2.24, applied to the function $\widehat{w}$ for $p = 1$ (this is the special case of the certainty effect), we get that $\frac{\widehat{w}(\lambda q)}{\widehat{w}(\lambda)} > \widehat{w}(q)$, so

$$\frac{w(\lambda q)}{w(\lambda)} \approx \frac{\widehat{w}(\lambda q)}{\widehat{w}(\lambda)} > \widehat{w}(q) \approx w(q). \quad (2.150)$$

For the condition $(1, q)$ we have, from (2.144),

$$\varphi(1, q) = \frac{w(q)}{w(1)} = w(q). \quad (2.151)$$

From (2.149)–(2.151) we get that

$$\varphi_I(\lambda, \lambda q) \approx \varphi_{II}(\lambda, \lambda q) > w(q) = \varphi(1, q). \quad (2.152)$$

Using (2.152) we get that the RHS in (2.143) is relatively higher in the condition $(\lambda, \lambda q)$ relative to the condition $(1, q)$ throughout the domain (even as we move across the two possibilities $(0, c_{t+k}) < (c_t, 0)$ and $(c_t, 0) < (0, c_{t+k})$). This feature is also present in the illustrative example as shown in the second column in Table 2.5. Thus, we expect average date t consumption, $\bar{c}_t$, to be relatively higher in the condition $(1, q)$, exactly as we have in the right panel of Figure 2.19. We leave it to the reader to show analogously that a comparison of the conditions $(q, 1)$ and $(\lambda q, \lambda)$ is also consistent with the right panel of Figure 2.19.

Thus, if we use the portfolio method, the results of Andreoni and Sprenger (2012b) are consistent with a decision maker who follows RDU in evaluating uncertain dated streams of consumption goods.

2.12.2 Intertemporal risk aversion

Epper and Fehr-Duda (2014b) make preliminary progress in operationalizing *intertemporal risk aversion*, based on the work of Richard (1975) and Bommier (2007). Intertemporal risk aversion refers to the desire of decision makers to diversify consumption across time; empirical evidence can be found in Cheung (2014) and Miao and Zhong (2012). Suppose that $c_t \in \{\underline{c}_t, \bar{c}_t\}$ and $c_{t+k} \in \{\underline{c}_{t+k}, \bar{c}_{t+k}\}$, where $\underline{c}_t < \bar{c}_t$ and $\underline{c}_{t+k} < \bar{c}_{t+k}$. Define the prospects

$$L^- = ((\bar{c}_t, \underline{c}_{t+k}), 0.5; (\underline{c}_t, \bar{c}_{t+k}), 0.5), L^+ = ((\bar{c}_t, \bar{c}_{t+k}), 0.5; (\underline{c}_t, \underline{c}_{t+k}), 0.5). \quad (2.153)$$

The prospect L^+ is not as well diversified as prospect L^- . This leads to the following definition.

Definition 2.57 (*Intertemporal risk aversion*): A decision maker is intertemporally risk averse if $L^+ \prec L^-$, where $\prec$ is the “strictly preferred to” binary relation.

Define

$$V_D(x_t, y_{t+k}) = D(t)v(x_t) + D(t+k)v(y_{t+k}),$$

and assume, as in (2.139), that

$$(x_t, y_{t+k}) \prec (x'_t, y'_{t+k}) \Leftrightarrow V_D(x_t, y_{t+k}) < V_D(x'_t, y'_{t+k}). \quad (2.154)$$

(2.154) allows us to express the preference $L^+ \prec L^-$ for decision makers who follow RDU. It is not clear if $(\bar{c}_t, \underline{c}_{t+k}) \prec (\underline{c}_t, \bar{c}_{t+k})$ or $(\underline{c}_t, \bar{c}_{t+k}) \prec (\bar{c}_t, \underline{c}_{t+k})$. Thus, in specifying the utility function of a RDU decision maker, V_R (Definitions 2.25 and 2.26), we need to consider two possibilities.

$$V_R(L^-) = \begin{cases} (1 - w(0.5)) V_D(\bar{c}_t, \underline{c}_{t+k}) + w(0.5) V_D(\underline{c}_t, \bar{c}_{t+k}) & \text{if } (\bar{c}_t, \underline{c}_{t+k}) \prec (\underline{c}_t, \bar{c}_{t+k}). \\ (1 - w(0.5)) V_D(\underline{c}_t, \bar{c}_{t+k}) + w(0.5) V_D(\bar{c}_t, \underline{c}_{t+k}) & \text{if } (\underline{c}_t, \bar{c}_{t+k}) \prec (\bar{c}_t, \underline{c}_{t+k}). \end{cases} \quad (2.155)$$

The ranking of outcomes under the prospect L^+ is unambiguous, hence, using Definitions 2.25 and 2.26

$$V_R(L^+) = (1 - w(0.5)) V_D(\underline{c}_t, \underline{c}_{t+k}) + w(0.5) V_D(\bar{c}_t, \bar{c}_{t+k}). \quad (2.156)$$

Comparing (2.155) and (2.156) we get

$$V_R(L^-) - V_R(L^+) = \begin{cases} (1 - 2w(0.5)) (v(\bar{c}_t) - v(\underline{c}_t)) D(t) & \text{if } (\bar{c}_t, \underline{c}_{t+k}) \prec (\underline{c}_t, \bar{c}_{t+k}). \\ (1 - 2w(0.5)) (v(\bar{c}_{t+k}) - v(\underline{c}_{t+k})) D(t+k) & \text{if } (\underline{c}_t, \bar{c}_{t+k}) \prec (\bar{c}_t, \underline{c}_{t+k}). \end{cases}$$

In each case, the condition for intertemporal risk aversion, $V_R(L^-) > V_R(L^+)$, holds if

$$1 - 2w(0.5) > 0, \text{ or } w(0.5) < 0.5. \quad (2.157)$$

This condition is satisfied for most inverse S-shaped probability weight functions; see Epper and Fehr-Duda (2014b). For instance it is satisfied by the Prelec function for commonly used values of α, β . In a nutshell, a decision maker who follows RDU also exhibits intertemporal risk aversion if the condition in (2.157) holds.

2.12.3 An agenda for the future

In an insightful paper, Epper and Fehr-Duda (2014b) lay out the agenda for future research that wishes to explore the link between risk preferences and time preferences. They propose a general theoretical framework designed to address a set of stylized facts that any aspiring theory in this area may wish to take into account. The stylized facts can be organized as follows.

1. Empirically observed risk tolerance and patience are *delay dependent*: People appear more risk tolerant for events that are to happen in the more distant future.⁹¹ An implication is that

⁹¹ See, for instance, Abdellaoui, Diecidue, and Öncüler (2011) and Coble and Lusk (2010).

people may be more tolerant in facing risks relating to global warming or insuring against natural hazards if the consequences are in the more distant future. Hyperbolic discounting captures delay dependence of behavior but does not incorporate risk.

2. Behavior under risk and time is *process dependent*. There is a very large number of process dependent behaviors that have been observed. For instance, it matters whether we first reduce lotteries in a one-shot manner or sequentially (see for instance, Figure 2.18). In other contexts, it matters if uncertainty is resolved gradually, or in the terminal period. For instance, a greater risk premium may be required on a portfolio whose returns are evaluated frequently rather than observed at the end of the period (Benartzi and Thaler, 1995; Barberis et al., 2001). More generally, in a range of contexts, people take on more risks if they are informed of the outcomes at the end rather than frequently over time.⁹² We shall also consider process dependence of time behavior in Part 3 of the book; in particular, we shall consider *subadditive discounting* in some detail.
3. The *timing of the resolution of uncertainty* affects behavior of subjects when they face situations that have a risk and time component: For instance, for negative news, people may prefer uncertainty to resolve later rather than sooner (e.g., is one affected by a genetic disorder that runs in the family?). For positive news, people may prefer uncertainty to resolve sooner rather than later (e.g., has the book that I ordered arrived?).⁹³
4. Time discount differs in the presence of risk: For instance, risky future incomes are discounted less than safe future incomes (Ahlbrecht and Weber, 1997a; Weber and Chapman, 2005). This is in contrast to the exponential discounting utility model that applies the same discounting to safe and risky returns. A literature also examines the dynamic inconsistency of choices in the presence of risk.⁹⁴
5. Behavior differs depending on whether we first reduce a lottery in the time dimension or in the risk dimension: It appears that the value assigned to prospects is lower if they are first reduced for risk and then for time (as compared to the reverse case). See, for instance, the recursive case in Figure 2.18 and for the empirical evidence, see Öncüler and Onay (2009).

We now consider some of these issues.

2.12.4 Delay dependence and probability weighting

Suppose that we have a binary lottery

$$L = (x, 1 - p; y, p), 0 \leq x < y. \quad (2.158)$$

A RDU decision maker evaluates the lottery L by

$$V_R(L) = [1 - w(p)]v(x) + w(p)v(y). \quad (2.159)$$

In (2.159), w is the “true” weighting function of the decision maker (i.e., it is a part of the decision maker’s preferences). We assume that w is subproportional (Definition 2.24).

⁹² See for instance, Gneezy and Potters (1997), Thaler et al. (1997), Gneezy et al. (2003), and Bellemare et al. (2005).

⁹³ For the empirical evidence, see Ahlbrecht and Weber (1997b), Lovallo and Kahneman (2000), Eliaz and Schotter (2007), and von Gaudecker et al. (2011). For papers that propose a theoretical framework, see Kreps and Porteus (1978), Chew and Epstein (1989b), and Grant et al. (2000).

⁹⁴ See, for instance, Cubitt et al. (1998) and Busemeyer et al. (2000).

Suppose that time is discrete, and the lottery L is paid out at time $t > 0$. Then we make two assumptions.

1. Utility is discounted back to time 0 using the exponential function

$$D(t) = e^{-\theta t},$$

where $\theta > 0$ is the discount rate.

2. There is uncertainty about whether the payments will be made in the future. Suppose that one is promised a payment at time 0 of an amount $z > 0$ to be received at time t . Then the probability of receiving the payment at time t is s^t , $0 < s \leq 1$, where s is the per-period probability of the payment being honored. Hence, if the prospect L in (2.158) is to be realized at time t , it is viewed at time 0 as the prospect⁹⁵

$$\tilde{L} = (0, 1 - s^t; x, (1 - p)s^t; y, ps^t), 0 \leq x < y, 0 < s \leq 1.$$

Under Assumptions 1 and 2, a decision maker, who uses RDU, evaluates at time 0 the prospect $\tilde{L}$ to be received at time t , as follows (see Definitions 2.25 and 2.26),

$$V_R(\tilde{L}; 0) = D(t) [(w(s^t) - w(ps^t)) v(x) + w(ps^t) v(y)], \quad (2.160)$$

where w is the *atemporal probability weighting function*. (2.160) can be rewritten as

$$V_R(\tilde{L}; 0) = D(t) w(s^t) \left[(v(y) - v(x)) \frac{w(ps^t)}{w(s^t)} + v(x) \right]. \quad (2.161)$$

Suppose that an observer does not realize that Assumption 2 applies; for instance, an experimenter (observer) may have no desire to withhold future payments to subjects (decision makers), but subjects may assign a positive probability to this event. Thus, the observer thinks that the decision maker evaluates lottery $\tilde{L}$ with $s = 1$, so, recalling that $w(1) = 1$, he imagines that

$$V_R(\tilde{L}; 0) = \hat{D}(t) [(v(y) - v(x)) \hat{w}(p) + v(x)], \quad (2.162)$$

where $\hat{w}$ and $\hat{D}$ are, respectively, the probability weighting function and the discount function, which the observer believes are used by the decision maker. Suppose that the observer uses (2.162) to estimate the true weighting function and the discount function of the decision maker. Then, comparing (2.161) and (2.162), the relation between his estimates, $\hat{w}$ and $\hat{D}$, and the true functions w and D is

$$\hat{w}(p) = \frac{w(ps^t)}{w(s^t)}. \quad (2.163)$$

$$\hat{D}(t) = w(s^t) D(t). \quad (2.164)$$

Epper and Fehr-Duda (2014b) use this simple framework to address the five stylized facts in Section 2.12.3, above. Consider the first hypothesis on delay dependence.

⁹⁵ This assumption is motivated by the work of Halevy (2008) and Walther (2010); see Part 3 of the book for more details.

Proposition 2.27 *Let the true probability weighting function w be subproportional (Definition 2.24), and $s < 1$. Then the properties of the estimated function, $\hat{w}$ (see (2.163)) are as follows.*

- (i) $\hat{w}$ is a probability weighting function.
- (ii) $\hat{w}$ is subproportional.
- (iii) $\hat{w}$ is more elevated than w . The degree of elevation of $\hat{w}$ increases with time delay, t , and survival risk, $1 - s$, but this increase is at a decreasing rate.
- (iv) $\hat{w}$ is less elastic than w .
- (v) The increase in risk tolerance is greater when w is more subproportional.

Proof: (i) Using (2.163), $\hat{w}(0) = \frac{w(0)}{w(s^t)} = 0$, $\hat{w}(1) = \frac{w(s^t)}{w(s^t)} = 1$, and $\hat{w}'(p) = s^t \frac{w'(ps^t)}{w(s^t)} > 0$, hence, $\hat{w}$ is a probability weighting function.

(ii) Using the definition of subproportionality in Definition 2.24, and noting that w is subproportional, we have for $0 < \lambda < 1$ and $q < p \leq 1$,

$$\frac{\hat{w}(p)}{\hat{w}(q)} = \frac{w(ps^t)}{w(qs^t)} > \frac{w(\lambda ps^t)}{w(\lambda qs^t)} = \frac{\frac{w(\lambda ps^t)}{w(s^t)}}{\frac{w(\lambda qs^t)}{w(s^t)}} = \frac{\hat{w}(\lambda p)}{\hat{w}(\lambda q)},$$

which proves the result.

(iii) Using the definitions of $w(p)$, $\hat{w}(p)$, and subproportionality, we get

$$\frac{1}{w(p)} = \frac{w(1)}{w(p)} > \frac{w(s)}{w(ps)} > \frac{w(s^2)}{w(ps^2)} > \dots > \frac{w(s^t)}{w(ps^t)} = \frac{1}{\hat{w}(p)},$$

hence, $w(p) < \hat{w}(p)$. Thus, $\hat{w}$ is more elevated than w . Differentiating $\hat{w}$ with respect to t and s gives the remaining results.

(iv) The elasticity, ε_w , of the weighting function, w is defined as $\varepsilon_w(p) = \frac{w'(p)p}{w(p)}$. It can be shown that if w is subproportional, then $\varepsilon'_w(p) > 0$. Thus, by the definition of $\hat{w}$, we have

$$\varepsilon_{\hat{w}}(p) = \frac{\hat{w}'(p)p}{\hat{w}(p)} = \frac{w'(ps^t)ps^t}{w(ps^t)} = \varepsilon_w(ps^t) < \varepsilon_w(p),$$

where the last inequality follows from $ps^t < p$ and $\varepsilon'_w(p) > 0$.

(v) We leave the proof of this part to the reader. ■

One of the central results here is Proposition 2.27(iii). The fact that the estimated probability weighting function, $\hat{w}$, is relatively more elevated means that the observed risk tolerance of the decision maker is also relatively higher (see Example 2.10). Thus, risk tolerance appears, to the observer, to increase with the time delay. Furthermore, this risk tolerance appears even larger as t increases. The empirical evidence is consistent with these findings. In particular, the utility function is not influenced by time delay, so the relevant channel operates through probability weighting;⁹⁶ see Figure 2.22 for an illustration of the delay effect that we discuss below.

From (2.164), $\hat{D}(t) = w(s^t)D(t)$. D takes the exponential form. Suppose that w is the inverse S-shaped Prelec function with $\alpha\beta = 1$ (Definition 2.22), so $w(s^t) = e^{(-\ln s^t)^\alpha}$. Thus, $\hat{D}(t) = e^{-\hat{\theta}t}$

⁹⁶ See Epper and Fehr-Duda (2014b) for the references.




Figure 2.20 A plot of the estimated discount rate $\hat{\theta}$ (vertical axis) and the time delay t (horizontal axis). The weighting function is a Prelec function with $\alpha = 0.5$ and $\beta = 1$.

Source: Epper and Fehr-Duda (2014b).

where $\hat{\theta} = \theta - \ln(s) > \theta$. Thus, the presence of uncertainty increases the observed discount rates and increases the degree of observed impatience. However, this does not account for the typical finding of decreasing discount rates as, say, in hyperbolic discounting (because θ is independent of t). However, this can be accommodated if the probability weighting function is subproportional, an idea that is closely related to Halevy (2008) that we consider in more detail in Part 3 of the book. We summarize the results in the next proposition that we leave the proof to the reader as an exercise.

Proposition 2.28 (Epper and Fehr-Duda, 2014b): Suppose that w is subproportional.

- (i) $\hat{D}$ is a discount function if $1 < s \leq 1$, i.e., it is decreasing in t (impatience), $\lim_{t \rightarrow \infty} \hat{D}(t) = 0$, and $\hat{D}(0) = 1$.
- (ii) When $s < 1$ we have $\hat{\theta} > \theta$.
- (iii) The observed discount rate $\hat{\theta}$ is declining in t .
- (iv) Greater risk of survival (lower s) results in more steeply declining discount rates.
- (v) An increase in the degree of subproportionality of the weighting function leads to more steeply declining discount rates.

For the Prelec probability weighting function with $\alpha = 0.5$ and $\beta = 1$, Figure 2.20 shows the effect of varying survival risk, s , on the observed discount rate, $\hat{\theta}$ (vertical axis), while time t is on the horizontal axis. Thus, the estimated discount rate is hyperbolic and the degree of hyperbolicity is increasing with an increase in risk (lower s).

2.12.5 Probability weighting and process dependence

Empirical evidence suggests that the probability weighting function is more elevated when future rewards are received with a longer delay, while the utility function appears stable under delay;




Figure 2.21 A choice between a compound lottery and its equivalent simple lottery.

Source: Reproduced with permission of *Annual Review of Economics*, Volume 4, Helga Fehr-Duda and Thomas Epper, "Probability and risk: foundations and economic implications of probability dependent risk preferences," © 2012 by Annual Reviews <http://www.annualreviews.org>

see Abdellaoui, Diecidue et al. (2011). Fehr-Duda and Epper (2012) give the example of a choice between a compound lottery and its equivalent simple lottery shown in Figure 2.21. The simple lottery is $L_s = (x_1, 1 - p; x_2, p)$, where $x_1 < x_2$, $p = rq$ and $1 - p = q(1 - r) + (1 - q) = 1 - rq$.

The RDU of the simple lottery is

$$V_R(L_s) = [1 - w(rq)] u(x_1) + w(rq) u(x_2). \quad (2.165)$$

Suppose that the compound lottery is reduced recursively. Then with probability q , the RDU of the decision maker is $[1 - w(r)] u(x_1) + w(r) u(x_2)$, and with probability $1 - q$ it is $u(x_1)$. Hence, at the initial node, the decision maker faces the lottery

$$L_c = (u(x_1), 1 - q; [1 - w(r)] u(x_1) + w(r) u(x_2), q).$$

The RDU of the compound lottery (reduced recursively) is

$$V_R(L_c) = [1 - w(q)] u(x_1) + w(q) ([1 - w(r)] u(x_1) + w(r) u(x_2)), \quad (2.166)$$

which can be rewritten as

$$V_R(L_c) = [1 - w(r)w(q)] u(x_1) + w(r)w(q) u(x_2). \quad (2.167)$$

Suppose that the weighting function is subproportional. Thus, the certainty effect (Definition 2.24) implies $w(qr) > w(q)w(r)$. Then, comparing (2.165) and (2.167) we get that $V_R(L_s) > V_R(L_c)$, so the decision maker tolerates less risk in the compound case. Thus, although the simple and compound lotteries are equivalent in the classical sense, the timing of the resolution of uncertainty does matter for a decision maker who follows RDU, if the weighting function is subproportional. As the number of stages in the compound lottery increase, the decision maker exhibits less and less risk tolerance.

Figure 2.22 shows the estimated Prelec weighting functions. Panel (a) shows the atemporal probability weighting function when there is one stage, i.e., a simple lottery ($m = 1$), two stages ($m = 2$), and twelve stages ($m = 12$). Elevation falls as the number of stages increases, so risk




Figure 2.22 Estimated Prelec weighting functions.

Source: Reproduced with permission of *Annual Review of Economics*, Volume 4, Helga Fehr-Duda and Thomas Epper, "Probability and risk: foundations and economic implications of probability-dependent risk preferences," © 2012 by Annual Reviews <http://www.annualreviews.org>

tolerance falls, and for $m = 12$, the probability weighting function becomes convex; in Section 2.3, we classified such a decision maker as exhibiting pessimism.

In panel (b) of Figure 2.22 we assume that the uncertainty in compound lotteries is realized in 12 months time. Comparing the corresponding curves in the two panels, we can compare the effect of time delay. With time delay, each of the curves is more elevated than the corresponding curve in the atemporal case, hence, they exhibit greater risk tolerance.

As Fehr-Duda and Epper (2012, pp. 585–86) put it: "Extreme risk aversion can thus result from several different underlying mechanisms: loss aversion, disappointment aversion, and rank dependence with convex probability weighting or frequent compounding of originally inverse S-shaped probability weights."

CHAPTER 3

Applications of Behavioral Decision Theory

3.1 Introduction

Prospect theory (PT) has been successful in explaining a wide range of problems in almost all areas of economics. In this chapter we consider several applications of PT. Many applications rely on a subset of the components of PT (reference dependence, loss aversion, concavity of utility in gains and convexity in losses, and non-linear probability weighting). The vast majority of the applications invoke reference dependence and loss aversion, but omit non-linear weighting of probabilities. While this explains much human behavior, which is a testament to the importance of these features of human behavior, sometimes this is undesirable, for two reasons. First, the calibration of such models to the data might not be very successful, particularly if one arbitrarily picks and chooses among the components of PT. Second, while the evidence supports the fourfold classification of risk under PT, the omission of non-linear probability weighting can induce, at best, twofold risk aversion, which may reduce the scope of the models' predictions.¹ There is a strong case for giving more attention to non-linear probability weighting in applications of PT, than has traditionally been the case (Fehr-Duda and Epper, 2012).

The general psychological principles behind reference dependence and loss aversion have excellent foundations in experimental and field data.² An important test of these principles has come from an examination of the disparity between *willingness to accept*, WTA, and *willingness to pay*, WTP, for the same good. This disparity is sometimes known as the finding of *exchange asymmetries*. Since this finding arises from mere physical ownership of an object, it is also known as the *endowment effect*.

In Section 3.2 we begin by describing the early and pioneering evidence for exchange asymmetries (Kahneman et al., 1990). The endowment effect arises for physical goods but typically not for money, intrinsically worthless items, or items held for exchange only. This finding, which has been replicated in a large number of experiments, survives controls for transaction costs of exchange, subject misconceptions, learning, and strategic concerns.

¹ For a classic example, see Wakker (2003).

² For experimental evidence on loss aversion, see, for instance, Kahneman and Tversky (1979), Tversky and Kahneman (1992), Thaler et al. (1997), Gneezy and Potters (1997), and Kahneman and Tversky (2000). For the field evidence, see, for instance, Camerer et al. (1997), Genesove and Mayer (2001), Fehr and Goette (2007), Odean (1998), and Mas (2006).

Section 3.2.1 considers the leading explanation of exchange asymmetries, loss aversion (Thaler, 1980). Parting with an object, when one is endowed with it, is coded as a loss by an individual. Hence, loss aversion applies to the act of parting with an object. But loss aversion does not apply to money, hence, the act of parting with the money to buy an object is not subject to loss aversion. In conjunction, this gives rise to the finding of $WTP < WTA$. Indeed, in a PT framework, it can be shown that loss aversion equals the ratio WTA/WTP . Kahneman et al. (1990) report an average value of $WTA/WTP \cong 2$, which is close to the typical values of loss aversion found in many experiments. One study reports the median estimate of WTA/WTP across 45 studies to be 2.6, although there is individual-level heterogeneity in estimates of loss aversion (Horowitz and McConnell, 2002). We also discuss, briefly, the dependence of loss aversion on a range of demographic, goods-specific, context-specific and emotional factors.

Typically, experiments on exchange asymmetries are conducted in riskless environments. In Section 3.2.2, we consider exchange asymmetries in risky environments (Novemsky and Kahneman, 2005), and pose a fundamental question: *Is there any risk aversion beyond loss aversion?* Loss aversion is found to be present in both conditions used to elicit exchange disparities: the risky and the riskless conditions. If there is any risk aversion beyond loss aversion, then one would expect, for instance, that WTA should be relatively higher in the risky situation. But this turns out not to be the case. Hence, there is no risk aversion beyond loss aversion.

Section 3.2.3 considers the effect of experience on the endowment effect. Some evidence suggests a slight reduction in the endowment effect from markets in sports cards memorabilia (List, 2003, 2004). However, contrary to common perception, this effect was found to be similar for lab subjects and market participants, so the conclusion that markets may reduce behavioral biases is not correct. Furthermore, it is not surprising that one may find lower loss aversion among people who hold goods from trade because they are interested in making a profit from the goods and do not necessarily derive satisfaction from owning the goods; this has been known since Kahneman et al. (1990). A comprehensive study based on more than 2.5 million data points from professional sports players shows that even experienced players exhibit loss aversion (Pope and Schweitzer, 2011). This ties in nicely with evolutionary explanations, which suggest that loss aversion may be hardwired in humans; see Section 3.3, below.

Section 3.2.4 considers the possibility that subject misinterpretation and the instructions themselves may be responsible for the endowment effect (Plott and Zeiler, 2005, 2007). However, using the altered instructions, and accounting for subject misperceptions in the suggested manner, it has not been possible to replicate these findings for the case of lotteries (Isoni et al., 2011), although they can be replicated for mugs. This suggests that it is not subject misperception that is the cause of the endowment effect but the particular context of the experiments.

Section 3.2.5 considers some implications and other explanations of exchange asymmetries. Indifference curves may cross over, depending on initial endowments, and the Coase theorem may be violated (Kahneman et al., 1990). We also consider some marketing implications of the endowment effect (Novemsky and Kahneman, 2005). *Stochastic reference points* may be a potential explanation of exchange asymmetries (Köszegi and Rabin, 2006); this is dealt with in more detail in Section 2.8.3. A potential explanation of exchange asymmetries is the *query theory of value construction* (Johnson et al., 2007), which is similar in spirit to the availability heuristic (see Part 7). This is a promising line of work, which shows that the nature of queries that sellers and buyers of objects pose may influence the extent of exchange asymmetries. Yet another possible explanation that we consider is the role of uncertainty in explaining the endowment effect (Engleman and Hollard, 2008).

Non-human primates share a common evolutionary history with humans. If they were to exhibit the endowment effect or PT preferences, then this might be the result of an evolutionary process that precedes the separation of chimpanzees and humans from a common ancestor. We consider this fascinating question in Section 3.3 (Chen et al., 2006; Brosnan et al., 2007; Lakshminarayanan et al., 2008, 2011). It turns out that capuchin monkeys have consistent preferences and they exhibit many of the features of PT that include reference dependence, loss aversion, and different risk attitudes in the domain of gains and losses. Chimpanzees exhibit an endowment effect for food items, but not for non-food items where they exhibit the reverse of the endowment effect (a preference for interacting with the experimenter makes them trade too much). This difference arises possibly because food is more salient for evolutionary success as compared to non-food items.

Section 3.4 considers the tendency of humans to evaluate gains and losses only over short horizons, even when the relevant horizons are longer. For instance, even if the returns on a portfolio are to materialize in two years, individuals may repeatedly check the market returns. Insofar as negative returns may lead to loss aversion, risky assets appear more aversive. Such a phenomenon is known as *myopic loss aversion*. This explains the *equity premium puzzle*, i.e., why the return differential on risky and riskless assets cannot be explained by any reasonable attitudes to risk (Benartzi and Thaler, 1995). Myopic loss aversion is supported in experimental data on individual subjects (Gneezy and Potters, 1997), team experimental data (Sutter, 2007), and in artefactual field data (Haigh and List, 2005).

Section 3.5 considers the only application in Chapter 3 that uses all the components of PT in order to address the *tax evasion puzzles*. The classic analysis of tax evasion, based on EU (Allingham and Sandmo, 1971), gives rise to qualitative and quantitative puzzles. The quantitative puzzle is that given the return from tax evasion for an amateur tax evader of 91%–98%, why should anyone bother to pay taxes?³ The qualitative puzzle arises because, under the reasonable assumption of decreasing absolute risk aversion, EU predicts that evasion should fall as the tax rate goes up (Yitzhaki, 1974). However, the bulk of the evidence indicates that people attempt to evade more if the government tries to tax away more of their incomes.

We show that PT is able to address the quantitative and qualitative tax evasion puzzles (Dhami and al-Nowaihi, 2007, 2010b). Although the probability of an audit and the penalty rates for tax evaders are low, loss aversion from being caught, and overweighting of small probabilities of detection, ensure that only about a third of those who can evade, do evade. Using parameters of risk from generic situations of risk under PT, we can explain well the evidence from tax evasion. This provides independent confirmation of PT in a two-horse race with EU. By contrast, the predictions of EU are wrong up to a factor of up to 100.

We considered Rabin's paradox for EU in Section 1.5.4 in Chapter 1. Section 3.6 shows that this paradox can be resolved by using loss aversion (Rabin, 2000a). Several alternative explanations for Rabin's paradox turn out to rely on reference dependence, which is central to PT (Wakker, 2005, 2010).

Under EU, the carriers of utility are final wealth levels, so *goals* or *targets* should not alter behavior. However, the evidence suggests that goals and targets do influence behavior and outcomes. We begin, in Section 3.7, by considering the role that *goals* play as reference points and propose a simple theoretical model of goals under PT, using unpublished notes by al-Nowaihi and Dhami.

³ The probability, p , of audit for an amateur tax evader is 1–5 % and, if caught, the taxpayer has to pay back the evaded tax, plus a fine, θ , at the rate of 0.5 to 2 times the evaded tax. Thus, the return on one dollar of evaded tax is $(1 - p) - p\theta$.

This model is able to address some of the stylized facts in the literature (Heath et al., 1999) that show the following. First, goals partition the space of outcomes into gains and losses, thus, they serve as reference points. Second, individuals often work harder when they have a goal. Third, individuals work even harder as they get closer to a goal. Finally, we make some preliminary progress in the direction of determining *optimal goals*.

Section 3.8 provides an application of PT to labor economics, as encapsulated in the following question. Why is it hard to find a taxi on a rainy day in New York (Camerer et al., 1997)? The behavioral explanation is that taxi drivers have a target daily level of earnings (reference dependence), which is possibly a response to issues of *self-control*, particularly for drivers who are not highly experienced (see Part 3). If they quit before achieving their target, then they suffer loss aversion. On a rainy day, they are more likely to achieve their target early because the average waiting times between successive jobs is low. Hence, they are more likely to quit early. Thus, the labor supply curve is backward bending, although it is less pronounced for more experienced drivers. We review the relevant empirical evidence (e.g., Fehr and Goette, 2007; Farber, 2005, 2008). We also sketch out models of labor supply with a single target (earnings) and the more recent work in which there are multiple targets (earnings and hours worked) (Crawford and Meng, 2011).

One implication of the inverse S-shaped probability weighting functions is that the decision maker places higher decision weight on more extreme outcomes (Fehr-Duda and Epper, 2012). In Section 3.9, we briefly consider applications of this idea, mainly to issues in finance. For additional applications to finance, see the chapter on behavioral finance in Part 7 of the book.

In Section 3.10, we consider sequential interaction between contracting parties who enter into a contract at date 0, and then engage in some economic transaction at date 1. At date 0, there is a competitive pool of contracting parties from which a pair (say, a buyer and a seller) are matched. The critical idea is that the date 0 contract may serve to define the *entitlements* of the two parties and, hence, serve as a reference point for decisions at date 1 (Hart and Moore, 2008). If the date 1 behavior of parties does not conform to their expected entitlements, then parties may shade their performance, and reduce the joint surplus. We consider the effect of these factors on the choice of optimal contracts. The empirical evidence is consistent with the model (Fehr et al., 2011).

Section 3.11 considers the effect of loss aversion on optimal contracts in the presence of moral hazard, in the standard principal–agent problem (De Meza and Webb, 2007). It turns out that loss aversion reduces the power of incentive schemes. If incentives schemes are too high powered, then loss aversion makes the satisfaction of the agent’s individual rationality constraint too expensive for the principal. Thus, it is optimal for the principal to offer relatively flat segments in the optimal wage schedule. We then examine heuristically how loss aversion may explain other observed contractual phenomena (Herweg et al., 2010).

Section 3.12 considers the issue of renegotiation of contracts under reference dependence and loss aversion (Herweg and Schmidt, 2013). The two main findings are that under loss aversion, renegotiation is sluggish and does not take place in all states of the world, although it would have, in the absence of loss aversion. The analysis also offers insights into the choice between long-term contracts and the use of spot markets.

This chapter explores only a very small set of the applications that have successfully applied PT. Some of the other successful applications include: prediction of complex states of health (Abellan-Perpiñan et al., 2009), life-cycle consumption savings models (Bowman et al., 1999), racetrack betting (Jullien and Selanié, 2000), crime and punishment (Dharm and al-Nowaihi, 2010a, 2014), asset market phenomena (Kliger and Levy, 2009; Gurevich et al., 2009; Barberis et al., 2001), the disposition effect (Shefrin and Statman, 1985), and asymmetric price elasticities (Camerer, 2000).

3.2 The endowment effect and exchange asymmetries

The idea behind exchange asymmetries is that ownership confers a psychic benefit, or an *endowment effect*. Kahneman et al. (1991) motivate the endowment effect through the following example. “A wine-loving economist we know purchased some nice Bordeaux wines years ago at low prices. The wines have greatly appreciated in value, so that a bottle that cost only \$10 when purchased would now fetch \$200 at auction. This economist now drinks some of this wine occasionally, but would neither be willing to sell the wine at the auction price nor buy an additional bottle at that price.” Neoclassical economics predicts that, depending on his valuation $v \gtrless 200$, either the seller would like to sell at a price of \$200 or prefer to buy it.

Let us term the amount of money that *owners* of an object are willing to accept in exchange for the object as the *willingness to accept* (WTA). The *willingness to pay* (WTP) for the object is the amount of money that individuals are willing to pay to buy an extra unit of the object. Thaler (1980) noted the presence of *exchange asymmetries*, i.e., $WTP < WTA$; a phenomenon also known as the *endowment effect*. The effects of ownership need not be immediate, but could be gradual, increasing over time as the duration of ownership increases (Strahilevitz and Loewenstein, 1998). Ariely and Simonson (2003) also demonstrate a similar pattern for the case of bidding in auctions. The highest bidder at any stage (who has not won the auction yet) become partially attached to the object (a pseudo-endowment effect).

Thaler (1980) invoked *loss aversion* to explain the endowment effect. The act of giving up the object is coded as a *loss* by the owner, so *loss aversion* applies to sales of items. However, the endowment effect has been found to apply to goods but not to money or to goods purchased for resale (Kahneman et al., 1990).⁴ Hence, loss aversion does not apply to the buyer's act of giving up cash, thus, $WTP < WTA$.

The endowment effect need not be limited to goods. It can also arise in the case of *perceived entitlements*, say, to a previously prevailing price. So, for instance, houses stay on the market longer in a period of falling prices. Sellers might be reluctant to sell at a lower price because they perceive that they are entitled to a previously prevailing, higher, price. There is also some evidence that the volume of stocks traded is lower for those stocks whose price is falling relative to those whose price is increasing (Shefrin and Statman, 1985).

We now describe a generic experiment in Kahneman et al. (1990), designed to test for the *endowment effect*. In order to control for factors such as subject misperception and transaction costs in potentially interfering with, or even causing, the endowment effect, Kahneman et al. (1990) first ran three markets in tokens. Tokens are intrinsically worthless and (in the absence of factors such as misperceptions and transactions costs) are not expected to induce an endowment effect, so $WTP = WTA$. The tokens can be redeemed at a pre-announced exchange rate from the experimenter.

We reproduce the instructions for sellers and buyers in some detail because these have played an important part in the debate on the explanation of the endowment effect. In the market for tokens, sellers receive the following instructions.

In this market, the objects being traded are tokens. You are an owner, so you now own a token which has a value to you of \$x. It has this value to you because the experimenter will give you this much money for it. The value of the token is different for different individuals. A price for the tokens will be

⁴ Using data from the TV show *Deal or no Deal*, Blavatsky and Pogrebna (2009) raise the possibility that the endowment effect might be weakened when the stakes are very high.

determined later. For each of the prices listed below, please indicate whether you prefer to: (1) Sell your token at this price and receive the market price. (2) Keep your token and cash it in for the sum of money indicated above. For each price indicate your decision by marking an X in the appropriate column.

The instructions for the buyers were as follows.

In this market, the objects being traded are tokens. You are a buyer, so you have an opportunity to buy a token which has a value to you of \$x. It has this value to you because the experimenter will give you this much money for it. The value of the token is different for different individuals. A price for the tokens will be determined later. For each of the prices listed below, please indicate whether you prefer to: (1) Buy a token at this price and cash it in for the sum of money indicated above. (2) Not buy a token at this price. For each price indicate your decision by marking an X in the appropriate column.

This is followed by a list of prices, ranging from \$0.25 to \$8.75, increased in steps of \$0.50. Sellers are asked to state their WTA, while buyers are asked to state their WTP. Subjects were made to alternate in the roles of buyers and sellers in three successive markets. The tokens were assigned a different individual redemption value in each trial. Once the subjects give their responses, the experimenter immediately computes the implied demand and supply schedules for each market as well as the market clearing price.

Three buyers and sellers are then selected and paid off in accordance with the market clearing price. The procedure described above is incentive compatible. The main finding was $WTP = WTA$. Hence, based on the trading in the market for tokens, one can rule out the presence of factors that might vitiate the endowment effect. These factors include subject misperception (e.g., confusing instructions), bargaining postures (sellers asking for more to strengthen their bargaining positions), and transactions costs of trade.

Kahneman et al. (1990) then gave Cornell coffee mugs to subjects on alternative seats (these mugs could be purchased from the University bookshop for \$6.00 each). Subjects were given an opportunity to examine the mugs. Then, the experimental procedure for tokens was applied by replacing mugs with tokens; all subjects found the mugs to be desirable items because they placed a positive valuation on them. Subjects were told that it was in their best interests to answer questions truthfully and that of the four mug markets conducted, one market, *the binding market trial*, would be selected at random and played for real. The initial assignment of subjects in the roles of buyers or sellers was maintained throughout. Following the four markets for mugs, four more markets were conducted for boxed ballpoint pens, which sold for \$3.98 in the University bookstore (the price tags were left visible to the subjects). This procedure is also incentive compatible.

The repetition of each of the markets for four times was designed to allow learning to take place. However, no learning effects were found in the experimental results (we comment more on these results below).

The results in the mugs and pens markets were very different from the results arising with tokens. In the markets for objects, Kahneman et al. (1990) found that the median selling price was more than twice the median buying price. The actual volumes of trade were only respectively 20% and 41%, respectively, of the predicted volumes.

Could sellers simply be inflating their WTA in order to get a higher equilibrium price? Kahneman et al. (1990) control for this objection by running another experiment in which the price is chosen randomly. Subjects are explicitly told that their decision will have no effect on the

actual price chosen. From Becker et al. (1964), it is known that this elicitation mechanism ensures truthful reporting of valuations and willingness to pay for the good. Since strategic postures on the part of sellers and buyers have no bearing on the actual price in this case, it is always at least a weakly dominant strategy to tell the truth. The experimental results still showed a large endowment effect and were nearly identical to the ones obtained when actual market clearing prices were used.

Experiments 6 and 7 in Kahneman et al. (1990) added an extra group of experimental subjects. Sellers in the experiment were given a mug and asked if they would sell the mug at prices ranging from \$0 to \$9.25. The *buyers* were asked if they would buy the mug at these prices. The third subset of subjects, the *choosers*, were asked at each price, say \$ x , if they would prefer to have a cash amount of \$ x , or the mug (which is exactly the same choice as faced by a seller). The only difference between choosers and sellers is that the latter have physical possession of the good.

If the behavior of the choosers and the sellers is similar, then the endowment effect does not have any bite, because, in this case, it is not physical possession of the good, per se, that creates a higher WTA. In contrast, it was found that the behavior of choosers was similar to that of buyers, but not the sellers, which confirms the endowment effect. Median valuations of the mug for sellers, choosers, and buyers, were \$7.12, \$3.12, and \$2.87, respectively.

By way of comparison, the earlier study of Knetsch (1989) used two objects (mugs and candy bars). Each object was, in turn, randomly allocated to some subjects (sellers) who had an opportunity to trade it for the other object. When endowed with the mug, 89% of the sellers preferred to keep the mug. When endowed with the candy bar, 90% of the sellers preferred to keep the candy bar. Since subjects are allocated randomly to the two treatments (i.e., mug ownership and candy ownership), in the absence of an endowment effect it is difficult to explain these results. An even earlier study by Knetsch and Sinden (1984) found the WTA to WTP ratio for lottery tickets to be about 4.⁵

The evidence for the endowment effect outlined in this section relied mainly on items such as mugs and pens. However, the endowment effect has also been documented for a range of other objects such as college basketball tickets (Carmon and Ariely, 2000), gift certificates (Sen and Johnson, 1997), choice of pizza toppings (Levin et al., 2002), hunting permits (Cummings et al., 1986), wine (Van Dijk and Van Knippenberg, 1996), and car attributes (Johnson et al., 2006). Kahneman et al. (1990) argue that the endowment effect is also likely to hold when the entities engaged in a transaction are firms rather than individuals. They conjecture that once firms obtain property rights over some economic asset (by accident of history, luck, or public policy) they might be reluctant to sell the asset, while at the same time being reluctant to buy such an asset at the market price.

The endowment effect suggests that individuals are relatively more reluctant to give up objects they own, so they appear to have a bias towards the current allocation, i.e., a *status-quo bias*. Samuelson and Zeckhauser (1988) gave subjects a choice between investing in one of four possible investments: moderate risk, high risk, treasury bills, and municipal bonds. In the neutral treatment, the endowment of the subjects is just money. In the non-neutral treatments, the endowment is one of the possible investments themselves. A status-quo bias (a larger fraction of the final portfolio consisting of the investment one is endowed with) is found in the non-neutral treatment relative to the neutral treatment. Furthermore, the status-quo bias is found to be more pronounced when there is a larger number of alternatives.

⁵ Perhaps one of the earliest studies to report a WTA-WTP gap was by Coombs et al. (1967). However, the experiment involved a hypothetical setting of lottery tickets without real exchange.

3.2.1 Loss aversion and the endowment effect

To understand exchange asymmetries in terms of loss aversion, consider the following piecewise-linear Köszegi–Rabin type utility function,

$$v(x) = \begin{cases} m + (x - r) & \text{if } x \geq r \\ m + \lambda(x - r) & \text{if } x < r \end{cases}, \quad (3.1)$$

where x is the actual value of the consumption of the good, r is some “reference point” for this value, m is the amount of money received (this is negative if money is paid out), and λ is the coefficient of loss aversion.⁶

In most exchange experiments, sellers are endowed with one unit of an object, say, a mug, hence, assume that the reference point of the seller is $r = 1$. In (3.1), the separability in goods and money is motivated by experimental results (see above), which suggest that there is no endowment effect for money. Suppose that the seller is offered a price p for the mug. Using (3.1), and noting that $m = p$, if he does not sell the mug, then his utility is $v(x = 1) = 0$. If he sells the mug, his utility is $v(x = 0) = p - \lambda$. Hence, he sells the mug if and only if

$$p - \lambda \geq 0. \quad (3.2)$$

Notice that the coefficient of loss aversion, λ , plays a critical role in the decision to sell the mug. Higher loss aversion impedes sales, unless the price is high enough.

Definition 3.58 (WTA): The WTA is the minimum price, $p = p_s$, that the seller is willing to accept in order to sell the mug. From (3.2)

$$p_s = \lambda. \quad (3.3)$$

Consider now the buyer who is not endowed with the mug, so his reference point is $r = 0$. Using (3.1), and noting that at any price that he pays, p , we have $m = -p$, if the buyer does not buy the mug, then his utility is $v(x = 0) = 0$. If he buys the mug, then, his utility is $v(x = 1) = 1 - p$. Hence, the buyer buys the mug if and only if

$$1 - p \geq 0. \quad (3.4)$$

Definition 3.59 (WTP): The WTP is the maximum price, $p = p_b$, that the buyer is willing to pay for the mug. From (3.4),

$$p_b = 1. \quad (3.5)$$

From (3.3), (3.5), we get that:

$$\frac{WTA}{WTP} = \frac{p_s}{p_b} = \lambda. \quad (3.6)$$

Hence, the ratio of WTA to WTP is simply the coefficient of loss aversion. In Kahneman et al. (1990), $WTA = 7.12$ and $WTP = 2.87$, so we get that $\lambda \approx 2.5$. Across 45 studies surveyed in Horowitz and McConnell (2002), the median estimate of WTA/WTP is 2.6.

⁶ Recall the first two lines of instructions above: “You are an owner, so you now own a token which has a value to you of \$ x . It has this value to you because the experimenter will give you this much money for it.” This is the value of x in (3.1).

What sort of factors does the endowment effect depend on? Insofar as the endowment effect is caused by loss aversion alone, one can equivalently ask, what factors does loss aversion depend on? Can the endowment effect be switched on and off, or be reduced/amplified? While more work remains to be done in this area, the literature already suggests several factors.

The endowment effect is reduced (equivalently loss aversion is reduced), when (1) the owned good and the unowned good are close substitutes (Chapman, 1998; Van Dijk and Van Knippenberg, 1996), (2) the duration of ownership is shorter (Strahilevitz and Loewenstein, 1998), (3) questions are framed in a way that directs sellers to the uses that money from sales can be put to and buyers to the benefits arising from ownership of the object (Carmon and Ariely, 2000), (4) a lesser proportion of money is earmarked for necessities (Wicker et al., 1995), (5) negative emotions such as disgust are invoked, prior to the elicitation of WTA and WTP,⁷ (6) subjects are relatively older or more educated individuals, and have greater knowledge of the attributes of a product (Johnson et al., 2006); the youngest subject had a loss aversion of 1.4, while the oldest subject had a loss aversion of 2.4, (7) there is a reduction in the ambiguity about the value of the good, and a reduction in the cost of gathering information (Kolstad and Guzman, 1999); explicit price information can reduce the ambiguity about value and reduce the WTA/WTP disparity (Zhao and Kling, 2001).

There is evidence of variability of loss aversion for an individual across different traits of a good or for different goods. Factors that could influence loss aversion for an individual include the following: whether the good is a health or a non-health good with greater disparity for health related goods (Brown and Gregory, 1999; Sayman and Oncüler, 2005), whether the good is an environmental or a non-environmental good (Irwin, 1994), whether trade in the good is legitimate or not (Sayman and Oncüler, 2005), and whether the good has hedonic or utilitarian attributes (Dhar and Wertenbroch, 2000).

To get a feel for how loss aversion might vary with the attributes to a product, consider the empirical study by Johnson et al. (2006). They conducted a survey of 360 German car buyers who had recently purchased a mid-sized car. Four different attributes of the car were considered. The aim was to measure loss aversion (using appropriate lotteries) for each individual and for each car attribute. Across all individuals, and attributes, the average loss aversion was 1.85. On average, for any individual, loss aversion for the four attributes was not correlated. There was also significant heterogeneity in measured loss aversion across individuals. The average coefficients of loss aversion across the four attributes were 1.66 (fuel consumption), 1.89 (comfort), 1.89 (safety), and 1.94 (information). One possible explanation is that people differ in their perceptions of the various attributes and loss aversion is context dependent. This supports a great deal of research elsewhere on the context dependence of preferences.

It is important to note that the same subjects who exhibit loss aversion in this study had recently engaged in a major purchase. This is a noteworthy point in the light of other studies (considered below) which push the alternative view that if subjects are sufficiently well trained, experienced, and briefed about strategies and equilibrium notions, then they might completely shed loss aversion. Johnson et al. (2006) conclude with the following. “Thus, we suggest that the question should not be whether or not loss aversion is important, but rather how important loss aversion is, and for which attributes and consumers.”

⁷ See, for instance, Lerner et al. (2004). These authors find that when disgust is invoked, the endowment effect gets switched-off. Peters et al. (2003) also point out the effect of emotionally difficult trade-offs on the WTA-WTP disparity. Boyce et al. (1992) find that the WTA-WTP disparity increases when the exchange of goods involves a degree of moral responsibility.

Individual-level heterogeneity is found in another study by Gächter et al. (2010) who compare loss aversion for risky and riskless choice. The correlation in loss aversion between these two choices is 0.635, which is statistically significant at the conventional levels. Averaged across all individuals they get the classic figure of loss aversion of around 2. However, across individuals they found that for 78% of the individuals, the loss aversion parameter is between 1 and 4; for 10 percent of the individuals, it is above 4; and for 22% of individuals, it is around 1.

3.2.2 Is there any risk aversion beyond loss aversion?

Novemsky and Kahneman (2005) extend Thaler's (1980) idea of using loss aversion as an explanation of exchange disparities to risky choice. Their most important conclusion is that *there is no risk aversion beyond loss aversion* when subjects are presented with lotteries that have balanced risks (i.e., equal probability of all outcomes). The idea that the finding of risk aversion could, at least partly, be caused by loss aversion is well recognized.⁸

Definitions 3.58, 3.59, respectively, define WTA and WTP. We now give another useful definition that relies on the concept of "choosers" in Kahneman et al. (1990) that we have defined above.

Definition 3.60 (*Choice equivalent*): The choice equivalent, CE, is defined as the minimum amount of money that "choosers" would accept in preference to the object.

Novemsky and Kahneman (2005) motivate their ideas through a series of hypotheses. Recall that choosers do not actually have possession of the object, and they exhibit no endowment effects. Empirically, they behave like buyers and not sellers. For this reason, they are willing to accept a lower amount of money in preference to the object (CE) relative to the WTA of sellers, hence, we get hypothesis H1, below.

H1: $\frac{WTA}{CE} > 1$ (endowment effect).

A second finding is that for intrinsically useless assets such as money and tokens, one does not find the presence of an endowment effect. In this case $WTA = WTP$, but $WTA = CE$ for choosers, thus, we get H2.

H2: $\frac{CE}{WTP} = 1$ (no loss aversion in parting with money).

In Novemsky and Kahneman (2005), the first three conditions are identical to Kahneman et al. (1990), namely, *sellers*, *buyers*, and *choosers*. In order to take account of risk, two additional conditions are introduced. In the first of these conditions, the *risky seller condition*, sellers are endowed with a mug. In addition, they are asked if they would be willing to accept the following voluntary *balanced risk* lottery with two outcomes. (i) Retain the mug and gain an amount of money, m , with probability 0.5. (ii) Lose the mug and receive nothing with probability 0.5.

Definition 3.61 (*RWTA*): The risky willingness to accept (RWTA) is the minimum sum of money, m_s , that the sellers require to accept the balanced risk gamble in the "risky seller condition."

In the second additional condition, the *risky buying condition*, buyers face the following balanced risk gamble with two outcomes. (i) Receive the object and pay nothing with probability 0.5. (ii) Pay some amount of money, m , and not receive the object with probability 0.5.

⁸ See, for instance, Kahneman and Tversky (2000), Kahneman (2003), List (2004), Marquis and Homer (1996), and Rizzo and Zeckhauser (2003).

Definition 3.62 (RWTP): The risky willingness to pay (RWTP) is the maximum amount of money, m_b , that the buyer will pay to accept the balanced risk gamble in the “risky buying condition.”

We use the classical utility function under PT (and not (3.11)) to calculate the relevant expressions for RWTA and RWTP. We use a probability weighting function $w(p)$ that has the same form for losses and gains. The decision maker’s utility from an outcome is given by (3.1).

In the *risky seller condition*, consider the problem faced by a seller who owns a mug, so $r = 1$. If the seller accepts the balanced risk gamble, then there is one outcome each in the domain of gains and losses, and his utility, v_s , is given by

$$v_s = w(0.5)(1 - 1 + m) + w(0.5)\lambda(0 - 1). \quad (3.7)$$

If the seller does not accept the gamble, then he retains the mug but gets no money, so his utility is $v_s = 1 - 1 = 0$. Hence, the seller accepts the gamble if and only if

$$w(0.5)m - \lambda w(0.5) \geq 0 \Leftrightarrow m - \lambda \geq 0.$$

Using Definition 3.61, the RWTA is given by

$$m_s = \lambda. \quad (3.8)$$

Now consider the *risky buyer condition*. Since the buyer does not own the mug to begin with, so his reference point $r = 0$. If the buyer engages in the balanced risk gamble, there is one outcome each in the domain of gains and losses, and his utility, v_b , is given by

$$v_b = w(0.5)(1 - 0 - 0) + w(0.5)\lambda(0 - 0 - m). \quad (3.9)$$

If the buyer does not engage in the gamble, then his utility is $v_b = 0$. Hence, in the *risky buyer condition*, the buyer accepts the gamble if and only if

$$w(0.5) - w(0.5)\lambda m \geq 0 \Leftrightarrow 1 - \lambda m \geq 0. \quad (3.10)$$

Using Definition 3.62 and (3.10), we get that the RWTP is given by

$$m_b = \frac{1}{\lambda} < 1. \quad (3.11)$$

(3.8) and (3.11), in conjunction with Definitions 3.58 and 3.59, allow us to formulate the next two hypotheses.

H3: For balanced risks, $\frac{RWTA}{WTA} \equiv \frac{m_s}{p_s} = 1$ (no risk aversion beyond loss aversion).

Hypothesis H3 simply follows from (3.3), (3.8). Its consequences are profound for economics, namely, that for balanced risks there is no risk aversion beyond loss aversion. Loss aversion is present in the risky and the riskless conditions. If the individual feels any risk aversion beyond loss aversion, then it must be the case that $RWTA > WTA$ (i.e., the seller would require greater compensation for parting with the object in the risky condition). But $RWTA = WTA$. Hypothesis

H3 also gives support to the notion that the parameter of loss aversion is the same in risky and riskless situations.

From (3.5), $p_b = 1$ and from (3.11) $m_b = 1/\lambda$, hence, we get our final hypothesis, H4.

H4: For balanced risks $\frac{WTP}{RWTP} \equiv \frac{p_b}{m_b} > 1$.

Hypothesis H4 assumes that in riskless conditions, buyers do not have loss aversion for the money that they are spending. However, in the risky condition, buyers are gambling away their money. If buyers are loss averse with respect to money under risk, then their risky willingness to pay should be relatively smaller, giving rise to H4.

To recap, there are five conditions: sellers, choosers, buyers, risky sellers, and risky buyers. The objects used are chocolates, pens, and mugs. The elicitation mechanism ensured that there would be no strategic responses. The median response by the appropriate subset of subjects is used to compute WTP, WTA. A bootstrapping method is used to estimate confidence intervals for the various ratios used in H1–H4. The aggregate estimates of these ratios were as follows,

$$\frac{WTA}{CE} = 1.85; \frac{CE}{WTP} = 1.07; \frac{RWTA}{WTA} = 0.91; \frac{WTP}{RWTP} = 2.31, \quad (3.12)$$

which supports all of the hypotheses H1–H4. The experimental results can thus be summarized as follows:

1. Because $\frac{WTA}{CE} > 1$ and $\frac{CE}{WTP} \cong 1$, as in earlier studies, there is no loss aversion for money in riskless situations.⁹ The 95% confidence interval for $\frac{WTA}{CE}$ included the number 2 (a commonly found magnitude for the parameter of loss aversion, λ) in five out of seven experiments.
2. Since $\frac{WTP}{RWTP} > 1$, there is loss aversion for money given up in risky situations. The ratio $\frac{WTP}{RWTP}$, however, turned out to be quite variable across experiments so there is need for more work in this area.
3. Since $\frac{RWTA}{WTA} \cong 1$ (the 95% confidence interval included the number 1 for all experiments), *there is no risk aversion beyond loss aversion*. This is a profound result for economics; indeed, economists who are detecting risk aversion in their empirical studies, might actually be picking out loss aversion.¹⁰

Schunk and Winter (2009) provide more indirect support for these results. They look at problems which require the computation of optimal stopping times. Such models typically postulate risk-neutral decision makers. The observed stopping times fall short of the predictions under standard theory with risk-neutral decision makers. Schunk and Winter (2009) find that the optimal stopping times are based on search heuristics. Furthermore, these search heuristics are not correlated with risk attitudes but with attitudes towards loss aversion.

3.2.3 Market experience and loss aversion

In their experiments, Kahneman et al. (1990) (see above) conducted two sets of market trials, respectively, with mugs and pens. Each set of trials consisted of four successive markets. At the

⁹ Unlike most other studies, Bateman et al. (1997) report loss aversion of buyers for money in riskless situations; they find that $\frac{CE}{WTP} = 1.67$, which is statistically different from 1. They also find little loss aversion in selling. Novemsky and Kahneman (2005) conjecture that the source of the difference lies in different subject pools (US versus UK students) as well as the perception of the UK subjects that they did not have budget reserves (hence leading to loss aversion for money).

¹⁰ An important condition for this result is that income effects must be small. Hence, for large stakes, this result might or might not hold.

end of each market trial, full feedback was provided to all the participants in order to facilitate learning. They found that as the market trials progressed (from 1 to 4) there was no evidence of learning, and the disparity between WTP and WTA remained. In contrast, some earlier studies report that market experience in practice trials (typically hypothetical and non-binding trials) significantly reduces the WTP/WTA disparity.¹¹ Kahneman et al. (1990) argue that the source of the difference between the two sets of results arises from the fact that each of their market trials was potentially binding, which is not the case with earlier empirical studies.

One may wonder if the time span required for learning is longer than what Kahneman et al. (1990) allow for in their experiments? List (2003) seeks to address these issues using field data. In his first market, he used two sports memorabilia items, A and B.¹² A distinction is made between dealers and non-dealers in the sports memorabilia market. Dealers have greater experience of trades as compared to non-dealers. Both dealers and non-dealers could be buying and/or selling objects. The pooled sample contained both dealers and non-dealers. There were also separate samples of dealers and non-dealers.

Dealers/non-dealers were given possession of one of the two objects. They were then given the opportunity to trade the object for the other object. Both items are somewhat non-standard in the sports memorabilia market. However, a pre-test suggested that the preference between the two items is nearly equally split, so we should expect trade to take place about 50% of the time. In the exit interviews, 95% indicated that they would like to keep the object for their own collection, rather than intending to use it for resale.

In the pooled sample of dealers and non-dealers, there was too little trading (i.e. trades below 50%) as predicted by the endowment effect. For items A and B, the actual trades were 32.8% and 34.6%, respectively, which suggest that $WTA > WTP$. Despite the random allocation of objects, a Fisher exact test was used to control for the problem that more people than the average could have received their most preferred item (and so do not need to trade). The test rejected the null hypothesis of no-endowment effect.

Splitting the pooled sample into dealers and non-dealers reveals a different picture. Dealers trade objects A, B, respectively, 43.6 and 45.7% of the time. For non-dealers, these percentages lie between 20–25. The Fisher test allows rejection of the null hypothesis (no-endowment effect) for non-dealers. However, the null hypothesis is accepted for dealers, i.e., dealers do not exhibit an endowment effect. When non-dealers are differentiated between experienced (who trade six or more items in a month) and inexperienced ones, it is found that for the latter category, the endowment effect is stronger. The nature of the results is similar when they are replicated with a different set of objects. A logit regression confirms that the propensity to trade is positively affected by trading experience, i.e., experience seems to mitigate the endowment effect.

An important issue is whether *selection effects* have been properly controlled for. Do dealers trade more because they are more experienced (treatment effect) or is it because the type of people who have lower loss aversion (and, so, a lower endowment effect) become dealers and trade more (selection effect)? To test for this effect, List followed up, a year later, the 148 people who

¹¹ See, for instance, Knez et al (1985), Brookshire and Coursey (1987), and Coursey et al. (1987). Loomes et al. (2003) also find a disappearance of the systematic part of the WTA–WTP disparity with experience, although they cannot pinpoint the mechanism that leads to this disappearance. Shogren et al. (1994) find that exchange asymmetries disappear with repeated experimental interaction when there are close substitutes for market goods. But this result does not hold for non-market goods with imperfect substitutes. However, Morrison (1997) raises doubts about these results.

¹² Item A is a ticket stub in which Cal Ripken Jr. broke the world record for consecutive baseball games played. Good B was a dated certificate to commemorate the game in which Nolan Ryan achieved the distinction of winning 300 baseball games.

participated in the earlier sports memorabilia experiment in 1999. Of these, he was able to meet 72 (less than half the original sample). The idea was that the subjects have had an extra year's experience and so would perhaps have shed their endowment effect even further. With this restricted set of subjects, two new A and B objects were introduced. The result was that the subjects, on average, engaged in slightly more trades in the year 2000 compared to the year 1999 (6.84 versus 5.66). List (2003) concludes that this demonstrates a reduction of the endowment effect with experience.

List (2004) seeks to address one possible criticism of List (2003). This is to take account of the possibility that no endowment effect should be observed if sellers hold a good for resale (e.g., Kahneman et al., 1990). Whilst the exit interviews in List (2003) suggested that subjects intended to keep the objects for personal consumption, List (2004) reverts back to the more traditional items used to test the endowment effect, such as mugs and Swiss chocolate bars. The expectation is that these items will not be used for resale. The other feature of List (2004) that is different from List (2003) is that experience in selling/buying one sort of item is expected to be transferred to selling/buying other objects (mugs, chocolates). The findings in List (2004) are similar to those of List (2003).

In conjunction, the findings in List (2003, 2004) have generated substantial interest. The skeptical view is that these findings have served to reject PT. A more balanced view is that the real contribution of List (2003) is to demonstrate that there are differences in the trading propensities of dealers and non-dealers. List (2003) is not alone in highlighting the role of experience in potentially affecting loss aversion. For instance, Camerer et al. (1997) find differences in loss aversion for experienced and inexperienced New York cab drivers. Similar evidence from the housing market comes from Genesove and Mayer (2001). However, we would urge caution in interpreting the List (2003) findings as a rejection of PT. We elaborate below.

1. Considerable experience is required in the market experiments of List for the endowment effect to diminish.¹³ However, in a range of plausible and non-trivial real-world situations such as selling a car, a house or a consumer durable, and making decisions such as marriage and divorce, one typically does not have the benefit of learning from a large number of repetitions. In these cases, one is likely to exhibit loss aversion, like non-dealers in the List experiments.
2. There are other possible interpretations of the results in List (2003). In 1999, his sample had 148 individuals, while the latter sample in 2000 had 72. It is quite likely that the average trades of the two groups could differ by a relatively small amount (6.84 versus 5.66, as it turns out) for a variety of other reasons, e.g., a more pessimistic trading environment in 2000 as compared to 1999. Further, the 72 people out of 148 is not a random sample. It is possible that the 72 who agreed to participate in the experiment, a second time, are the more committed traders who regularly go to the annual sports trade shows. Hence, there is a selection issue involved. Others are even more skeptical about the attempts to control for selection issues in the context of List's work. For instance, Plott and Zeiler (2007, p. 1451) write: "One can formulate other candidate theories, based on various features of List's experiments, to explain his observation . . . Conducting experiments in the field makes it difficult, if not impossible to control for selection effects."

¹³ In List (2003), experienced dealers carry out 6 or more trades per month, while in List (2004) the requirement increases to 11 or more trades per month. Furthermore, such intense trading is carried out by these dealers over several years.

3. Several other studies that we review below suggest that in other important contexts, with higher stakes, and much larger datasets (in one case, more than 2.5 million data points), experienced subjects are also loss averse.
4. List (2003) also finds a reduction in the endowment effect for students who participate in a WTA/WTP experiment for four successive weeks. Students exhibit experience effects that are similar to these effects for dealers in the field. In the first week, 12% of the students trade their endowed good, but 26% do so in the fourth week, suggesting that experience may mitigate the endowment effect. In the light of these experiments, the title of List's paper "Does market experience eliminate market anomalies?" is misplaced. Going by the results of the paper, market experience appears to reduce loss aversion both in the lab and in the field. Yet the misleading inference that only market experience in the field is somehow effective is routinely drawn. Consider, for instance, the following quote from Haigh and List (2005, p. 525): "In light of some recent studies (e.g., List, 2002, 2003, 2004) that report market anomalies in the realm of riskless decision-making are attenuated among real economic players who have intense market experience, the current lot of experimental studies with MLA [myopic loss aversion] may be viewed with caution."¹⁴

A recent field study by Pope and Schweitzer (2011) finds that even experienced players on the PGA golf tour exhibit loss aversion. The stakes in this study, in contrast to the work by List (2003, 2004) are very large. In golf, on the PGA circuit, players play 4 rounds of 18 holes, for a total of 72 holes, over 4 days. The objective is to minimize the number of shots required to complete the 72 holes. In order to complete a hole, the golfer begins with a drive (or a *tee*) and then a series of shots that takes the golfer to the "green." At the green, the golfer engages in a series of sequential shots, until eventually the ball rolls into the targeted hole. Technically, the shots on the green are known as *putts* and the aim is to *sink the putt into the hole*.

Each hole in a golf course is associated with a number, called the *par*. The par is the minimum number of shots that a skilled golfer would take in order to sink the putt, starting with the drive. The par score can differ among the holes because of the distance from the tee to the hole, or the level of difficulty involved. On the PGA golf tour, the par scores for holes are typically 3, 4, or 5. The scores are recorded on score cards that measure the number of shots taken to sink the putt relative to the par score. Thus, if one sinks a putt using one stroke less than the par score, then one is said to have shot a *birdie*, and two strokes below par is an *eagle*. In these cases, the golfer is said to be *under par*. On the flip side, sinking the putt one stroke above par is called a *bogey* and two strokes above par is called a *double bogey*. In this case, the golfer is said to be *over par*. The third possibility is that the golfer might be *on par* (matching the par score).

Since all scores are measured relative to the par score, the authors make the plausible assumption that the par score for a hole acts as the reference point for that hole. It follows that if the golfer is *under par* (say, on a birdie or an eagle), he/she faces a different situation as compared to being *on* or *over par*. When on par, there is the possibility of missing the par score (reference point) so shots ought to be hit with relatively greater precision and accuracy, otherwise the golfer might be in the domain of losses and suffer loss aversion. Thus, we expect more successful par putts as compared to birdie putts.

Pope and Schweitzer (2011) explore this idea by using a sample in excess of 2.5 million data points with laser measurements of initial and final ball placements. They control for a large number of potentially relevant factors. These include the following. (i) The greater difficulty of birdie

¹⁴ MLA in this quote refers to *myopic loss aversion*; see Section 3.4.




Figure 3.1 Fraction of successful par and birdie putts by distance to the hole.

Source: Pope and Schweitzer (2011) with permission from the American Economic Association.

putt positions of the ball relative to the par putt positions. (ii) Information revealed by a birdie putt that helps the par putt. (iii) Player and tournament specific factors. The findings are as follows.

First, controlling for other factors, and consistent with a loss aversion account, golfers are less successful with their birdie putts, relative to comparable par putts, by 2–3 percentage points; see Figure 3.1. Second, the contrast between the birdie and par putts is even more noticeable on the first hole, where the par score is even more salient relative to subsequent holes.¹⁵ Third, if the top 20 golf players on the PGA tour in 2008 were to shed the trait of loss aversion, then they would each stand to gain \$1.2 million. Fourth, players are more risk averse in the gains domain than in the loss domain. In conjunction, these results make a persuasive case that loss aversion persists even with experience when large stakes are involved.¹⁶

Using field data from school teachers in Chicago, Fryer et al. (2012) demonstrate that teachers, including experienced teachers, exhibit loss aversion. They distinguish between two treatments. In one treatment, teachers were promised an “end of year bonus” conditional on student grades. In a second treatment, teachers were given financial incentives in the form of money, up-front. However, they were told that the money would be clawed back from them at the end of the year if the students did not achieve the desired educational levels. In EU, RDU, the results should

¹⁵ In subsequent holes, other factors such as competitor’s scores also become important.

¹⁶ Readers, particularly golf enthusiasts, may wonder about the role of confounding factors in these results. Here we quote from the authors (p. 130): “Beyond controlling for distance, consider and rule out several competing explanations for this finding. First, prior to hitting a par putt, players may have learned something about the green (by having already attempted a birdie putt). Second, birdie putts may start from a more precarious position on the green than par putts due to a longer approach shot. Third, player or Tournament-specific differences may bias our results. Using detailed data, we are able to rule out competing explanations with control methods and matching estimators. For example, we can match par and birdie putts attempted within one inch of each other on the exact same hole in the same Tournament. We are also able to rule out other psychological explanations. For example, we consider whether or not players become more nervous or overconfident when they shoot birdie putts relative to par putts.”

be identical under the two treatments. However, if teachers follow PT, then in the second treatment, teachers who are already endowed with the money would be loss averse with respect to giving back the money at the end of the year. So, as expected from a loss aversion based explanation, the effect of the financial incentives on student grades was significantly larger in the second treatment.

Kliger and Levy (2009) put EU, RDU, and PT to the test, using real financial assets market data. Their findings provide strong support for the individual components of PT, which include reference dependence, diminishing marginal sensitivity, loss aversion, and non-linear weighting of probabilities in the domain of gains and losses. Gurevich et al. (2009) also find support for PT using US stock market data, although they find loss aversion that is asset specific and, on average, lower than that found in student populations.

The work of Abdellaoui et al. (2013) nicely complements these findings using as their subjects, finance professionals (private bankers and money managers handling \$300 million on average). Unlike Gurevich et al. (2009) who use more aggregate data, Abdellaoui et al. (2013) are able to get data at the level of the individual finance professional, and observe if these “individual choices” are in accordance with PT. Almost all elements of PT find strong support (with the exception of loss aversion), and there is significant preference heterogeneity among the professionals. Furthermore, these professionals violate EU. These are significant findings because finance professionals are experienced and well trained in diversifying risk, and provided with extensive feedback on their performance.

3.2.4 *Subject misinterpretation as explanation of exchange asymmetries*

Plott and Zeiler (2007) (henceforth, PZ) conjecture that disparities between WTP and WTA are caused by the “procedure” that is used to elicit WTA and WTP. In the traditional experiments for exchange asymmetries, it was typical for the experimenter to announce the following; “I am giving you this object. It is a gift. You own it.” The *language* used in these instructions potentially creates problems.¹⁷ In order to rectify this potential problem, PZ tell subjects that a coin is tossed to determine their endowment, which could be one of two objects. In the typical exchange asymmetry experiment, subjects make their choices publicly. To ensure that there are no spillover effects of such announcements, PZ asked subjects to report their choices privately, in writing.

In a control group under the standard procedure, they detected the presence of exchange asymmetries. However, having made the procedural changes, listed above, they find that the endowment effect is weakened or eliminated (depending on the treatment). PZ take their results as a challenge to the traditional explanations of exchange asymmetries. The explanations offered by PZ of their results are behavioral in nature, such as framing effects, reciprocity, and gift-exchange arguments. Some of these results have been replicated by Knetsch and Wong (2008).

In an earlier paper, Plott and Zeiler (2005) (henceforth, PZ1), examine the effect of subject misperceptions in generating exchange asymmetries. Subjects are extensively *tutored* in economic logic, prior to the conduct of actual experiments.¹⁸ In a baseline experiment, absent these controls

¹⁷ They also raise other potential confounds. One potential problem is asymmetric information about the quality of the object used. However, the objects are typically quite simple such as mugs and pens, so asymmetric information might not be a serious issue. Another potential confound is that sellers may have other-regarding preferences, so they could reciprocate a gift from the experimenter by placing a high value on it.

¹⁸ As part of the tutorial, they are taught how and why economic mechanisms designed to elicit truthful preferences, such as the Becker DeGroot and Marshak (1964) mechanism, work. They are also taught about the cost

for subject misrepresentation, the phenomenon of exchange asymmetry is confirmed. However, in the presence of the controls, the means of the ordered (WTA , WTP) pairs for three different experiments were (5.69, 5.60), (5.71, 7.88), (5.06, 7.29). The mean WTP exceeded the mean WTA for two out of three experiments. The authors conclude that the observed gaps are the result of subject “misconceptions about the nature of experimental tasks.”

Like PZ, the major contribution of PZ1 is to bring out the importance of framing effects in the area of exchange asymmetries. Subjects in experiments alternate in the roles of sellers and buyers, which might alter attitudes that produce, say, the ownership effect that is important for exchange asymmetries. Furthermore, as PZ1 recognize, their procedures might have created a feeling on the part of subjects that the experimenter wishes to remove notions of ownership of the object. If this is the case (which cannot be rejected by the data from the experiment), then there could be a gap between the subject’s stated and true preferences.

Isoni et al. (2011) question the results of PZ1 along the following lines. PZ1 first run lottery experiments for training purposes and then run their experiments using mugs. PZ1 find that controlling for subject misperceptions, no exchange asymmetries are found for the case of mugs. But they do not report the experimental results for the case of lotteries because these were subjected to two forms of contamination. The first problem was that the selling task preceded the buying task and the second problem was that the experimenter publicly corrected mistakes and gave public answers to questions. The result was that PZ1’s influential results were based only on coffee mug experiments with 36 subjects reporting WTP and 38 reporting WTA .

The aim of the experiments by Isoni et al. (2011) is to report the results for the WTA/WTP for not only coffee mugs but also experiments based on lotteries that were not subjected to any contaminating influences. To determine the WTA , subjects were asked to report their WTA for the lottery $(x, p; y, 1 - p)$ and to determine the WTP , subjects reported their WTP for the lottery $(x + c, p; y + c, 1 - p)$, where $c > 0$. Under EU and constant absolute risk aversion, we have $WTP - WTA = c$, while the alternative hypothesis is $WTP - WTA < c$; this allowed for a within-subjects test of exchange disparities. The experimental format replicated PZ1, so it was intended to control for subject misrepresentations. The main finding is that no WTA/WTP disparities were found for mugs (as in PZ1), but significant exchange disparities were found in the case of lotteries.

The exchange disparity is not mitigated as subjects gain more experience in successive rounds. If the elicitation procedure holds the key to explaining exchange asymmetries, then it should apply to both lotteries and mugs. But since exchange asymmetries are found for lotteries but not mugs, Isoni et al. (2010) conclude that the elicitation procedure is possibly a red herring in explanations of exchange asymmetries. It is an open question as to why exchange asymmetries are found for lotteries but not mugs. The authors offer some conjectures using the mental accounting literature, such as the *house money effect* caused by show-up fees.¹⁹ Based on their results, the authors also conjecture that the difference between the results from lotteries and mugs are not due to the certainty/uncertainty dimension but due to the nature of money and consumption goods.

and benefits of pursuing alternative strategies. Subjects could ask questions, which the experimenter would then clarify. Numerical examples with concrete illustrations were also provided. Each experimental session consisted of a detailed training session, 2 practice rounds (unpaid), followed by 14 paid rounds using lotteries and finally 1 paid round with mugs that tested for exchange asymmetries.

¹⁹ The house money effect refers to the tendency to take greater risks in the presence of prior gains (Thaler and Johnson, 1990). For a form of the house money effect, based on a loss aversion explanation, and using stock market data, see Zhang and Semmler (2009).

Fehr et al. (2015) suggest that if subject misconception is the cause of WTA/WTP anomalies, then we ought to separate subjects into those who understand the elicitation mechanism used in the experiments, and those who do not. When they make this separation, and use the experimental protocol in PZ1 and the insights of Isoni et al., they find that the WTA/WTP anomaly arises for both lotteries and mugs; the findings for mugs are also replicated for USB sticks. Hence the title of their paper: The Willingness to Pay-Willingness to Accept Gap: A Failed Replication of Plott and Zeiler.

3.2.5 Other implications/explanations of the endowment effect

PROPERTY RIGHTS AND EXCHANGE ASYMMETRIES

A basic lesson from microeconomics is that indifference curves can be specified in the absence of any information on endowments and the budget constraint. However, one implication of WTP and WTA disparities is that an individual's indifference curves might cross if their endowment changes.

Figure 3.2 shows indifference curves in the space of mugs (a commonly used object in experiments for which the endowment effect is found), and money (for which no endowment effect arises). For any individual in the respective roles of buyer (no mugs, no money) and seller (one mug, no money), the respective endowment points are $e_b = (0, 0)$ and $e_s = (0, 1)$. Starting from the endowment point, the evidence on exchange asymmetries shows that a seller places a greater marginal value on another unit of mug (in terms of money) as compared to a buyer. Thus, the indifference curves for a seller are relatively flatter, as shown in Figure 3.2; $I_b I_b$ is an indifference curve for a buyer and $I_s I_s$ for a seller. The implication is that the indifference curves of the seller and the buyer, who only differ in their endowment, can cross.

The Coase theorem states that in the absence of any frictions/transaction costs, the initial allocation of property rights is irrelevant to the efficiency of the eventual outcome. However, if the endowment effect is accepted, then the act of assigning ownership (i.e., property rights) alters one's valuation (for sellers) or willingness to pay for the good (for buyers). This, in turn, alters the efficient outcome before and after the allocation of property rights. Hence, the Coase theorem, as it is traditionally stated, does not hold, or has to be modified to take account of ownership effects.




Figure 3.2 Crossing of indifference curves when there is an endowment effect.

The increase in the value of an object that arises purely from the act of ownership is also recognized in common law. For instance, the common perception that “possession is nine-tenths of the law” exemplifies this. Kahneman et al. (1991) credit the US Supreme Court justice, Oliver Wendell Holmes with the following quote. “It is in the nature of a man’s mind. A thing which you enjoyed and used as your own for a long time, whether property or opinion, takes root in your being and cannot be torn away without your resenting the act and trying to defend yourself, however you came by it. The law can ask no better justification than the deepest instincts of man.”

SOME MARKETING IMPLICATIONS OF EXCHANGE ASYMMETRIES

Novemsky and Kahneman (2005) give several implications of the endowment effect for marketing that are based on two premises.

1. When the benefits of the good whose ownership is about to be given up are similar to the benefits of the good to be acquired, then loss aversion with respect to the owned good is not likely to present. In this view, loss aversion operates along the benefits dimension of the good rather than along the attributes dimension.
2. Goods that are intended to be used only for exchange are not evaluated as losses. This assertion has two dimensions. First, firms or merchants that acquire a good with the sole intention of reselling it for money, do not experience loss aversion from the act of selling. This also potentially explains why List (2003, 2004) finds a lower degree of loss aversion for dealers relative to non-dealers. Second, some consumers/households may have a budget for miscellaneous items in order to cover for unforeseen expenditures. The goods acquired by spending money from this account might not be close substitutes for other goods. However, there is no loss aversion with respect to the money spent from this miscellaneous account. Other consumers may not have a miscellaneous account for unanticipated purchases, which typically are not substitutes for other goods. These consumers may suffer loss aversion for the money spent on making such unanticipated purchases.

Novemsky and Kahneman use these premises to offer the following applications. *Trial-offers*, which give consumers the opportunity to use a good for a limited period without making a payment, induce consumers to include the good as part of their endowment. Having used the good, the consumer experiences loss aversion by giving up the good at the end of the trial period. Such consumers may be willing to buy the good at the end of the trial period, but in the absence of the trial period, they might never have decided to buy. Yet another effect of the trial period is that the expenditure on the good could be transformed from being an unanticipated item, to which loss aversion applies, to being a part of the consumer’s main budget for intended purchases. Thus, the money spent on such goods once the trial period expires may not be subjected to loss aversion.

Yet another marketing device is to trade in old cars for new ones. Consumers of older cars might be loss averse to selling their old cars (endowment effect), creating a trade friction. However, by trading-in their old car for another car, whose benefits are very similar, loss aversion is mitigated (premise 1 above). Consumers are then more likely to buy newer cars.

Another implication of the differing behavior of sellers and buyers in studies of the endowment effect is the following. Buyers consider the *opportunity cost* of buying objects such as mugs. Not buying the mug is a *reduction in gain*. If the endowment effect is accepted, then sellers on the other hand are worried about *out-of-pocket costs*; parting with the mug is an *actual loss*. Hence, the disparity between WTP and WTA can alternatively be viewed as a disparity between (1) opportunity costs versus out-of-pocket costs or, equivalently, as (2) reduction in gain versus actual loss.

Kahneman et al. (1986) test these ideas using telephone interviews conducted on Canadian subjects. Successive questions, which were designed to draw out the contrast between “opportunity costs versus out-of-pocket costs” or “reduction in gain versus actual loss” were posed to the subjects. These questions were as follows.

Q.1a: (134 respondents) A shortage has developed for a popular model of automobile, and customers must now wait two months for delivery. A dealer has been selling these cars at list price. Now the dealer prices this model at \$200 above list price.

71% found the dealer’s action to be unfair and 29% found the action to be acceptable.

Q.1b: (123 respondents) A shortage has developed for a popular model of automobile, and customers must now wait two months for delivery. A dealer has been selling these cars at a discount of \$200 below list price. Now the dealer sells this model only at list price.

42% find the dealer’s action to be unfair and 58% find it acceptable.

In Q.1b, the elimination of the discount is viewed as a “reduction in gain.” This, we know, from the endowment effect, is less troubling for the individual relative to the imposition of the surcharge in Q.1a, which is viewed as an “actual loss.” Hence, a larger percentage of subjects think of the dealer’s action in Q.1a to be unfair. Along these lines, Thaler (1980) explains why firms that charge customers differently depending on whether they pay by cash or credit, frame the cash price as a discount (rather than saying that the credit price contains a surcharge).

Similar considerations might explain why it is easier to cut real wages during inflationary periods as the next two questions illustrate.

Q.2a: (129 subjects) A company is making a small profit. It is located in a community experiencing a recession with substantial unemployment but no inflation. The company decides to decrease wages and salaries 7% this year.

63% consider the company’s actions to be unfair, while the remaining 27% find the actions acceptable.

Q.2b: (129 subjects) A company is making a small profit. It is located in a community experiencing a recession with substantial unemployment and inflation of 12%. The company decides to increase salaries only 5% this year.

22% find the company’s actions to be unfair, while 78% find those actions acceptable.

The only difference in questions Q.2a and Q.2b is that in the latter, inflation takes place at 12%. In both questions, the final outcome for the workers is a 7% cut in real wages. In Q.2a, the outcome is framed as a nominal wage decrease (which subjects in experiments interpret as an actual loss). In Q.2b, on the other hand, it is framed as a reduction in (wage) gains on account of inflation, which subjects find relatively more acceptable.²⁰

STOCHASTIC REFERENCE POINTS

The exchange disparities between WTA and WTP can be explained if reference points are stochastic, say, in the Köszegi and Rabin (2006) framework; we do this in Section 2.8.3.

²⁰ While loss aversion can explain the outcome of these sets of questions, there are other competing behavioral explanations such as “money illusion” which is particularly relevant for Q.2a, Q.2b, and theories of reciprocity that take account of the role of intentions; see Parts 2 and 4 of the book.

A QUERY THEORY OF VALUE CONSTRUCTION

Another explanation of the endowment effect, called the *query theory*, has been proposed by Johnson et al. (2007). The theory is based on four main premises, all grounded in well known psychological principles.²¹

(1) *Queries*: When subjects in experiments are asked how much they would accept, or pay, for an object, say, a mug, they decompose the question into a series of *queries* such as the following. What are the advantages of my status quo (mug for the seller, and money for the buyer)? What are the advantages of pursuing the alternatives (in this case, arising from trade)? Advantages and disadvantages are encapsulated by means of positive and negative *aspects* of a situation. For a seller, the positive or value increasing aspects could include the following aspects. “Nice memento from the experiment; can be used as a Christmas present; you can never have too many mugs.” The negative, or value decreasing aspects could be things like the following. “This mug is ugly, I don’t need a coffee mug, or this mug looks dirty.” For a potential buyer, on the other hand, there are also positive aspects such as: “I am a low-income student, a few extra dollars will be very valuable”; and negative aspects: “I could hardly buy anything with this amount of money.” Queries could be automatic responses to questions that are unobservable to an outsider. Unlike queries, positive and negative aspects of a decision can be observed.

(2) *Sequential execution of queries*: Queries are executed *sequentially*. The order in which the queries are executed depend on the conditions in the experiment.

(3) *Competing memories*: Suppose that one is trying to recall some items from a list. Then the directed recall of some items in the list reduces the available memory for the unrecalled items in the list. In particular, when there are two main queries, then the recall of answers to one query suppresses recall of the answers to the second query; this is termed as *output interference*. In the particular context of this example, output interference implies that the retrieval of one source of information such as the advantages of the mug, reduces the retrieval of information about what could be done with that money.

(4) *Dependence on the endowment state*: This premise implies that the order in which the queries are executed depends on one’s endowment of the mug. So, if one is an owner of the mug, then the first executed query is the advantages of the mug. Alternatively, if the endowment is no-mug, say, the subject is a buyer, then the first executed query is alternative uses of the money.

The first three of these premises have reasonably good psychological bases, while the fourth, which is specific to the exchange asymmetry problem, needs to be tested. An important motivation of Johnson et al. (2007) is to test this premise.

In order to perform these tests, the authors perform an exchange asymmetry experiment using mugs. However, in addition, they use the method of *aspect listing*. While queries used and the order in which they are executed is not observable, the *aspects* (positive or negative) used by subjects in experiments are observable and can be used to form inferences about the queries used. *Aspect listing* accomplishes this by asking subjects to list the aspects they consider when making their WTA and WTP decisions. Subjects are also asked to note down the order of these aspects as well as the salience of various aspects. Trading in mugs was real, actual money was exchanged, and tests were carried out to check for subject miscomprehension. Subjects were randomly assigned as either *sellers*, *choosers*, or *buyers* and an incentive compatible elicitation mechanism was used.

²¹ We omit the references. The interested reader can directly look up the references in Johnson et al. (2007).

The results are as follows.

1. On average, the WTA for the mugs is \$5.71, while the WTP is \$3.42; the difference is statistically significant.
2. Sellers of the mugs produced more *value increasing aspects of mugs* that highlighted advantages of holding on to the mug. By contrast, buyers and choosers produced more *value increasing aspects of the alternative uses of money*. Furthermore, the order of queries was also in the expected sequential order. Sellers of mugs first produced value increasing aspects of the mugs before thinking about the value decreasing aspects of the mugs. In contrast, buyers first produced aspects relating to the alternative uses of the money.
3. Value increasing aspects increase the price and value decreasing aspects reduce the price of mugs.
4. Even stronger support for the theory is provided by an experiment in which the natural order of queries was reversed. Sellers were asked to first produce value decreasing aspects of the mugs, followed by the value increasing ones. Buyers were asked to produce value increasing, followed by value decreasing aspects of the mugs. This reversal of the natural order of queries eliminated exchange asymmetries entirely.

UNCERTAINTY AND EXCHANGE ASYMMETRIES

Engleman and Hollard (2008) distinguish between two alternative sources of uncertainty. *Choice uncertainty* refers to uncertainty about the object at the time of making the choice. For instance, the value of an object might be uncertain, or the individual might have incomplete preferences over objects in a set. On the other hand, *trade uncertainty* refers to uncertainty arising from the trading procedure and aspects of the trading partners. For instance, does the trading partner have “other-regarding preferences”? Is it morally correct to trade this object?

The authors conjecture that the endowment effect is caused by trade uncertainty. If true, then the presence of trade uncertainty causes people to trade less. In the words of the authors: “The market would thus be a poor teacher because traders will avoid those trades that would teach them the crucial lessons.” In order to examine these ideas, two treatments are considered. In the first treatment, subjects are free to trade, as in the traditional experiments that test for exchange asymmetries. However, in the second treatment, subjects were forced to trade; if they did not trade, they lost their endowments. This allowed subjects to experiment in uncertain situations, where natural caution would have prevented them from trading. In subsequent rounds of the second treatment, subjects in experiments do not exhibit the endowment effect. The authors provide one possible resolution of the endowment effect by artificially inducing more trade.

Roca et al. (2006) find that the status-quo bias arising on account of the endowment effect persists when an individual is endowed with an ambiguous lottery that can be exchanged with a lottery that has no ambiguity.²²

3.3 Prospect theory preferences in primates

Humans and other primate species share a common evolutionary history. Humans (genus *Homo*) and chimpanzees and bonobos (genus *Pan*) split from a common ancestor, the genus *Hominini*,

²² The ambiguous lottery is simply a ticket that allows one to bet on the unknown urn in Ellsberg’s experiments (see Chapter 1.4). The unambiguous lottery is a ticket to bet on the color of the ball from the known urn in Ellsberg’s experiments.



Figure 3.3 Tests of capuchin monkey preferences.

Source: Lakshminarayanan et al. (2008).

about 5 to 6 million years ago. The two immediately previous temporal evolutionary splits from common ancestors gave rise to, respectively, the gorillas, orangutans, and the gibbons. Humans differ experientially from the other primates in having specific market experience, however, evolution is very slow and conservative. Hence, common evolutionary features are likely to be exhibited in shared behavior between humans and other primates in several domains; for a survey, see Santos and Platt (2014).

We focus on studies that have been done on two species of monkeys, capuchin and macaque.²³ Macaques are extremely spatially dispersed Old World monkeys that are found mainly in Africa and Asia. The most common species of macaques in neurophysiological studies is the rhesus macaque. In contrast, the most common species of monkeys used in experiments on economic preferences and risk is capuchin monkeys. The habitat of these monkeys is in South and Central America; these are also sometimes known as New World monkeys. There are several reasons why capuchin monkeys are particularly suitable for these studies. The ratio of their brain to body size is high, they live in social groups, are socially sensitive, can learn from others, share food, engage in using tools, and quickly learn to trade tokens for food with experimenters (Fragaszy et al., 2004). Humans, chimpanzees, bonobos, and gorillas are part of the genus *Hominidae*, but monkeys belong to a different genus. Macaques separated from the evolutionary line that led to humans about 25 million years ago, while capuchins separated about 35 million years ago.

Chen et al. (2006) is one of the first studies to test elements of prospect theory on capuchin monkeys. They used monkeys who were trained to exchange tokens for food rewards. Under certainty, the monkey offers a token to the experimenter, who simultaneously holds a food item in one hand; see Figure 3.3. On receiving the token, the experimenter gives the food item to the monkey. Under risk, the experimenter offers a probabilistic reward, as explained below. The monkeys were placed in an enclosure and given a budget of tokens. Two experimenters wearing different colored clothing then positioned themselves on different sides of the enclosure. The monkey could trade tokens for food with both experimenters using openings in the enclosure.

²³ We ignore studies on pigeons and rats because they are more distant in evolutionary terms.

Two initial tests of consistency of capuchin preferences were conducted. In the first test, the monkeys' original allocation of budget on two food items (jello pieces and apple slices) was established after repeated trials. Then the monkeys were given a new budget but the price of one of the food items, in terms of tokens, was dropped. The monkeys substituted consumption towards the relatively cheaper item by spending a greater fraction of their budget on it. In fact, the choices of monkeys reliably satisfied the generalized axiom of revealed preference. As the authors note, many naive models and even some forms of random behavior can satisfy GARP.

In the second consistency test, experimenter 1 always offered and traded one apple slice in return for one token. However, experimenter 2 made a risky offering; he initially held out two apple slices to the monkey, but was equally likely to offer either two slices or one slice after receiving the token from the monkey. Thus, on average, experimenter 2 offered 1.5 slices, which is higher than the number offered by experimenter 1. Clearly, the offering of experimenter 2, first order stochastically dominates the offering of experimenter 1. Over 60 trials, spread over five sessions, 87% of the trading choices were directed towards experimenter 2. Thus, the vast majority of capuchins satisfied stochastic dominance.

In order to test for reference dependence, the offerings made by the two experimenters were as follows. In each case, the initial offer made by any experimenter, say, offering x slices of apples in return for a token, may be considered as the reference point. However, once a token is received from the monkey, the experimenter offers a trade in the following lottery in apple slices $(y, 0.5; z, 0.5)$. For experimenter 1, $x = y = 1, z = 2$ for an average offering of 1.5, while for experimenter 2, $x = y = 2, z = 1$, for an average offering of 1.5. The reference point is low for experimenter 1 and half the time his offering leads to a gain relative to the reference point. In contrast, the reference point for experimenter 2 is high and half the time his offering leads to a loss. Pooled across all monkey subjects, in the last five sessions, 71% of the monkeys traded with experimenter 1; the results are statistically different from an equal split in trading between the two experimenters. One cannot explain this result in a theory that ignores reference dependence of preferences. For instance, under EU this implies that $0.5u(1) + 0.5u(2) > 0.5u(2) + 0.5u(1)$, which is absurd.

Finally, in order to test for loss aversion, experimenter 1 always initially offered two pieces of apple but after receiving a token, always traded one piece in return. Experimenter 2 always initially offered one piece and always traded one piece in return for a token. The final choice is identical for both experimenters, however, under reference dependence, the offering of experimenter 1 invokes loss aversion. As expected, in nearly 80% of the trades, monkeys preferred to trade with experimenter 2, suggesting an explanation based on loss aversion.

Lakshminarayanan et al. (2011) test for different risk attitudes in the domain of gains and losses with capuchin monkeys. In one experiment, both experimenters framed their offerings as a gain. Experimenter 1, the safe experimenter, initially offered one piece of food, but on receiving the token always traded two pieces. Experimenter 2, the risky experimenter, initially offered one piece of food, but then was equally likely to trade either one or three pieces in exchange for a token. Most trades occurred with the safe experimenter. However, when the offerings were framed as losses, most trades occurred with the risky experimenter. These results are consistent with prospect theory preferences but are not immediately consistent with theories that do not have a reference point.

Brosnan et al. (2007) wonder if the endowment effect arises from the shared evolutionary history of humans and chimpanzees, on account of a common ancestor? To test this conjecture, they perform standard endowment effect experiments, using food and non-food items on 33 chimpanzees. They ran six trials with items that are familiar to chimpanzees. The food items were a




Figure 3.4 Choices made by chimpanzees in different conditions. Hatched bars represent the percentage of the population that preferred the object in a choice condition, and solid bars represent the percentage of the population that chose to maintain possession. Vertical bars show standard errors of means.

Source: Reprinted from *Current Biology*, 17(19): 1704–7. Brosnan, S. F., Jones, O. D., Lambeth, S. P., Marengo, M. C., and Schapiro, S. J., “Endowment effects in chimpanzees” Copyright © 2007, with permission from Elsevier.

frozen fruit juice stick and a PVC pipe filled with peanut butter (henceforth, juice and PB). The non-food items were a rubber-bone dog chew toy and knotted-rope dog toy (henceforth, bone and rope).

Three trials were run for each of the food and non-food items. The trials were similar to those run for human subjects with the difference that the exchange of an item by a chimpanzee takes place with a human at the other end (and not with another chimpanzee). Prior to the experiment, chimpanzees are trained to exchange items with an experimenter. If they wish to exchange an item for another, they must do so within two minutes, otherwise no exchange takes place. The results are shown in Figure 3.4.

In the food trial, it was found that 58% of the chimpanzees preferred PB to juice. When endowed with the PB, however, 79% of the chimpanzees preferred to keep it rather than exchange it for juice. In another trial, when chimpanzees were endowed with juice, 58% of the chimpanzees preferred not to trade their juice for the PB.

For non-food items, the outcome was opposite to that predicted by the endowment effect: 74% of the chimpanzees preferred the bone over the rope. Despite this, when chimpanzees are given ownership of the bone, they prefer to keep it only 16% of the time. When endowed with the rope, they preferred to keep it only 10% of the time (rather than the predicted 26%). There are two possibilities for this difference in result between the food and non-food items. First, chimpanzees simply do not have a strong enough preference for non-food items. Second, they like interacting with the experimenter.

To test for the possibility that chimpanzees simply like interaction with the experimenter, further tests were carried out. For food items, no chimpanzee (with one exception) exchanged one food item with an identical item offered by the experimenter. However, for non-food items, the opposite result was found: 82% traded a bone for a bone, while 79% traded a rope for a rope.

In conjunction, these results suggest that chimpanzees exhibit the endowment effect for evolutionary salient items such as food, but there is no endowment effect for non-food items. The evolutionary explanation is that food has a greater bearing on the fitness (i.e., reproductive success) of the animal relative to non-food items. Since humans share a common evolutionary history with chimpanzees, the authors conjecture that it is also likely that humans too possess the endowment

effect. Lakshminarayanan et al. (2008) showed that the results on endowment effect survive when one controls for transactions costs and timing issues.

3.4 Myopic loss aversion

Yet another determinant of loss aversion is the horizon over which a risky situation is to be evaluated. *Myopic loss aversion* (MLA) relies on two concepts. The first is loss aversion. The second is *mental accounting*, due to which individuals may evaluate the *return* on lotteries over very short horizons, even when the *actual return* accrues over a longer horizon.

Example 3.24 (Gneezy and Potters, 1997): Consider the following PT utility function that exhibits loss aversion.

$$v(x) = \begin{cases} x & \text{if } x \geq 0 \\ 2.5x & \text{if } x < 0 \end{cases}. \quad (3.13)$$

Suppose that the individual is offered two identical prospects on two separate days. The prospect is (200, 0.5; −100, 0.5). Assuming linear weighting of probabilities, the utility of the prospect under PT is

$$0.5(200) - 0.5(250) = -25 < 0. \quad (3.14)$$

Now consider two individuals who use two different mental accounting procedures. Individual *S* uses a short horizon and evaluates the prospect on each day separately. From (3.14), his utility on each day is negative, so he refuses the prospect on each day. The other individual, *L*, who uses a long horizon, combines the two prospects on the two days to give the composite prospect (400, 0.25; 100, 0.5; −200, 0.25). Using the utility function in (3.13), the utility of the combined prospect for individual *L* is

$$0.25(400) + 0.5(100) - 0.25(500) = 25 > 0.$$

MLA applies to individual *S*, who does not choose to invest in the gamble. Individual *L*, who does not use MLA, in contrast, prefers to invest in the gamble.

Gneezy and Potters (1997) provide direct evidence for MLA. They gave individuals an endowment of $e = 200$ cents; a part $x \in [0, 200]$ of the endowment could be invested. The individual was then offered a lottery $(-x, 2/3; 2.5x, 1/3)$ and asked to make an investment decision over nine rounds. The same level of endowment, e , is given at the beginning of each round. In any round, subjects could only bet using their endowment for that round.²⁴

In one treatment, treatment *S* (short horizons), subjects were asked to make an investment decision, x , in every period, having observed the outcome of the lottery in the previous period. Thus, a total of nine investment decisions had to be made. In the second treatment, treatment *L* (long horizons), investment decisions are made in three rounds 1, 4, 7 and, in each case, the investment decision also applies for the following two periods. Individuals can observe their earnings in rounds 1–3 when they choose their investment in round 4, and observe earnings in rounds 4–6 when they choose their investment in round 7.

²⁴ Three more rounds were run after the 9th round. In each of these rounds, the money won in the first nine rounds divided by 3 became the new endowment level.

The average level of investment in treatment *S* is found to be significantly lower as compared to treatment *L*. On average, subjects bet about half their endowment in treatment *S* and about two-thirds of their endowment in treatment *L*. The basic result on MLA has proven to be robust to several changes in the experimental design, such as changing the number of rounds, and the number of lotteries.²⁵

Haigh and List (2005) examine MLA for professional traders in finance, relative to a control group of undergraduate students. There were 32 students and 27 traders. They elicited, over several rounds, the amount of money each of the groups was willing to bet in two treatments. In one treatment (Treatment I) subjects were given “infrequent” feedback on the returns to their investment, while in the other (Treatment F) “frequent” feedback was given. The prediction of MLA that there should be less betting on the risky investment in Treatment F relative to Treatment I was confirmed for both subject groups. However, professional traders exhibited statistically significantly greater MLA as compared to students; see, Figure 3.5. Traders bet an average of 75 units in Treatment I but only 45 units in Treatment F. The corresponding numbers for students are 62.50 and 50.89. These results are in contrast to the ones by List (2003, 2004), reviewed above, that seek to demonstrate elimination or reduction in loss aversion with real-world experience.

Sutter (2007) uses the framework of Gneezy and Potters (1997) to determine if membership of a team weakens or strengthens MLA. Subjects were asked to make their decisions individually, or as part of a team of three individuals. The endowment was $e = 100$ euro-cents. The two conditions were short (*S*) and long (*L*) planning horizons. The results are shown in Figure 3.6. Two main findings emerge from Sutter’s (2007) experiments. First, individuals invest more when they are part of a team. Second, individuals and teams invest more when the horizon is longer.




Figure 3.5 Average bets of students and professionals.

Source: “Do professional traders exhibit myopic loss aversion? An experimental analysis,” Haigh, M. and List, J. A. *The Journal of Finance*, 60(1): 523–34. Copyright © 2005, John Wiley and Sons.

²⁵ See, for instance, Thaler et al. (1997), Barron and Erev (2003), Langer and Weber (2008), and Bellamare et al. (2005). But for evidence to the contrary, see Blavatskyy and Pogrebna (2009).




Figure 3.6 Amount invested in each treatment.

Source: Reprinted from *Economics Letters*, 97(2): 128–32. Sutter, M., “Are teams prone to myopic loss aversion? An experimental study on individual versus team investment behaviour.” Copyright © 2007, with permission from Elsevier.

The frequency of the evaluation of a lottery and the horizon over which the return from the lottery accrues, play a similar role in determining the effect of MLA on the amount invested. So what is the relative contribution of each of these factors. The experimental results of Fellner and Sutter (2009) indicate that each of these factors contributes equally. They find that making subjects opt for either longer horizons or a lower frequency of evaluation leads to an increase in investment.

Benartzi and Thaler (1995, p. 73) make the following observation: “Since 1926 the annual real return on stocks has been about 7 percent, while the real return on treasury bills has been less than 1 percent.” The difference in returns between the two assets is not explained by differences in the standard deviations; the annual standard deviation of stock returns over the period 1926–1997 is about 20%. This is known as the *equity-premium puzzle*.

The main distinction between stocks and bonds in terms of their returns is the greater variability of returns on stocks. Risk averse individuals will naturally demand a greater return on stocks relative to bonds. However, Mehra and Prescott (1985) find that for EU to explain the equity premium puzzle, an unreasonably high coefficient of relative risk aversion of about 30 is required.²⁶ The actual coefficient of risk aversion is close to 1 (see, Bombardini and Trebbi, 2012), and in some estimates, up to about 5.

Benartzi and Thaler (1995) seek to explain the equity premium puzzle by PT and narrow bracketing (stock returns are considered in isolation from all other returns and sources of wealth). They first ask, what planning horizon must investors be choosing so that they are indifferent between stocks that give a 7% return and bonds which give a 1% return? Their empirical results give a best estimate of one year. This is plausible, given that individuals prepare tax returns, and receive company annual newsletters over the same horizon. They term a one year planning horizon as *myopic investor behavior*. Furthermore, most stock returns fluctuate a lot over a one year horizon, hence, the investor potentially experiences repeated gains and losses. The resulting loss aversion

²⁶ Mankiw and Zeldes (1991) illustrate the absurdity of such a high level of risk aversion. A decision maker with a relative risk aversion coefficient of 30 evaluates the certainty equivalent of the lottery (100,000, 0.5; 50,000, 0.5) to be 51209.

accentuates the downward movements in returns, and ensures that the return on stocks must be sufficiently high if they are to be held in a portfolio that contains bonds. Under these assumptions, it turns out that a 7% return on stocks is indeed required to explain the observed compositions of portfolios between stocks and bonds.

Example 3.25 Suppose that initial wealth I is 1 and the decision maker uses PT with linear probability weighting. The reference income, r , is the status quo, so it also equals 1. There are two assets, each with a maturity period of 1 year. A safe asset offers a return of 0.01. The value function is piecewise linear so it takes the form

$$v(x) = \begin{cases} x & \text{if } x \geq 0 \\ -\lambda(-x) & \text{if } x < 0 \end{cases}.$$

The coefficient of loss aversion $\lambda = 2.25$. The value function from holding the safe asset is

$$V_S = (1 + 0.01)I - r = (1 + 0.01) - 1 = 0.01. \quad (3.15)$$

Assume that the return on the risky asset, θ , can take either of two values, a low level of return, $\theta_L = -0.06$ (with probability 0.8), or a high level of return θ_H (with probability 0.2). Hence, there is one outcome each in the domain of gains and losses. The value function in this case is given by:

$$V_R = 0.8(-2.25(0.06)) + 0.2\theta_H.$$

What level of return, θ_H , is required for the decision maker to be indifferent between the two assets? Setting V_S equal to V_R , we get that $\theta_H = 0.59$. Hence, the average return on the risky asset is

$$0.2(0.59) - 0.8(0.06) = 0.07 = 7\%,$$

which is seven times higher than the return on the riskless asset. Thus, relatively high average returns are required on the risky asset in the presence of loss aversion. This is the main insight of Benartzi and Thaler.

The insights in the Benartzi–Thaler model were formalized more fully in Barberis et al. (2001). While there are no direct empirical tests of these insights, Barberis (2013) cites the empirical work of Dimmock and Kouwenberg (2010) that found a relation between loss aversion and stock market participation.

3.5 Why do people pay taxes?

So far, we have seen applications where researchers selectively choose the components of PT that they feel are most essential for explaining their research questions. We now consider one of the few applications that uses all the components of PT. Issues of tax evasion are extremely important and losses to society from tax evasion are huge.²⁷ The existing literature on tax evasion, since

²⁷ For the US, for example, Slemrod and Yitzhaki (2002) give a conservative estimate of the *tax gap* (difference between taxes paid and taxes owed) of the order of \$300 billion per year.

Allingham and Sandmo (1972) and Yitzhaki (1974), had largely assumed that taxpayers use EU. Dhimi and al-Nowaihi (2007, 2010b) used, instead, all components of PT to model the tax evasion decision; a simplified version of the model is given below.²⁸

An EU model of tax evasion leads to several puzzles.

1. Quantitative puzzle: Suppose that tax is not withheld at source, and taxpayers have a choice to report or evade their income. Consider a taxpayer who follows EU and evades \$1. With probability $1 - p$, he is not caught, so he can keep \$1. With probability p , he is caught, so he surrenders \$1 to the tax authorities, and, in addition, pays a fine at the rate $\theta > 0$ on \$1. Thus, under EU, the taxpayer will evade taxes if the expected return per dollar on evading the tax,

$$1 - p - p\theta > 0. \quad (3.16)$$

In actual practice, for the amateur tax evader, the audit probability, p , ranges from 0.01 to 0.03, while the penalty rate, θ , on the amount evaded, ranges from 0.5 to 2.0. Using these values of p and θ , the expected return on tax evasion lies between 91 and 98.5%; it would be hard to find assets with such high rates of return. So, why do people pay taxes? This contradicts the empirical evidence in two ways. First, only about 40% of taxpayers who have a chance of evading, do evade. Further, experimental evidence suggests that at least 20% of taxpayers fully declare all their income, even when $p = 0$. Second, to make these results consistent with EU, taxpayers are required to be risk averse to an absurd degree.²⁹

2. Qualitative puzzles: In addition to the quantitative puzzle, there are also qualitative puzzles arising from the application of EU to tax evasion. First, Yitzhaki (1974) showed that using EU, and under the reasonable assumption of decreasing absolute risk aversion (DARA), an increase in the tax rate leads to a decrease in tax evasion. However, most empirical evidence rejects this result; people try to hide more income if the government tries to extract greater tax revenues. Second, obligatory advance tax payments should not influence the taxpayer's evasion decision under EU. This is because under EU, the carriers of utility are final levels of wealth (see Remark 1.1); final levels are identical under advance tax payment or post tax payment. In contrast, obligatory advance payments alter tax evasion; see, for example, Yaniv (1999) for the evidence and an explanation based on PT.

Dhimi and al-Nowaihi (2007) run a two-horse race that pits EU against PT. They show that these qualitative and quantitative paradoxes can be explained by PT.

THE BASIC SET UP

The government moves first and levies a tax at the constant marginal rate $t \in [0, 1]$ on the declared income of the taxpayer, D . Each taxpayer who has exogenous taxable income $W > 0$ can then choose to declare an amount $D \in [0, W]$; we focus only on those individuals for whom tax is not withheld at source. An exogenous fraction $p \in [0, 1]$ of the taxpayers are randomly audited,

²⁸ We shall keep references to a minimum. The interested reader can look up Dhimi and al-Nowaihi (2007, 2010b) for the references, the justification behind the assumptions, robustness tests, etc.

²⁹ In one example, Skinner and Slemrod (1985) calculate that in order to square the observed tax evasion rates a coefficient of relative risk aversion (CRRA) of 70 is required, while realistic magnitudes of CRRA lie between 1 and 5.

subsequent to the filing of tax returns, and the audit reveals the true taxable income; Dhimi and al-Nowaihi (2007) consider the more general case where this probability is *endogenous* and it depends negatively on the amount of income declared. If caught, the dishonest taxpayer must pay the outstanding tax liabilities $t(W - D)$ and a penalty on the amount evaded, $\theta t(W - D)$, where $\theta > 0$ is the penalty rate. If evasion is discovered, the taxpayer also suffers some stigma $s(W - D)$, where $s \in [0, 1]$ is the stigma rate on evaded income.³⁰ There is a distribution of stigma rates over the taxpayers with density $\phi(s)$. Thus, our treatment of the stigma rate parallels the tax rate; both lie in the interval $[0, 1]$.

Denoting by Y_A and Y_{NA} , respectively, the income of the taxpayer when he is audited and when he is not, for any $D \in [0, W]$ we have

$$\begin{cases} Y_{NA} = W - tD, \\ Y_A = W - tD - t(1 + \theta)(W - D) - s(W - D). \end{cases} \quad (3.17)$$

Under EU, individuals derive utility from their final wealth levels. In order to facilitate comparison with the utility function under PT (see (2.20)) the utility of any outcome for an EU decision maker is defined as $v(Y_i) = Y_i^\beta$; $\beta \in (0, 1)$.³¹ Thus, the EU of a taxpayer is given by

$$EU = (1 - p)Y_{NA}^\beta + pY_A^\beta. \quad (3.18)$$

Under PT, take the *legal after-tax income*, $R = (1 - t)W$, as the reference point of the taxpayer; this is in the spirit of the status-quo justification of reference points. A simple calculation shows that (provided R is independent of p) this is the *unique* reference income, that has the following desirable property. For any D , if the taxpayer is caught, he is in the domain of losses, and if he is not caught, he is in the domain of gains. Thus, it is an eminently sensible reference point.

The taxpayer's income relative to the reference point, in the two states, is $X_i = Y_i - R$, $i = A, NA$. Using (3.17), we get

$$\begin{cases} X_{NA} = t(W - D) \geq 0 \\ X_A = -(\theta t + s)(W - D) \leq 0 \end{cases}. \quad (3.19)$$

The taxpayer is in the *domain of gains* in the No Audit state, and the magnitude of the gains is the amount of the tax evaded. In the Audit state, the taxpayer is in the *domain of losses*; the magnitude of the loss equals the penalty on the evaded tax plus stigma costs. The taxpayer faces the following lottery in incremental form (or *prospect*), denoted by L .

$$L = \{-(\theta t + s)(W - D), p; t(W - D), 1 - p\}.$$

³⁰ The UK government initiative of naming and shaming tax offenders, spearheaded by the HMRC, is at least partly motivated by tapping into such stigma costs of the taxpayer.

³¹ From (3.17) it is possible that $Y_A < 0$. The simplest way to incorporate $Y_A < 0$ into EU is to define $v(Y_A) = 0$ whenever $Y_A < 0$ as we have done. Other possibilities are $v(Y_A) = -v(-Y_A)$ and $v(Y_A) = -\infty$. None of our results depend on which specification we use for $v(Y_A)$ when $Y_A < 0$.

The *value function* under PT is given by (2.24), which in this case is

$$V(L) = \pi(1-p) v(X_{NA}) + \pi(p) v(X_A),$$

where v is the utility function under PT. Substituting (3.19) in the value function, using the power form of utility in (2.20), and the decision weights in Definition 2.33, we get the objective function of the taxpayer under PT as

$$V(L) = w(1-p) [t(W-D)]^\gamma + w(p)(-\lambda) [(\theta t + s)(W-D)]^\gamma. \quad (3.20)$$

THE TAX EVASION DECISION UNDER EXPECTED UTILITY

Maximizing (3.18) by a choice of D , the solution to the declared income is

$$D^* = W \left[1 - (1-t) \left[(\theta t + s) + t \left(\frac{(\theta t + s)p}{t(1-p)} \right)^{\frac{1}{1-\gamma}} \right]^{-1} \right]. \quad (3.21)$$

Direct differentiation of (3.21) gives rise to all the comparative static results under EU, summarized in Proposition 3.29.

Proposition 3.29 *Under EU, tax evasion is decreasing in the probability of detection and conviction, p , the penalty rate, θ , and the rate of stigma, s . If preferences exhibit decreasing absolute risk aversion, then tax evasion is decreasing in the tax rate, t (Yitzhaki (1974) puzzle).³²*

Simple calculations give the following result.

Proposition 3.30 *Suppose that the taxpayer follows EU. Consider the set of (θ, p) points implicitly generated by the inequality*

$$\theta \geq \frac{1-p}{p} - \frac{s}{t}. \quad (3.22)$$

In (θ, p) space, (3.22) gives the loci of points such that a taxpayer with stigma $s \in [s, \bar{s}]$ reports the full amount of income. At all points below this loci, the taxpayer chooses to evade some strictly positive fraction of income.

Using (3.22), in the absence of stigma costs, $s = 0$, the taxpayer will evade income when $1 - p - p\theta > 0$, as claimed in (3.16).

THE TAX EVASION DECISION UNDER PROSPECT THEORY

The taxpayer maximizes the value function in (3.20) by an appropriate choice of D , given the probability weighting function $w(p)$. Rewriting (3.20),

$$V(L) = (W-D)^\gamma h(\theta, p, s, t, \lambda, \alpha, \gamma), \quad (3.23)$$

³² For the assumed utility function under EU, the coefficient of absolute risk aversion (for EU we need only consider the domain of gains) is given by $\frac{1-\beta}{\gamma}$ which is decreasing in γ . Hence, the Yitzhaki (1974) result applies.

where,

$$h = h(\theta, p, s, t, \lambda, \alpha, \gamma) = t^\gamma \left(w(1-p) - \lambda \left(\theta + \frac{s}{t} \right)^\gamma w(p) \right) \geq 0. \quad (3.24)$$

Since the value function in (3.23) is monotonic in D , the solution is given by:

$$D = \begin{cases} 0 & \text{if } h > 0 \Leftrightarrow \left(\frac{1}{\lambda} \right) \frac{w(1-p)}{w(p)} > \left(\theta + \frac{s}{t} \right)^\gamma \\ \in [0, W] & \text{if } h = 0 \Leftrightarrow \left(\frac{1}{\lambda} \right) \frac{w(1-p)}{w(p)} = \left(\theta + \frac{s}{t} \right)^\gamma \\ W & \text{if } h < 0 \Leftrightarrow \left(\frac{1}{\lambda} \right) \frac{w(1-p)}{w(p)} < \left(\theta + \frac{s}{t} \right)^\gamma \end{cases}. \quad (3.25)$$

(3.25) depicts a *bang-bang solution*, which might be descriptive of several forms of tax evasion that take the form of hiding certain activities completely from the tax authorities. For the general case considered in Dhimi and al-Nowaihi (2007), in which the probability of detection is endogenous and depends on the amount declared, one gets continuous dependence of the solution on the parameters of the model. The bang-bang solution arises in the special case of an exogenous and fixed probability of detection.

In (3.25), a large enough change in the parameters can push the solution from one case to the other. Starting from the case $D = 0$, the following factors are conducive to completely eliminating tax evasion ($D = W$). (1) Low tax rate, t , (2) high levels of stigma, s , and high penalties for tax evasion, θ , (3) high levels of loss aversion, λ , and (4) overweighting of the probability of a loss, p , which leads to a decrease in $\frac{w(1-p)}{w(p)}$.

Under EU and PT, tax evasion decreases with an increase in p , θ , and s , although, as we shall see below, the quantitative effects under PT are in accord with the evidence, while those under EU are not. There is also no counterpart in EU to the effect on tax evasion of (i) loss aversion, and (ii) overweighting of small probabilities. The main qualitative difference between EU and PT is in the predicted effect of the tax rate on the amount evaded. The result under EU is given in Proposition 3.29. However, PT predicts the empirically supported result that tax evasion increases with an increase in the tax rate. This is formally stated in Proposition 3.31 below (the simple proof is omitted).

Proposition 3.31 *Ceteris paribus, $\exists t = t_c \in [0, 1]$ such that the individual does not evade taxes if $t < t_c$ but evades taxes if $t > t_c$.*

A rough but useful intuition about the effect of the tax rate can be given as follows (the actual results are more subtle). An increase in the tax rate makes the taxpayer poorer. Hence, under decreasing absolute risk aversion, EU predicts a taxpayer would become more risk averse, and so evade less. By contrast, under PT the utility function is convex for losses. So, if the reference point is such that an audited taxpayer finds himself in the domain of losses (as is indeed the case), then an increase in the tax rate causes him to be poorer and, hence, more risk seeking. So he evades more.

A comparative static result similar to the one in Proposition 3.31 can be stated with respect to the stigma rate s (the simple proof is omitted). If stigma is high enough, then people do not evade any taxes; a similar result can easily be seen to hold for EU.

Proposition 3.32 *Ceteris paribus, $\exists s = s_c \in [0, 1]$ such that all individuals characterized by $s > s_c$ do not evade taxes, while all those with $s < s_c$ evade taxes.*

MODEL CALIBRATION

The analysis below provides simulation results in a calibrated model.³³ We use the Prelec probability function (see Definition 2.22) with $\alpha = 0.35$ and $\beta = 1$. For the utility function, we use the power form of utility in (2.20). From estimates in Tversky and Kahneman (1992), $\gamma \approx 0.88$ and $\lambda \approx 2.25$. Let us use a realistic tax rate of 30%, i.e., $t = 0.3$.

Based on TCMP (taxpayer compliance measurement program) data, Andreoni et al. (1998) report that about 40% of taxpayers report their incomes incorrectly in the US; see also Bernasconi (1998). Since these figures include those who unintentionally report incorrectly (about 7%), let us use the more conservative estimate that only about 30% of taxpayers who have a chance, do evade taxes. Assume that the density of the stigma rate, $\phi(s)$, is uniform. From Proposition 3.32, we know that under PT all individuals characterized by $s > s_c$ do not evade taxes, while all those with $s < s_c$ evade taxes. Given the assumption that the stigma rate is uniformly distributed over $[0, 1]$, the hypothesis that approximately 30% of taxpayers evade taxes corresponds to $s_c = 0.3$, which is the value we shall use for the calibration exercise.

We ask the following question. Given realistic magnitudes of tax evasion, about 30%, how close are the predicted θ, p values under EU and PT, respectively, to their actual values?³⁴

Under EU, using (3.22), and the calibration values, the locus of (θ, p) combinations that need to be consistent with the actual data is given by

$$\theta^{EU} = \frac{1-p}{p} - \frac{0.3}{0.3}, \quad (3.26)$$

where θ^{EU} may be taken to be the predicted value of θ under EU, conditional on all the other parameters of the model. Using (3.25), and the Prelec weighting function, a similar (θ, p) locus under PT is given by

$$\theta^{PT} = \left(\frac{1}{2.25} \right)^{\frac{1}{0.88}} \Gamma(p)^{\frac{1}{0.88}} - \frac{0.3}{0.3}, \quad (3.27)$$

where $\Gamma(p) = \exp [(-\ln p)^{0.35} - (-\ln(1-p))^{0.35}]$ and $\alpha = 0.35$. θ^{PT} may be taken to be the predicted value of θ under PT, conditional on all the other parameters of the model.

The numerical magnitudes of (θ, p) that correspond to the two loci in (3.26), (3.27) are shown in Table 3.1; in the top row, we use probabilities ranging from 0.005 to 0.050. For most realistic audit probabilities, i.e., $p \in [0.01, 0.03]$ the prediction of the penalty rate, θ , ranges from 0.66 to 1.21 for PT. This is consistent with actual values for the audit probabilities which range from 0.5 in the USA to 2.0 in Italy with several other Western countries in the middle (Bernasconi, 1998). On the other hand, EU predicts a penalty rate that is up to 100 times larger in this range and, on average, about 60 times higher.

It is sometimes argued that sufficiently high stigma can align the predictions of EU with the evidence. This is not correct. For instance, choose a middle probability of around 0.03 and increase the stigma rate from 30% (which forms the basis of the calculations in Table 3.1) to the maximum

³³ For robustness tests with respect to changes in the calibration values, such as the parameters of the Prelec function and alternative assumptions about the rate of stigma, s , the reader can consult Dhami and al-Nowaihi (2007). All the results go through.

³⁴ We could have equivalently posed our question in many different ways. For example, given actual (θ, p) combinations, what level of tax evasion does each theory predict?

Table 3.1 Comparison of the predicted value of penalty rates under EU and PT.

	.005	.010	.015	.020	.025	.030	.035	.040	.045	.050
θ^{EU}	197.7	97.7	64.3	47.7	37.7	31.0	26.2	22.7	19.9	17.7
θ^{PT}	1.6	1.2	1.0	0.9	0.8	0.7	0.6	0.5	0.5	0.4

possible rate, 100%. The new value of $\theta^{EU} = \frac{1-0.03}{0.03} - \frac{1}{0.3} = 29.0$, which is nowhere near the 0.5–2 range that is observed in actual practice.

On the other hand, in PT, all components of the theory (reference dependence, diminishing sensitivity, loss aversion, non-linear probability weighting) were needed in order to explain the qualitative and quantitative tax evasion puzzles. While the penalty rate and the probability of detection may appear low, PT decision makers overweight the small probability of being caught and suffer loss aversion, which magnifies the penalty rate. This provides strong vindication for the choice of PT in explaining actual parameters of policy choice relevant for the tax evasion problem. It is remarkable that the predicted magnitudes of (θ, p) are so close to the actually observed values when one considers that the parameter values used for the calibration exercise were obtained from independent experimental evidence applied to generic situations of risk. This model also shows that PT is quite simple and tractable to apply.

3.6 Explanation of Rabin's paradox using PT

Rabin's paradox for EU was described in Section 1.5.4. We now show that the concepts of reference dependence and loss aversion, that are central to PT, can explain Rabin's results, an explanation that Rabin (2000a) favored. Recall the three lotteries in Section 1.5.4, L_0, L_1, L_2 :

$$L_0 = (w, 1), L_1 = (w - 10, 0.5; w + 11; 0.5), L_2 = (w - 100, 0.5; w + m, 0.5), \quad (3.28)$$

where m is any positive reward level, no matter how high. Under PT, lotteries are defined in incremental form. Assume that initial wealth, w , is the reference point. Redefine the three lotteries, L_0, L_1, L_2 , in incremental form (or prospects) as follows

$$L_0 = (0, 1), L_1 = (-10, 0.5; 11; 0.5), L_2 = (-100, 0.5; m, 0.5). \quad (3.29)$$

Let us continue with Example 1.4 from Section 1.5.4 and show that Rabin's paradox can be explained with PT. For any level of wealth x , let $y = x - w$ be the wealth relative to the reference point (initial wealth level, w). Adapt the utility function in (1.44) to define the reference dependent utility function of the individual as

$$v(y) = \begin{cases} 1 - e^{-y/1000} & \text{if } y \geq 0 \\ -\lambda(1 - e^{y/1000}) & \text{if } y < 0 \end{cases}, \quad (3.30)$$

where $\lambda > 1$ is the parameter of loss aversion. Figure 3.7 gives a plot of this function for $\lambda = 2$, which is the parameter value we shall use in the calculations below. We wish to show that Rabin's paradox can be explained by using reference dependence and loss aversion only, so we do not




Figure 3.7 Plot of (3.30) for the case of $\lambda = 2$.

invoke non-linear probability weighting (but see Wakker (2005) who shows that the paradox can also be explained with non-linear probability weighting).

1. Let us first compare the two prospects in incremental form $L_0 = (0, 1)$ and $L_1 = (-10, 0.5; 11, 0.5)$ when the decision maker follows PT. The preference $L_1 \prec L_0$ (see Assumption A1 in Section 1.5.4) implies that

$$V(-10, 0.5; 11, 0.5) < V(0) = 0, \quad (3.31)$$

where V is the value function under PT. Using (3.30) and (3.31), we get

$$-(1 - e^{-.010}) + \frac{1}{2}(1 - e^{-.011}) < 0 \Leftrightarrow -4.4803 \times 10^{-3} < 0,$$

which is true. Thus, at all levels of initial wealth, a decision maker who follows PT prefers the prospect L_0 over L_1 .

2. Let us now compare the two prospects $L_1 = (-10, 0.5; 11, 0.5)$ and $L_2 = (-100, 0.5; \infty, 0.5)$ when the decision maker follows PT. Rabin's paradox will arise if one now observes the preference $L_2 \prec L_1$. We now show that the preference $L_2 \prec L_1$ cannot arise. Using (3.30), we get that

$$\begin{aligned} L_2 \prec L_1 : -(1 - e^{-0.1}) + \frac{1}{2}(1 - e^{-\infty}) &< -(1 - e^{-.010}) + \frac{1}{2}(1 - e^{-0.011}) \\ \Leftrightarrow 0.40484 &< -4.4803 \times 10^{-3}, \end{aligned}$$

which is false. Thus, a decision maker who follows PT prefers the prospect L_2 over L_1 , hence, Rabin's paradox does not arise for a decision maker who follows PT.

The Rabin paradox has elicited several responses. For a behavioral counter response, which argues that these responses are unsatisfactory, see Wakker (2005, 2010). Palacios-Huerta and

Serrano (2006) and Watt (2002) argued that the preferences $L_1 \prec L_0$ need not hold at all wealth levels (see Assumption A1 in Section 1.5.4) if the decision maker follows EU. They do so by estimating the parameters of the power form of utility that allow for Assumption A1. They find that some of the parameters are unreasonable. But this only illustrates the difficulty in accounting for plausible preferences under EU, a point that Rabin had already made.

Cox and Sadiraj (2006), Watt (2002), Rubinstein (2006), and Samuelson (2005) propose to solve Rabin's paradox, not by abandoning EU but by defining utility over income. But doing so immediately imposes reference dependence because income arises over and above current wealth and so this simply assumes the reference point to be current wealth. For this reason, Wakker (2005, 2010) argues that these authors provide the same explanation as Rabin (2000a).

Consider, for instance, the explanation by Cox and Sadiraj (2006). They assume the following utility function.

$$v(y) = \begin{cases} 0.9y + 0.1 & \text{if } y < 1 \\ y^{0.9} & \text{if } y \geq 1 \end{cases} \quad (3.32)$$

If we did not subtract the wealth level w from all the outcomes, then the relevant lotteries are given in (3.28). The variable y would then be the final wealth level. A simple calculation shows that this gives rise to Rabin's paradox under EU (as in Section 1.5.4). However, now suppose that utility in (3.32) is defined over income levels. So essentially, we subtract wealth, w , from all outcomes. This gives rise to the relevant lotteries in incremental form as in (3.29), exactly as under PT. Then, under EU, with the utility function in (3.32), the preference $L_1 \prec L_0$ implies that

$$0.5(0.9(-10) + 0.1) + 0.5(11^{0.9}) < v(0) = 0.1.$$

But $0.5(0.9(-10) + 0.1) + 0.5(11^{0.9}) = -0.12264 < 0.1$. Hence, the original preference $L_1 \prec L_0$ holds. Yet, a decision maker who has the utility function of income given in (3.32) will accept the lottery $(-1680, 0.5; 3412, 0.5)$ over the lottery $(0, 1)$, so $L_1 \prec L_0$ does not hold. But this means that Assumption A1 in Section 1.5.4 does not hold and Rabin's paradox does not arise in this case. Cox and Sadiraj propose an amendment to the expected utility function in which wealth directly enters the utility functions and attitudes to risk are dependent on initial wealth.

The Rabin paradox and the Cox and Sadiraj explanations are both results of the form, if A then B, and both results are technically sound. As we point out in the introduction to Part I, reference dependence is not inconsistent with the axioms of EU, so the Cox and Sadiraj explanation is a feasible one. However, once again, it illustrates the difficulty in accounting for plausible preferences under EU without imposing reference dependence. Furthermore, simply imposing reference dependence in EU without a theory of how the behavior of decision makers differs in the domains of gains and losses is not entirely satisfactory either.

Neilson (2001) and Safra and Segal (2008) show that Rabin's puzzle also applies to other decision theory models based on final wealth levels. However, it is not clear that these two papers are based on empirically plausible assumptions; see Wakker (2010, p. 245).

3.7 Goals as reference points

Under EU, one typically derives utility from final levels of outcomes such as wealth, or if there are reference effects, then there is no difference in behavior in the domain of gains and losses

(Remark 1.1). Thus, setting wealth targets or other goals should not, in the absence of auxiliary assumptions, have any influence on behavior. Yet, as we shall see below in Section 3.7.2, targets and goals influence human actions. The reason is that goals may serve as reference points and the decision maker may feel that he is in the domain of losses if he falls short of the goal. This may induce the decision maker to exert even greater effort to achieve his goal. These effects have been found for human and non-human subjects and for cognitive and non-cognitive tasks.³⁵ Such reasoning is consistent with PT but not necessarily with other theories that lack a reference point.

3.7.1 A theory of goals in prospect theory

We first organize these ideas into a formal model, using unpublished notes by al-Nowaihi and Dhami.³⁶ An individual can exert effort $e \in [0, \bar{e}]$ at a cost $c(e)$ to produce output according to the production function $y(e)$. The properties of these functions are:

$$c(e) \geq 0, c(0) = 0, c'(0) \geq 0, c''(e) > 0; y(e) \geq 0, y(0) = 0, y'(e) > 0, y''(e) \leq 0. \quad (3.33)$$

Thus the cost function is positive, increasing, and convex, while the production function is positive, increasing, and concave.

The exogenously given *goal* of the individual is to produce a level of output $y = g \geq 0$. We assume that:

$$\text{for each } g \geq 0, y(e) = g \text{ for some level of } e. \quad (3.34)$$

The utility function of the individual, v , is separable in output and effort, and is defined separately for the domain of gains and losses, depending on whether the outcome $y(e)$ is larger or smaller relative to the goal, g .

$$v = \begin{cases} v_+(e, g) = u(y(e) - g) - c(e) & \text{if } y(e) \geq g \\ v_-(e, g) = -\lambda u(g - y(e)) - c(e) & \text{if } y(e) < g \end{cases}; \lambda > 1. \quad (3.35)$$

We assume that the functions u, y, c have continuous second order partial derivatives. We make the following standard assumptions.

$$u(x) \geq 0 \text{ for } x \geq 0; u(0) = 0; u'(x) > 0 \text{ for } x > 0; u''(x) < 0 \text{ for } x > 0. \quad (3.36)$$

Further assumptions on the function u are as follows.

$$u'(x) \rightarrow \infty, \text{ as } x \rightarrow 0; u'(x) \rightarrow 0, \text{ as } x \rightarrow \infty. \quad (3.37)$$

We consider separately, the domain of gains and losses.

³⁵ See, for instance, Locke and Latham (1990) for a wide-ranging review of the laboratory and field studies in support of these stylized facts. For a more updated list of references see Heath et al. (1999).

³⁶ This model was constructed circa 2010 and has been used by one of our Ph.D. students in his empirical research.

I. Domain of gains: In the domain of gains, $y(e) \geq g$. Using (3.35), we get

$$\frac{\partial v_+}{\partial e} = y'(e) u'(y(e) - g) - c'(e), \quad (3.38)$$

$$\frac{\partial^2 v_+}{\partial e^2} = y''(e) u'(y(e) - g) + [y'(e)]^2 u''(y(e) - g) - c''(e) < 0. \quad (3.39)$$

Thus, v_+ is continuous and strictly concave in e . Since $e \in [0, \bar{e}]$ is a compact interval, so v_+ has a unique (local and global) maximum. Let $e_+^*(g)$ be the (unique) function that assigns to each $g \geq 0$, the unique local maximum, e_+^* , corresponding to g in the domain of gains. In the domain of gains, $y(e) \geq g$, and from (3.37), (3.38), a corner solution cannot exist, hence, $e_+^* > 0$. We summarize these results in the next lemma.

Lemma 3.2 Given $g \geq 0$, $\frac{\partial v_+}{\partial e} = 0$ at a unique maximum $e_+^* > 0$ and $y(e_+^*) > g$.

Implicitly differentiating (3.38) we get

$$\frac{\partial^2 v_+}{\partial e \partial g} = -y'(e) u''(y(e) - g) > 0, \quad (3.40)$$

and, hence,

$$\frac{de_+^*(g)}{dg} = - \left[\frac{\partial^2 v_+}{\partial e \partial g} / \frac{\partial^2 v_+}{\partial e^2} \right]_{e_+^*(g)} > 0. \quad (3.41)$$

II. Domain of losses: We now turn to an analysis of the domain of losses, $y(e) < g$. Using (3.35), $v_-(e, g) < 0$, and we would like to maximize this function. We get

$$\frac{\partial v_-}{\partial e} = \lambda y'(e) u'(g - y(e)) - c'(e). \quad (3.42)$$

$$\frac{\partial^2 v_-}{\partial e^2} = \lambda y''(e) u'(g - y(e)) - \lambda [y'(e)]^2 u''(g - y(e)) - c''(e). \quad (3.43)$$

$$\frac{\partial^2 v_-}{\partial e \partial g} = \lambda y'(e) u''(g - y(e)) > 0.$$

Unlike the domain of losses, we cannot sign the second derivative of v_- . Using (3.37), we get

$$\lim_{y(e) \uparrow g} \frac{\partial v_-}{\partial e} = \infty. \quad (3.44)$$

Since the domain of losses is defined by $y(e) < g$, and since y is continuous, it follows that the domain of losses is open, i.e., v_- is not defined over a compact domain. In conjunction with the fact there is no guarantee that v_- is concave or convex, the analysis of the optimal solution in the domain of loss is less direct.

Let e_-^* be a local maximum in the domain of losses. It then follows that, necessarily, $\left[\frac{\partial v_-}{\partial e} \right]_{e=e_-^*} = 0$. Since $\frac{\partial v_-}{\partial e}$ is continuous, it follows that the set of all e such that $\frac{\partial v_-}{\partial e} = 0$ is closed. Finally, since y is strictly increasing, the set of feasible effort levels is bounded (by 0 below and

$y^{-1}(g)$ above). Thus the set of all effort levels, e , such that $\frac{\partial v_-}{\partial e} = 0$ is closed and bounded, hence compact. It follows that, if this set is non-empty, then it contains a maximum. Choose the largest effort level associated with this maximum, and call it $e_-^*(g)$. If, in addition, $\left[\frac{\partial^2 v_-}{\partial e^2}\right]_{e=e_-^*(g)} < 0$, then

$$\frac{de_-^*(g)}{dg} = - \left[\frac{\partial^2 v_-}{\partial e \partial g} / \frac{\partial^2 v_-}{\partial e^2} \right]_{e=e_-^*(g)} < 0. \quad (3.45)$$

The existence of $e_-^*(g)$ is not guaranteed but $e_+^*(g)$ always exists. So we have two cases.

(1) For each $g \geq 0$, let $e^*(g) = e_+^*(g)$, if no $e_-^*(g)$ exist.

(2) Suppose one, or more, optimal effort levels exist in the domain of losses. By definition, each of these optimal effort levels give the same value of the objective function, hence, we need some tie breaking rule, which, given our convention, is to choose the highest of these, $e_-^*(g)$. In this case, let $e^*(g) = e_+^*(g)$ if $v_+(e_+^*(g), g) \geq v_-(e_-^*(g), g)$. Otherwise, set $e^*(g) = e_-^*(g)$.

From (3.41) we know that $\frac{de_+^*(g)}{dg} > 0$, and from (3.45) we know that $\frac{de_-^*(g)}{dg} < 0$. Define g^* , such that

$$v_+(e_+^*(g^*), g^*) = v_-(e_-^*(g^*), g^*),$$

in which case, given our convention, we would choose $e^*(g^*) = e_+^*(g^*)$. The next proposition shows that g^* is an optimal goal if the objective is to maximize the output y . The proof is left to the reader.

Proposition 3.33 *There exists a unique goal, g^* , such that the following holds.*

- (a) For $g < g^*$, $v_+(e_+^*(g), g) > v_-(e_-^*(g), g)$ and so $e^*(g) = e_+^*(g) < e^*(g^*)$.
- (b) For $g > g^*$, $v_+(e_+^*(g), g) < v_-(e_-^*(g), g)$ and so $e^*(g) = e_-^*(g) < e^*(g^*)$.
- (c) g^* is the goal that maximizes output, $y(e)$, when one restricts attention to the set of effort levels that maximize utility of the individual.

A number of lessons can be drawn from Proposition 3.33.

1. If the objective is to maximize the output produced by an individual, then there exists an optimal goal, g^* .
2. For goals in the interval $[0, g^*]$, the optimal effort level of the individual is $e_+^*(g)$, which places the individual in the domain of gains, so $y(e_+) - g \geq 0$. Thus, the individual exceeds the goals that he sets.
3. Over the range $[0, g^*]$, $\frac{de_+^*(g)}{dg} > 0$, so the individual responds to an increase in goals by increasing effort.
4. For any goal that is too ambitious, i.e., in the range, $g > g^*$, the individual flips from being in the domain of gains to being in the domain of losses, i.e., $y(e_-) < g$. Indeed, $e_-(g)$ rather than $e_+(g)$ is optimal for goals in this range. Thus, the individual falls below the goal that he sets. Furthermore, because $\frac{de_-^*(g)}{dg} < 0$, the individual (at least locally in the neighborhood of $e_-^*(g)$) responds to an increase in goals by putting in even less effort.

Table 3.2 Some simulations for optimal goals and effort levels.

g	$e_+^*(g)$	$e_-^*(g)$	$e^*(g)$
0	0.908 66	n.a.	0.908 66
1	1.116 2	n.a.	1.116 2
2	2.000 3	n.a.	2.000 3
$g^* = 2.391113282$	2.391 2	1.956 4	2.391 2
2.5	2.5+	1.891 8	1.891 8
3	3+	1.762	1.762
4	4+	1.6528	1.6528

Source: unpublished notes by al-Nowaihi and Dhami.

Remark 3.9 (*Predictions under EU and RDU*): Under EU and RDU, the utility function is defined over final levels of wealth (in this case $y(e)$) and is independent of g .³⁷ So goals are not predicted to have any effect on the effort levels of individuals. If goals do turn out to have an effect, then this finding would be inconsistent with EU and RDU.

We now report a few simulations in Table 3.2 that illustrate the conclusions of the model. Let us assume that the production function is linear, i.e., $y = e$. Let the utility function for gains and losses be defined respectively by the following.

$$v_+(y, g) = (y - g)^\beta - \frac{1}{2}y^2, y \geq g,$$

$$v_-(x, g) = -\lambda (g - y)^\beta - \frac{1}{2}y^2, \lambda > 1, x < g.$$

For simulation purposes, we use the values $\beta = 0.9$, $\lambda = 2$. We report for each value of the goal, g , the corresponding values of $e_+^*(g)$, $e_-^*(g)$, and $e^*(g)$. The reader can verify observations 1–4 above.

3.7.2 The evidence on goals

We now consider the evidence on goals, based on experimental results from Heath et al. (1999). The results in Table 3.2 were results based on the calculation of optimal effort levels. The experiments, on the other hand, present hypothetical information that is not necessarily based on a calculation of the optimal solution.

The number of subjects in an experiment is denoted by N . One implication of goals as reference points is that it is possible that an individual with a higher goal performs better than an individual with a lower goal, yet feels less happy. This is supported by the following problem.

Problem 1. [$N = 60$] Alice typically scores around 80 on weekly 100-point quizzes in one of her classes. One week she sets a goal of scoring 90. She scores 87. Betty typically scores around 80 on weekly 100-point quizzes in one of her classes. One week she decides to do her best. She scores 83. Overall, who will be more satisfied with her performance?

³⁷ Even if we allowed for reference dependence under EU and RDU, then identical behavior in gains and losses (Remark 1.1) ensures that this is not a fruitful exercise.

Alice's performance is better than Betty's but because of her higher goal, a statistically significant majority of people respond that Alice should be less satisfied with her performance. This indicates that goals partition the domain of outcomes into gains and losses akin to the role of reference points in prospect theory. This result, and the next experiment which illustrates the role of loss aversion, is in conformity with the framework of our theoretical model.

Problem 2. [$N = 48$] Sally and Trish both follow workout plans that usually involve doing 25 sit-ups. One day, Sally sets a goal of performing 31 sit-ups. She finds herself very tired after performing 35 sit-ups and stops. Trish sets a goal of performing 39 sit-ups. She finds herself very tired after performing 35 sit-ups and stops. Who is experiencing more emotion?

Of the subjects 29% answer, Sally, while the remaining 71% answer, Trish. One plausible interpretation is that because Trish falls below the goal of 39 she experiences more emotion, on account of loss aversion.

Problem 3. [$N = 73$] Charles and David both follow workout plans that usually involve doing 25 sit-ups. One day, Charles sets a goal of performing 30 sit-ups. He finds himself very tired after performing 34 sit-ups and, at most, has the energy to perform 1 more. David sets a goal of performing 35 sit-ups. He finds himself very tired after 34 sit-ups and, at most, has the energy to perform 1 more. Who will work harder to perform the 35th sit-up?

Of the respondents 82% give the answer as David, which supports the prediction of PT. This can be seen from (3.42). As $y(e) \rightarrow g$, the marginal returns to effort, $\lambda y'(e) u'(g - y(e))$ increase rapidly and in the limit, (3.37) ensures that this marginal return is infinite. In other words, by doing just one more sit-up David can avoid falling into the domain of losses, so he should work harder.

Problem 4. [$N = 76$] Charles and David both follow workout plans that usually involve doing 25 sit-ups. One day, Charles sets a goal of performing 30 sit-ups. He finds himself very tired after performing 42 sit-ups and, at most, has the energy to perform one more. David sets a goal of performing 40 sit-ups. He finds himself very tired after 42 sit-ups and, at most, has the energy to perform one more. Who will work harder to perform one final sit-up?

Of the respondents 69% give David as the answer.

Problem 5. Same as problem 4, with 42 sit-ups replaced by 28 sit-ups

Now 86% of the respondents give Charles as the answer.

In problem 4, both individuals exceed their goal. However, David has exceeded his goal by less than Charles. Since the effort level is held fixed for both (42 sit-ups), the marginal benefit for David of performing an extra sit-up is greater (because $u'(x) > 0$ $u''(x) < 0$). People exert greater effort when they are closer to their goals. The key is to invoke diminished sensitivity to gains and losses as one moves away from the reference point, a feature that is already incorporated into various versions of PT. The same reasoning explains the responses to Problem 5.

There are other emerging effects of goals that we have not touched on. For instance, there appears to be a strong effect of non-binding goals on performance and this effect appears to be much stronger for men relative to women (Smithers, 2015).

3.8 Why is it so hard to find a taxi on a rainy day in New York?

Consider the standard model of labor supply. Suppose that there is a small current transitory increase in wage. The decision maker would then smooth the transitory wage change over all future time periods so that income effects are negligible. In this case, we can concentrate solely on the substitution effects. Since wages are temporarily high, and so the opportunity cost of leisure is high, the individual engages in intertemporal substitution of labor, working relatively harder today. Hence, we can state the following well-known result.³⁸

Remark 3.10 *Transitory increases in the real wage are predicted, under neoclassical economic theory, to produce an increase in hours worked.*

In the labor economics literature, there is little empirical consensus on the magnitude and sign of the labor supply elasticity. There are two main problems with estimating labor supply elasticities. First, actual wage changes are often not transitory in nature. Second, individuals often have limited control over hours worked. The modern literature has tried to address these issues.

Camerer et al. (1997) ask: *Why is it harder to find a taxi on a rainy day in New York?* The situation for cab drivers in New York is particularly interesting. First, it is possible to identify transitory wage changes. Second, by deciding when to quit for the day, cab drivers make a real “hour’s worked” decision. The rate that cab drivers can charge per mile is fixed by law. However, on busy days, drivers have to spend less time searching for customers and, so, their (effective) real wage rate is higher.

Let h_{it} and Y_{it} denote hours worked and total daily earnings by driver i at date t . Hourly wage w_{it} is computed as Y_{it}/h_{it} . Let X_{it} denote a set of other independent variables; some are driver-specific, and others such as rain, temperature, and shift dummies, control for the difficulty of driving conditions and, hence, the opportunity cost of leisure. Camerer et al. (1997) estimate regression equations of the following form.

$$\log h_{it} = \beta_0 + \beta_1 w_{it} + \beta_2 X_{it} + \varepsilon_{it}, \quad (3.46)$$

where β_1 is a measure of labor supply elasticity and β_2 is a vector of coefficients. For the three sets of data in the sample, the inclusion of driver fixed effects gives the estimated labor supply elasticities to be respectively, -1.86 , -0.618 , -0.355 . These estimates suggest the presence of a *downward sloping labor supply curve* over the relevant range.

The construction of wages as $w_{it} \equiv Y_{it}/h_{it}$ raises two sets of issues. First, in the econometric model in (3.46), the labor supply decision responds to the average wage over the day, rather than fluctuations in the wage over the day. If the wage is negatively correlated over the day, then the intertemporal labor substitution model predicts that workers will work harder during the early part of the day when wages are relatively higher, and quit early. However, the authors find a positive autocorrelation in wages over the day, ruling out this problem. Second, h_{it} appears on the LHS as well as on the RHS because $w_{it} \equiv Y_{it}/h_{it}$. Hence, any model mis-specification or measurement errors will bias the estimates of β_2 downwards.

³⁸ For a simple, tractable, two-period model and discussion of these issues, see, for instance, Blanchard and Fischer (1989, section 7.2).

To control for these issues, the authors use the average daily wage rate of other drivers as an instrument. However, they find that this makes the estimate of β_1 even more negative. Another econometric issue is an identification issue. Are the authors picking out points along a downward sloping labor demand curve because of relatively greater variation in the labor supply curve? The authors report that they conducted personal interviews with cab drivers that led them to believe that the marginal cost of driving does not alter in a manner that might have an effect on the wage. However, the authors do not directly introduce an instrument for labor supply.

How does one explain the estimated labor supply elasticities? The authors favor an explanation that is based on reference dependence. Cab drivers have a daily target level of earnings. The target earnings level could be some fixed number (say, \$150 per day) or be linked to the costs of leasing the cab (such as, twice the daily leasing fee). Drivers dislike falling below the daily target and the probability of quitting for the day increases as they approach the target.

Explanations based on daily income targeting are supported by interviews that the authors conducted with the fleet managers. For instance, six fleet managers chose the following response: “Drive until they make a certain amount of money.” Five fleet managers chose: “Fixed hours.” Only one fleet manager chose a response that is consistent with the intertemporal substitution model: “Drive a lot when doing well; quit early on a bad day.” Daily income targeting is not a feature of standard labor economics models. However, it appears to be a sensible heuristic and is reminiscent of the short horizon of portfolio holders (1 year) in the Benartzi and Thaler (1995) model and of the literature on goals and targets. A daily income target is also simple and computationally easy, say, relative to a monthly or a weekly income target.

Does experience make a difference? For instance, experienced drivers might realize that it pays to work longer on busy days. It is possible to split up the cab drivers’ data into experienced drivers (greater than 3–4 year’s experience), and inexperienced drivers. The elasticity of labor supply for inexperienced drivers is generally close to -1 (which is consistent with a daily income target). The corresponding elasticity for experienced drivers is much larger, and positive.

We now present a simple model of reference dependent preferences in labor supply, due to Farber (2008). Section 2.8.2 considers the optimality of income targeting in a similar model with Köszegi–Rabin preferences. Suppose that Y is income, $\bar{Y}$ is reference income, h is hours worked, and w is the wage rate (such that $w = Y/h$). Suppose that the reference dependent utility function is given by

$$U(Y, \bar{Y}) = \begin{cases} (Y - \bar{Y}) - \frac{\theta}{1+\nu} h^{1+\nu} & \text{if } Y \geq \bar{Y} \\ -\lambda(\bar{Y} - Y) - \frac{\theta}{1+\nu} h^{1+\nu} & \text{if } Y < \bar{Y} \end{cases}, \quad (3.47)$$

where λ is the parameter of loss aversion, ν is the inverse of the elasticity of labor supply, θ is a parameter related to disutility of hours worked. Noting that $Y \equiv wh$, the conditions which separate the domain of income in (3.47) into losses and gains are

$$\begin{aligned} Y > \bar{Y} &\Leftrightarrow w > \bar{Y}/h \\ Y < \bar{Y} &\Leftrightarrow w < \bar{Y}/h \end{aligned} \quad (3.48)$$

When $Y > \bar{Y}$, a simple calculation gives the interior solution to the hours worked,

$$h = \left(\frac{w}{\theta} \right)^{1/\nu}. \quad (3.49)$$

Table 3.3 A summary of the solution to the simple labor supply problem in Farber (2008).

Wage	Hours	Elasticity
$w < w_1 \equiv \left(\frac{\theta \bar{Y}^\nu}{\lambda}\right)^{1/(1+\nu)}$	$h = \left(\frac{\lambda w}{\theta}\right)^{1/\nu}$	$1/\nu$
$w_1 \leq w \leq w_2 \equiv \left(\theta \bar{Y}^\nu\right)$	$h = \bar{Y}/w$	-1
$w_2 < w$	$h = \left(\frac{\lambda w}{\theta}\right)^{1/\nu}$	$1/\nu$

Analogously, when $Y \equiv wh < \bar{Y}$, from the second row of (3.47), a simple calculation gives the interior solution to the hours worked as

$$h = \left(\frac{\lambda w}{\theta}\right)^{1/\nu}. \quad (3.50)$$

When $Y \equiv wh = \bar{Y}$, the solution can be directly found as follows

$$h = \bar{Y}/w. \quad (3.51)$$

From (3.48), (3.49), (3.50), (3.51), we can summarize the solution in Table 3.3.

The labor supply curve has the following shape. For wage rates that are less than w_1 and those that are greater than w_2 , the labor supply curve is upward sloping with elasticity of labor supply given by $1/\nu$. However, for intermediate ranges of wage rates, between w_1 and w_2 , the individual finds that his income is at the reference point, so, as in (3.51), it is optimal simply to offset, one for one, an increase in the wage rate by the hours worked. Thus, the labor supply elasticity in the intermediate range is -1 .

Having fixed the basic ideas, we now give a brief overview of the empirical evidence. Oettinger (1999) examines the labor supply decision of vendors in baseball stadiums. The vendors have a choice of working during a particular game by giving prior notice of their decision to their employers. Vendors receive a commission for their sales during the game. The effective wage depends on the attendance at a particular game; greater attendance reduces the time needed to search for the next customer. The level of attendance is a “demand shock,” while the number of other vendors who have chosen to work at a particular game is the relevant “supply shock” for a vendor. Using instrumental variables for the attendance at any game, they find positive and statistically significant wage elasticities (ranging from 0.53 to 0.64). This evidence does not contradict either the standard neoclassical model (see Remark 10 above) or the reference dependence model because most of the equilibrium points could be along the positive sloping part of the labour supply curve.

Fehr and Goette (2007) look at the labor supply decision of bicycle messengers in Zurich who are paid a commission for each successfully delivered message. Bicycle messengers have to work a given number of fixed shifts during a month, but can voluntarily sign up for more shifts. There are two groups of messengers (groups A and B). Group A members are first offered a 25% increase in the commission rate for a period of one month after which they revert back to the old rate. In the following month, group B members are offered the same temporary increase in the commission rate for a month. Fehr and Goette (2007) find that the probability of working on a given day

is higher for the group that is currently a recipient of the increase in the commission rate. The wage elasticities are positive and have similar magnitudes to those in the Oettinger (1999) study, ranging from 0.72 to 0.82.

Fehr and Goette (2007) then regress the log of revenues earned by a messenger in each shift on a dummy for the treatment group that receives the 25% increase in commission rates. They also use other control variables that take account of daily fixed effects and individual fixed effects. They find that the revenues within a shift are lower for the group that has been offered the commission increase, and the implied labor elasticities are negative, ranging from -0.23 to -0.29 . These results are consistent with the reference dependence model.

The Fehr–Goette results are not only consistent with those of Camerer et al. (1997), they also control for some of the objections that were raised against those results. For instance, the source of the variation in wage (25% increase for the control group) is exogenous, so the issue of instruments for labor supply is not a concern. Similarly, messengers sorted themselves into one of the two groups (A or B) prior to the announcement of the increase in the wage, hence, selection issues are controlled for. In an insightful extension, the authors perform a follow-up experiment, designed to elicit the loss aversion parameter of the bicycle messengers. It is found that the effort elasticity is negative and statistically significant only for those messengers who also exhibit loss averse behavior (others do not exhibit an elasticity that is statistically different from zero). This provides much stronger confirmatory evidence of reference dependent preferences in labor supply.

Farber (2005) examines the within-day labor supply decision of New York cab drivers. He uses a hazard model of quitting to test the prediction of the reference dependence model that higher wages during the early part of the day induce earlier quits. In the simplest model, Farber finds that holding fixed the hours worked, the probability of quitting early increases when the earnings for the day up to that point are higher. This is consistent with the reference dependence model. Next, Farber adds controls for driver fixed effects, which alters the result of the simple model; the probability of quitting now does not depend on earnings so far during the day. However, one of the controls used is clock hours, which is the most predictable source of variation in hourly wages; average hourly wages are positively correlated over the day, peaking at evening rush time. This control eliminates an important source of information that could have identified the reference dependence model; for a discussion of these points, see Goette et al. (2004).

Goette and Huffman (2003), like Farber (2005), consider the within-day labor supply decisions of bike messengers based on US and Swiss data. They test the following prediction of the reference dependence model. Suppose that there is a windfall gain in the morning (e.g., a bike messenger gets lucky). Then, reference dependence should predict an increasing effort, followed by decreasing effort. The reason is that when messengers are below the target earnings, they have an increasing incentive to provide effort. However, once this target is exceeded, their incentive to provide effort, falls. In order to isolate just the effect of luck in the morning, they introduce messenger fixed effects, fixed effects to control for the effects of firm, day of the week, and clock hour. The residual variation in earnings of the messenger is on account of luck. The authors verify the prediction of reference dependence that they start out to prove.

Some of the concerns with Farber (2005) that prevented reference dependence to be identified are addressed in Farber (2008) in a model of reference dependent preferences and a reference point that is stochastic. The evidence supports a finding of reference dependence. The model in which income is targeted, performs much better than the standard neoclassical model. However, the reference point is found to fluctuate too much from day to day. The author concludes that this limits the power of reference dependence theory to explain labor supply.

Crawford and Meng (2011) raise two objections to the Farber (2008) results. (1) Farber finds that the probability of stopping increases when the income target is reached. While the individual decisions seem to have these kinks, the aggregate relation between stopping probability and income is too smooth. How can these results be reconciled? (2) The randomness in reference points is large but imprecisely estimated.

Crawford and Meng (2011) propose a model of reference dependent preferences that relies on the Köszegi–Rabin framework (see Section 2.8). In this model (and unlike Faber’s model) there are targets for both hours worked and income. These targets are determined endogenously as rational expectations equilibria. The rational expectations of the wage is operationalized by simply taking the average wage. As before, we denote hours worked by h , income by Y , the actual hourly wage rate by w and reference income by $\bar{Y}$. The new variables are average or expected wage (w^e), and reference numbers of hours worked ($\bar{h}$). Conditional on the reference points for income and hours ($\bar{Y}, \bar{h}$), which are themselves endogenous, the utility of the individual is given by a weighted combination of *consumption utility* and *gain–loss utility*, as follows.

$$U(Y, h, \bar{Y}, \bar{h}) = (1 - \eta) [u(Y) + v(h)] + \eta R(Y, h, \bar{Y}, \bar{h}); \eta \in [0, 1]. \quad (3.52)$$

In (3.52), consumption utility (first term) is additively separable in income and hours and given the weight $1 - \eta \geq 0$. There is positive marginal utility from consumption $u'(Y) > 0$, negative marginal utility from hours worked, $v'(h) < 0$, and $u(Y)$, $v(h)$ are both concave. Gain–loss utility, denoted by R , is the second term in (3.52), and it has weight $\eta \in [0, 1]$. This term is defined by partitioning all (Y, h) pairs into four domains. The income domain and the hours domain are both in gains when $Y > \bar{Y}$ and $h < \bar{h}$; both in loss when $Y < \bar{Y}$ and $h > \bar{h}$; one in loss and the other in gain if either of $Y < \bar{Y}$ and $h < \bar{h}$ or $Y > \bar{Y}$ and $h > \bar{h}$ holds. Formally

$$R = \begin{cases} [u(Y) - u(\bar{Y})] + [v(h) - v(\bar{h})] & \text{if } Y > \bar{Y}, h < \bar{h} \\ \lambda [u(Y) - u(\bar{Y})] + \lambda [v(h) - v(\bar{h})] & \text{if } Y < \bar{Y}, h > \bar{h} \\ \lambda [u(Y) - u(\bar{Y})] + [v(h) - v(\bar{h})] & \text{if } Y < \bar{Y}, h < \bar{h} \\ [u(Y) - u(\bar{Y})] + \lambda [v(h) - v(\bar{h})] & \text{if } Y > \bar{Y}, h > \bar{h} \end{cases}. \quad (3.53)$$

When drivers decide to work an extra hour, they know that their earnings are random. To simplify, it is assumed that drivers simply focus on the expected or average wage, w^e , and use it to compute their hourly earnings. w^a is determined by rational expectations. The objective of the driver is to maximize (3.52), subject to the budget constraint

$$Y = w^e h. \quad (3.54)$$

We motivate a diagrammatic solution. In each of the four domains listed in (3.53), the marginal rate of substitution (MRS) between Y and h is computed using (3.52) and (3.53); Table 3.4 summarizes these calculations.

Since the coefficient of loss aversion $\lambda > 1$, we get that $1 + \eta(\lambda - 1) > 1$. From Table 3.4, when both income and hours are either in the domain of gains or losses, then the MRS is given as in the standard neoclassical model. However, in the other two cases, when both outcomes are in different domains, this is no longer true. When income is in gains and hours are in losses, then

Table 3.4 Summary of the marginal rates of substitution between Y and h in various domains.

Income/Hours	Gain ($h < \bar{h}$)	Loss ($h > \bar{h}$)
Gain ($Y > \bar{Y}$)	$\frac{-v'(h)}{u'(Y)}$	$\frac{-v'(h)}{u'(Y)} (1 + \eta(\lambda - 1))$
Loss ($Y < \bar{Y}$)	$\frac{-v'(h)}{u'(Y)} \frac{1}{1 + \eta(\lambda - 1)}$	$\frac{-v'(h)}{u'(Y)}$

the MRS is magnified (i.e., the indifference curves are steeper). When income is in loss and hours in gain the MRS shrinks (indifference curves are flatter). Equilibrium labor supply can be found from the usual condition that equates MRS with the marginal rate of transformation, which from (3.54) is w^e , hence, $MRS = w^e$. Denote the actual wage of the taxi driver on trip by w^a .

Lemma 3.3 (a) Suppose that $w^a > w^e$, i.e., the actual wage is greater than the expected wage. Then the income target, $\bar{Y}$, is reached first, followed by the hours target, $\bar{h}$.
 (b) Conversely, if $w^a < w^e$, then the hours target is reached first.

Proof: By definition, on any trip, $w^a = Y/h$. Since $w^a > w^e$, we get $Y/h > w^e$. On account of rational expectations, $\bar{Y} = hw^e$, i.e., the endogenous reference income equals the expected income. Suppose that the hours target is reached, i.e., $h = \bar{h}$, then from $\bar{Y} = hw^e$ we get that

$$h = \bar{h} = \frac{\bar{Y}}{w^e}. \quad (3.55)$$

From (3.54), actual income is given by $Y = w^a h$. Substituting (3.55) we get

$$Y = w^a \frac{\bar{Y}}{w^e}.$$

But since $w^a > w^e$ we get that $Y > \bar{Y}$. Hence, when the hours target is reached, the income target is already exceeded. Part (b) of the lemma is analogously derived. ■

Figure 3.8 gives insights into the equilibrium in (Y_t, h_t) space, where the subscript t denotes the time of the day during the trip. On the vertical axis, we measure Y_t (increasing upwards) and along the horizontal axis, h_t (increasing leftward). The space is truncated into four quadrants I, II, III, IV, based on the reference income and reference hours level, $(\bar{Y}, \bar{h})$, as the origin.

The driver begins at his endowment point of $(Y, h) = (0, 24)$, assuming that the total time endowment is 24 hours. He begins with an expected wage w^e and then moves in the northeast direction along the line $Y = w^e h$. There are two possibilities, $w^a > w^e$ and $w^a < w^e$, both shown in the diagram. Having completed any single trip, the driver begins another trip starting from the last level of (Y, h) , again moving in a northeast direction along the line $Y = w^e h$.³⁹ The key insight comes from Lemma 3.3; when $w^a > w^e$, the driver first reaches the income target, and when $w^a < w^e$ the driver first reaches the hours target.

³⁹ In practice, the budget line is random (because the wage is random) but for exposition, we consider that the driver takes account of only the wage w^e .




Figure 3.8 An illustration of the equilibrium in the Crawford–Meng (2011) model.

Source: Crawford and Meng (2011) with permission from the American Economic Association.

Consider first the case $w^a > w^e$. We know from Table 3.4 that in quadrants I and III, the slope of the indifference curves is identical in neoclassical theory and in prospect theory with targets $(\bar{Y}, \bar{h})$ (PTT). In quadrant II, the indifference curves are relatively steeper under PTT as compared to neoclassical theory, while in quadrant IV, these indifference curves are relatively flatter. Let us concentrate only on quadrants I and IV.

Suppose that in the neoclassical model, the equilibrium is at point 1 (it could also be at other points such as 2, 3, and 4). However, because in this quadrant the indifference curves under PTT are flatter, so the equilibrium under PTT could be at point 2, as shown in the figure. Starting at point 2, the individual will not like to enter quadrant I if the marginal rate of substitution, $\frac{-v'(h)}{u'(Y)}$ (see (1, 1)th element in Table 3.4) exceeds the opportunity cost of leisure, w^e . Should this happen, then equilibrium is at point 2 and the income target is met before the hours target (see Lemma 3.3).

Conversely, when $w^a < w^e$, the relevant budget constraint is the lower of the two budget constraints in Figure 3.8. In this case, starting from the endowment point $(Y, h) = (0, 24)$, an individual using PTT may similarly find it optimal to be at point 5 (this turns out to be the empirically relevant case). But in this case, the hours target is reached before the income target.

Crawford and Meng (2011) estimate the probability of stopping, as in Farber (2005), with two main differences. First, targets are not treated as latent variables but rather as rational expectations (proxied by average sample realizations). Second, the sample is split into two parts $w^a > w^e$ and $w^a < w^e$, i.e., depending on whether the realized income is lower or higher than the (proxied) expected income. The authors find clear evidence of reference dependence.

The probability of stopping is heavily influenced by realized income when $w^a > w^e$. Conversely, when $w^a < w^e$, the probability of stopping is heavily influenced by hours worked. The cases $w^a > w^e$ and $w^a < w^e$ have almost equal weight in the sample. In the first case, the wage elasticity of hours is strongly negative, while in the second it is zero. Thus, in the aggregate, wage elasticity is negative, as in Camerer et al. (1997).

3.9 Some implications of inverse S-shaped weighting functions

Figure 2.5 in Section 2.3 considered the following important implication of inverse S-shaped probability weighting functions under rank dependent models. Extreme events (or tail events in finance) are more salient in the sense that such events are accorded higher decision weights. This is one possible explanation for why positively skewed portfolios and assets with positively skewed returns may be more attractive to investors, hence, driving up their prices (Barberis and Huang, 2008).⁴⁰ A related implication of a preference for positively skewed portfolios is that individuals may hold portfolios that are under-diversified; see Mitton and Vorkink (2007) for confirmatory evidence.

Another implication of the increased salience of tail events is that the decision maker is more averse to negatively skewed returns, which accords well with the actual shape of the average over-all returns in the market (Fehr-Duda and Epper, 2012). This produces an even higher level of the equity premium, beyond that which can be explained by loss aversion alone (De Giorgi and Legg, 2012). This illustrates, once again, the importance of using the entire machinery of PT, particularly when one wishes to explain quantitative phenomena.

Salience of tail events may also explain why the volume of retirement annuities purchased is lower than the predicted volume in neoclassical theory. Suppose that the retirement age is 65. Hu and Scott (2007) present the problem as one in which if someone dies early, within a few years of 65 (small perceived probability of the event, hence, this is a tail event), then the annuity is not very useful. However, if tail events are overweighted, then the desirability of annuities falls, in line with the evidence.

A basic feature of the inverse S-shaped weighting functions is that low probabilities are overweighted. This has been used to explain why people choose low deductibles in insurance⁴¹ even when the high deductibles are several times cheaper in expected value. In Sydnor (2010), potential insurees are given a menu of choices between four different levels of home insurance deductibles, \$100, \$250, \$500, \$1000. The most common choice was \$500. However, in order to reduce the deductible from \$1000 to \$500, consumers typically paid an extra cost, in terms of the premium, of \$100. The claim rate of this group of consumers was 5%. Thus, consumers are willing to pay an extra \$100 to reduce the expected loss by \$25. The implied coefficient of relative risk aversion is in triple and quadruple digits, which is unreasonably high. Sydnor's preferred explanation is that individuals must overweight the probability of tail events significantly, which is consistent with inverse S-shaped weighting functions.

The empirical tests in Barseghyan et al. (2013) discriminate among the Köszegi–Rabin framework, Gul's theory of disappointment aversion, and non-linear probability weighting. They show that it is non-linear probability weighting that accounts for the observed choice of low deductibles. Overweighting of low probabilities also explains the favorite long-shot *bias* in betting markets (e.g., in horse racing). This takes the form of over-betting on horses which have a small probability of winning and under-betting on horses which have a high chance of winning.⁴²

⁴⁰ For related empirical evidence, see Ang et al. (2006), Boyer et al. (2010), Green and Hwang (2012), and Conrad et al. (2013).

⁴¹ In the event of a loss, the loss of the consumer is capped at the level of the deductible that is chosen.

⁴² See, for instance, Jullien and Salanié (2000) who show that prospect theory outperforms expected utility in explaining this phenomena. Snowberg and Wolfers (2010) are able to distinguish between the role of convex utility and probability weighting and find the latter to be the main source of the explanation.

A quick look at the literature suggests that loss aversion based explanations are more predominant in the literature as compared to those based on non-linear probability weighting. Fehr-Duda and Epper (2012) make a strong case for redressing this imbalance. They argue that (p. 589): “Probability dependence is the more fundamental of the two [loss aversion and non-linear probability weighting] because it is observed for all types of prospects: pure gain prospects as well as pure loss and mixed prospects. In contrast, loss aversion is effective for mixed prospects only. It is our impression that probability-dependent risk preferences have not received as much attention as loss aversion has, however.”

3.10 Contracts as reference points

Suppose that a buyer and a seller write a contract at an initial stage in which the future cannot be perfectly anticipated. Hence, any contracts that they write have to be necessarily incomplete. At the initial stage, both parties (or just one party) sink relation-specific investment. Once the relevant uncertainty is resolved at a future date, then the two parties can renegotiate the terms of the contract. Under the assumptions of Coasian bargaining, renegotiation allows them to achieve an efficient outcome.

Renegotiation entails sharing, ex-post, the benefits of the relation-specific investment. Anticipating this hold up problem in the renegotiation stage, at the initial stage, each party invests an inadequate amount. This highlights an important *transaction cost* of using the market. The solution to the problem is the allocation of ownership rights and *residual control rights* in such a way that facilitates *efficient investment* (Grossman and Hart, 1986; Hart and Moore, 1990). Indeed, one possible solution is the *formation of firms* and entering into *long-term contracts*.

This account of the theory of the firm is restrictive in scope, because ex-ante investment in non-contractible relation-specific investments is the main driver in the formation of firms. In contrast, we often observe the existence of firms when such factors are absent. This explanation also does not shed light on the *internal organization of the firm* and the explanation of *authority, hierarchy, and delegation*.

In this section, we consider an alternative view of the firm, one that is based on the idea of reference dependence in prospect theory.

In a reference dependence-based explanation of firms, contracts establish a *reference point of entitlements* for the contracting parties. Initial contracts are typically written among parties chosen from a competitive pool of buyers and sellers. Hence, each party to the contract is likely to believe that the terms of the contract capture the market beliefs about its entitlement. Competition is critical to establishing entitlements. Lack of competition does not establish entitlements, hence, contracts that are signed in non-competitive conditions may not serve as a reference point; see Fehr et al. (2009) for the empirical evidence.

At a future date, once ex-ante uncertainty is resolved, parties to the contract evaluate the current outcomes, relative to the ex-ante entitlements established by the initial contract. If the outcomes are inferior to the entitlements, parties feel *aggrieved* and may *shade* their ex-post performance. Both parties need not share the same beliefs about their respective entitlements. For instance, there could be a *self-serving bias*, which leads each party to confound fairness with its selfishly preferred outcome. This self-serving bias in entitlements is analyzed in the context of bargaining models in Chapter 12 (Babcock and Loewenstein, 1997).

Why is it typically difficult to specify ex-post performance in the ex-ante contract? Hart and Moore (2008) distinguish between *perfunctionary performance* and *consummate performance*.

The former refers to conformity with the *letter* of the contract (e.g., number of hours worked) and the latter with the *spirit* of the law (e.g., care and attention exercised during working hours). Perfunctionary performance is observable and verifiable to a third party, so it can be specified in the ex-ante contract. By contrast, consummate performance is observable to both parties, but not verifiable to a third party, so it cannot be written down in the contract. Hence, the ex-post actions of the parties are only *partially contractible*. If, ex-post, parties to a contract feel aggrieved, based on a comparison with their entitlements based on the initial contract, then they can lower their consummate performance. This is the sense in which *shading of performance*, to the detriment of the other parties to a contract, can take place. In this framework, there is no loss aversion and there is no renegotiation.

In a buyer–seller framework, shading by the seller can take the form of reducing quality, or not agreeing to reasonable requests by the buyer that greatly enhance the utility of the buyer at zero or very low cost to the seller. Buyers can also engage in shading by, say, delaying payments, engaging in discourteous behavior, and writing bad references for future potential buyers.

These elements of the theory can be used to determine the choice between *rigid contracts* and *flexible contracts*. A rigid contract, such as an *employment contract*, rigidly fixes the wage at the initial stage in a competitive environment, and allows no ex-post flexibility in altering the wage. Hence, ex-post, parties have no opportunity to be aggrieved. However, such a contract lacks the ex-post flexibility to adapt to changing states of nature. For instance, demand might turn out to be surprisingly low, but the initially fixed wage cannot be lowered. By contrast, a flexible contract offers ex-post flexibility in adapting to the realized state of the world. However, it also creates an opportunity for aggrievement. As one of the parties adapts to the state of the world, the other party may feel unfairly treated relative to the entitlements and retaliate by reducing its consummate performance. This trade-off governs the choice between the rigid contract and the flexible contract. Under certain conditions that we specify below, the rigid contract may be optimal.

3.10.1 A simple model

Hart and Moore (2008) consider the following simple model of a long-term relation between a buyer, B, and a seller, S. A competitive pool of buyers and sellers meet at date 0. A matched buyer–seller pair then sign a contract at date 0, which structures their interaction at date 1. Once the two parties are matched, the competitive pool of market participants reduce to a bilateral monopoly situation; this is Williamson’s well-known *fundamental transformation*. Both parties are risk neutral, all information is symmetric, there are no wealth constraints, and there are no non-contractible relation-specific investments.

There is potential uncertainty when parties sign contracts at date 0, but all uncertainty is resolved by date 1. At date 1, the parties act, in each state, as per the instructions in the date 0 contract. For instance, if the date 0 contract is a rigid contract, then parties make no active decisions. On the other hand, if the date 0 contract is a flexible contract, it might require one of the parties to the contract to choose a date 1 price, depending on which state of the world is realized. Trade then occurs and the two parties choose their consummate performance. In order to measure the degree of aggrievement, we make the simplifying assumption that each party feels entitled to the best outcome that can arise for it in the date 0 contract (say, on account of self-serving bias).

Denote by Π_B, Π_S , the set of gross payoffs that are contractually feasible and let $\bar{u}_j = \max \Pi_j$ be the maximum possible payoff to party j as dictated by the date 0 contract. Suppose that at date 1, the gross payoffs of the buyer and seller are, respectively, u_B and u_S . Let the buyer and the seller be indexed, respectively, by $j = B, S$. The aggrievement of player j is given by

$$a_j = \bar{u}_j - u_j, u_j \in \Pi_j. \quad (3.56)$$

Being aggrieved causes a drop in consummate performance that causes a monetary loss to the other party (effort shading). Denote the monetary loss caused by the buyer and the seller, respectively, to the other party, on account of aggrievement, by l_B and l_S . The respective utilities of the buyer and the seller, U_B and U_S , are given by

$$U_B = u_B - l_S - \max\{\theta a_B - l_B, 0\}, \quad (3.57)$$

$$U_S = u_S - l_B - \max\{\theta a_S - l_S, 0\}, \quad (3.58)$$

where $\theta \in (0, 1]$. The third term on the RHS in each of (3.57) and (3.58) shows that being aggrieved leads to a loss in utility, the extent of which depends on θ . However, causing a monetary loss to the other party shades this loss in utility; this is viewed as “just retribution” inflicted on the other party when one perceives oneself to be unfairly treated. However, shading can never be used to increase one’s own utility, although one can shade enough such that all aggrievement disappears (i.e., if $\theta a_B = l_B$ or $\theta a_S = l_S$). Each party simultaneously chooses l_j , $j = B, S$. Clearly, the utility maximizing choice of each party is $l_B = \theta a_B$ and $l_S = \theta a_S$. Substituting this solution in (3.57) and (3.58), the equilibrium utilities of the buyer and seller are, respectively,

$$U_B = u_B - \theta a_S \text{ and } U_S = u_S - \theta a_B. \quad (3.59)$$

Let us use this framework to find the optimal contracts for the following simple case. A buyer would like to enter into a date 0 contract with a seller to supply 1 unit of an object at date 1. It is common knowledge that the buyer’s utility is $u_B = 100 - p$, where p is the date 1 price at which the object is exchanged. The seller’s cost is zero, so $u_S = p$. The total surplus is $u_B + u_S = 100$. There are no ex-ante non-contractible investments. The results for this simple case are listed in the next proposition.

Proposition 3.34 (i) *In the neoclassical framework, the solution is identical in the absence or presence of an ex-ante contract.*

(ii) *In the “contracts as reference points” framework, the total surplus is given by $(1 - \theta) 100$, which is independent of p and there is a deadweight loss of 100θ .*

(iii) *Ex-post Coasian renegotiation does not alter the result in (ii).*

(iv) *A mechanism that gives residual control rights to the buyer to choose a date 1 price does not alter the result in (ii).*

(v) *Suppose that there is a competitive market for sellers or buyers. Consider a rigid contract at date 0 that fixes a date 1 price, p , determined by competitive contracts. Such a contract achieves a social surplus of 100.*

(vi) *Suppose that the date 0 market is imperfectly competitive and the contracting parties agree on some price interval $[\underline{p}, \bar{p}]$ for the date 1 trading price, such that $0 < \underline{p} < \bar{p} < 100$. Then the shading costs of aggrievement are $\theta(\bar{p} - \underline{p}) < 100\theta$.*

Proof: (i) Consider the neoclassical framework in which there is absence of aggrievement and shading. Assume that there is a competitive supply of sellers but only one buyer, and lump-sum transfers are allowed. Suppose that the buyer writes a contract with one of the sellers from the competitive pool of sellers at date 0. All parties know that at date 1, there is bilateral monopoly between the contracting parties, so that the date 1 price is determined by bargaining. The total

surplus from trading is $u_B + u_S = 100$, arising entirely from the buyer's valuation, because the seller's cost is zero.

Assume symmetric Nash bargaining among the two parties at date 1, thus, $p = 50$ and the gains from trade are shared equally. However, at date 0, since there is a competitive pool of sellers, the single buyer must get all the surplus. Anticipating ex-post opportunism, the seller who gets the contract at date 0, pays a lump-sum transfer to the buyer, equal to 50. Thus, the date 1 payoffs of the seller and the buyer are $u_B = 100$, $u_S = 0$. On the other hand if the parties had waited for date 1 to write a contract, then competition among the sellers would have driven the price to $p = 0$ so that $u_B = 100$, $u_S = 0$, which is identical to the outcome under a date 0 contract.

This result can also be shown to hold for other market configurations such as one buyer and one seller or many competitive buyers and one seller.

(ii) Consider one buyer and one seller. Suppose that the two parties do not write a date 0 contract and both realize that they will agree on some price p at date 1. At date 1, the maximum possible payoff of the buyer, as dictated by the date 0 contract, is $\bar{u}_B = 100$, so the buyer feels aggrieved if $p > 0$. Analogously, for the seller, $\bar{u}_S = 100$, so the seller feels aggrieved if $p < 100$. Thus, for any $0 < p < 100$ both parties will shade their consummate performance.

Since $a_j = \bar{u}_j - u_j$, $j = B, S$, we have that $a_B = 100 - (100 - p) = p$ and $a_S = 100 - p$. Thus, the buyer shades by $l_B = \theta a_B = \theta p$, and the seller shades by $l_S = \theta a_S = \theta(100 - p)$. Using (3.59), the final utilities of the buyer and the seller are, thus,

$$U_B = 100 - p - \theta(100 - p) = (1 - \theta)(100 - p).$$

$$U_S = p - \theta p = (1 - \theta)p.$$

The total surplus is $W = U_B + U_S = (1 - \theta)100$, which is independent of p . Thus, there is a deadweight loss of 100θ that is increasing in the weight that the two parties put on aggrievement.

(iii) Suppose that at date 1, the buyer pays an amount x to the seller to reduce the seller's shading. The seller will find that a_S falls (see (3.56)), so he indeed shades less. However, the buyer will find that a_B goes up, so he feels more aggrieved, thus, he will shade more. On net, given the linear, additive, structure of the model, there is no change to the result in (ii).

(iv) Part (ii) holds for any arbitrary date 1 price, p and there is a deadweight loss of 100θ , which is independent of p . In particular, this result holds for $p = 0$. In part (iv) we have a particular price p chosen by the buyer who makes a take it or leave it offer to the seller, hence, $p = 0$. So we get the same deadweight loss and same total surplus as in part (ii).

(v) Suppose that there is a competitive pool of suppliers and one buyer, so that the buyer has all the bargaining power. Then a rigid contract at date 0 must fix a date 1 price of $p = 0$. Conversely, if there are a large number of buyers and one seller, this price must be $p = 100$. Once the price is fixed at date 0, none of the parties feels aggrieved at date 1 from the implementation of the date 0 price. Since the date 0 price was competitively determined, each party realizes that this price fixes an objective, market determined, entitlement. Thus, no shading of consummate performance occurs and the total surplus is $100 - p + p = 100$.

(vi) In this case, given the definition of $\bar{u}$, the buyer feels entitled to $\bar{u}_B = 100 - \underline{p}$ and the seller feels entitled to $\bar{u}_S = \bar{p}$. Suppose that trade takes place at some $p \in [\underline{p}, \bar{p}]$. Then $a_B = (100 - \underline{p}) - (100 - p) = p - \underline{p}$. Thus, the buyer shades his effort by $\theta a_B = \theta(p - \underline{p})$. For the seller, $a_S = \bar{p} - p$, thus, the seller shades by $\theta a_S = \theta(\bar{p} - p)$. Total shading is given by $\theta(p - \underline{p}) + \theta(\bar{p} - p) = \theta(\bar{p} - \underline{p}) < 100\theta$ because $0 < \underline{p} < \bar{p} < 100$. Hence, the deadweight loss in this case is lower and the social surplus is higher. ■

Proposition 3.34 shows that rigid contracts that determine the date 1 terms of trade at date 0 are preferable to flexible contracts that leave price determination to date 1. This is because they lead to lower aggrievement and lower shading of effort. An extension of these results to the case of non-contractible ex-ante investment can be found in Hart (2009).⁴³

In general, Proposition 3.34 need not hold when either the buyer's valuation of the object, v , or the seller's cost, c , are uncertain. Suppose that trading at date 1 is voluntary and a date 0 contract specifies a no-trade price p_0 (paid by the buyer to the seller in the event of no-trade) and a trade price, p_1 , if trade occurs at date 1. We allow lump-sum transfers between the two parties and no renegotiation of contracts at date 1. Let $q = 1$ if trade occurs at date 1, and $q = 0$ if no trade occurs. The buyer and the seller voluntarily engage in trade at date 1 if, respectively, $v - p_1 \geq 0 - p_0$ and $p_1 - c \geq p_0$, which can be summarized by

$$q = 1 \Leftrightarrow v \geq p_1 - p_0 \geq c. \quad (3.60)$$

We normalize $p_0 = 1$ because lump-sum transfers are allowed. Thus, the relevant trade condition is $v \geq p_1 \geq c$. One may compare this with the first best condition for trade:

$$q = 1 \Leftrightarrow v \geq c. \quad (3.61)$$

Since condition (3.60) is more restrictive than (3.61), thus, in the second best, trade takes place less often as compared to the socially optimal level.

Suppose that we consider a slightly more general contract in which:

- (i) p_1 is not a point price but it can be chosen from the interval $[\underline{p}, \bar{p}]$, and
- (ii) the contract gives the buyer the right to choose the date 1 price, once any uncertainty has resolved.

In this case, a generalized version of (3.60) can be written as

$$q = 1 \Leftrightarrow \exists p_1 \in [\underline{p}, \bar{p}] : v \geq p_1 - p_0 \geq c. \quad (3.62)$$

Normalizing $p_0 = 0$, a sufficient condition for (3.62) to hold for all $p_1 \in [\underline{p}, \bar{p}]$ is

$$q = 1 \Leftrightarrow v \geq \bar{p} \text{ and } c \leq \underline{p}. \quad (3.63)$$

This condition holds irrespective of whether the buyer or the seller has residual control rights to choose $p_1 \in [\underline{p}, \bar{p}]$.

Assume that each party recognizes that due to the voluntary trade assumption, they cannot get more than 100% of the trade surplus. Thus, each party feels an entitlement up to the maximum surplus within the contract, conditional on trade taking place. So long as they achieve this level of entitlement, they feel no aggrievement.

Definition 3.63 *An 'optimal contract' maximizes ex-ante expected surplus, net of shading costs.*

⁴³ For an extension of the Hart-Moore framework to renegotiation, see Halonen-Akatwijuka and Hart (2013) and for an application to the scope of the firm, see Hart and Holmstrom (2010).

	s1	s2
v	9	20
c	0	10

Figure 3.9 Valuations and costs in two states of the world.

	s1	s2	s3
v	9	20	20
c	0	10	0

Figure 3.10 Valuations and costs in three states of the world.

Example 3.26 Suppose that there are two possible states of the world at date 1, s_1 and s_2 , but at date 0 it is not known which state will arise. The buyer's valuation, v , and the seller's cost, c , are given in Figure 3.9. Suppose first that p_1 is a point price. Then, we need to find a single value for p_1 at date 0 such that $v \geq p_1 \geq c$ in each state (see (3.60) when $p_0 = 0$). From Figure 3.9 there exists no such price.

Now consider a contract that specifies $p_1 \in [\underline{p}, \bar{p}] = [9, 10]$. Recall that the buyer has the right to choose the date 1 price, once the state of the world is known. In state s_1 , the buyer picks $p_1 = 9$. Since this is the lowest possible price in the range of feasible prices $[9, 10]$, the buyer does not feel aggrieved. The seller too does not feel aggrieved because any higher price will exceed the buyer's valuation and prevent trade from taking place. When the state of the world is s_2 , the buyer picks $p_1 = 10$, which is the upper bound of the feasible set of prices. The buyer does not feel aggrieved because any lower price will lead to no trade. The seller does not feel aggrieved because he gets his maximum possible entitlement within the contract. Thus, in each state, we obtain an efficient outcome.

The outcome under uncertainty about v, c need not always be efficient (although it is efficient in Example 3.26). Consider, for instance, Figure 3.10, which augments an extra state to Example 3.26.

We leave the proof of the following proposition as an exercise.

Proposition 3.35 Consider Figure 3.10 and suppose that the three states, s_1, s_2 , and s_3 occur with respective probabilities π_1, π_2 , and π_3 . In all states, $v > c$ so it is socially optimal to trade. Consider three possible contracts A, B, and C. In contract A, $\underline{p} = \bar{p} = 9$. In contract B, $\underline{p} = \bar{p} = 10$. In contract C, $\underline{p} = 9, \bar{p} = 10$.

- (i) In contracts A and B, there is inefficient amount of trade.
- (ii) Contract A is optimal if π_2 is small, contract B is optimal if π_1 is small, and contract C is optimal if π_3 is small or θ is small.

In Proposition 3.35, a contract is more likely to be optimal if the probability of the state in which it does not allow trade is also low. Of the available contracts, contract C allows trade in all states, yet in state 3, there are aggrivement costs, hence, contract C is likely to be optimal if state 3 is less likely and if aggrivement costs are small.

3.10.2 Empirical evidence for contracts as reference points

Fehr et al. (2011) test if ex-ante competition serves to create entitlements via the reference point effect, using a version of the Hart and Moore (2008) model with payoff uncertainty. Their aim is not to test if players choose their contracts optimally, but rather to test if the behavioral assumptions of the Hart–Moore framework are empirically relevant.

Groups of 2 sellers and 2 buyers interact over 15 periods; groups are randomly rematched in each round to minimize reputational concerns. Each seller can sell up to 2 units, while each buyer needs only 1 unit at date 1, so there is seller competition. Parties are matched at date 0. At date 1, the contracted seller not only supplies the contracted quantity of 1 unit, but also determines the quality, q , which can take two possible values, normal (q_n) or low (q_l). There are two states of the world, the good state ($s = s_g$) that occurs with probability 0.8 and the bad state ($s = s_b$) that occurs with probability 0.2.

The production cost of the seller at date 1 is given by $c(q, s)$, where s is the state of the world and $c(q_n, s) < c(q_l, s)$, i.e., it is more costly to produce low quality in any state. The intuition is that it might be costly to alter the production process to produce an altered, and shaded level of consummate effort that results in a lower quality. The seller's costs are higher in the bad state, relative to the good state, so $c(q, s_g) < c(q, s_b)$. The payoffs of the buyer and the seller are given, respectively, by

$$u_B = v(q) - p \text{ and } u_S = p - c(q, s),$$

where v gives the buyer's valuation, and $v(q_n) > v(q_l)$, i.e., the buyer prefers normal to low quality. The outside option of the seller for any unit that he is unable to sell is $x_s = 10$. The payoff relevant data used in the experiments is shown in the self-explanatory Figure 3.11.

The buyers choose among contracts. They can choose either a rigid contract (r), or a flexible contract (f). These contracts are then auctioned off to sellers. In a *rigid contract*, the contract price $p_r \in [c(q_l, s_g) + x_s, 75] = [35, 75]$ is determined in an auction where the sellers bid against each other. The lower bound of 35 ensures that in the good state, even with the low quality, the

	s_g		s_b	
	q_n	q_l	q_n	q_l
c	20	25	80	85
v	140	100	140	100

Figure 3.11 Quality dependent valuation and cost data in two states of the world.

Source: Fehr et al. (2011) with permission from the American Economic Association.

seller does not make losses by trading, relative to getting an outside option price of $x_s = 10$ for unsold units. The upper limit ensures that the seller is always better off exercising his outside option price of $x_s = 10$ if the state is bad, rather than engaging in trade. This captures the idea that contract rigidity does not allow the parties to adapt well, ex-post, to changes in circumstances.

A *flexible contract* vests residual control rights in the hands of the buyer to choose the price p_f at date 1, once the state of the world is revealed. In the good state, s_g , the price $p_f \in [\underline{p}, \bar{p}]$, where $\underline{p} \in [35, 75]$ as in the rigid contract, and $\bar{p} = v(q_n) = 140$. In the bad state, s_b , the buyer chooses a price $p_f \in [c(q_l, s_b) + x_s, v(q_n)] = [95, 140]$. This ensures that in the bad state, the buyer's price does not let the seller make a loss by trading relative to exercising the outside option of $x_s = 10$. Unlike a rigid contract, a flexible contract permits trade to take place in both states of the world.

In each of the two contracts, sellers observe the price and then choose their quality. If they decide to choose a lower quality, they incur an extra cost of 5 units (see Figure 3.11).

Let us first summarize the results for the case of purely self-regarding preferences. In this case, bygones are bygones, so each seller will always choose the normal quality to avoid reducing their own payoff by 5 units. Anticipating that the seller never engages in shading of consummate performance, buyers choose the smallest price allowed in their respective contracts. In the rigid contract, buyers choose a price of 35 in each state of the world. Thus, the buyer's payoff equals 105 in the good state but 0 in the bad state because at a price of 35, the seller prefers to exercise the outside option. By contrast, in the flexible contract, in the good state, the buyer chooses 35 (so his payoff is 105) but in the bad state the buyer is restricted to choose no lower than 95, so his payoff is 45. Clearly, the flexible contract dominates the rigid contract because the payoffs are identical in the good state but relatively higher in the bad state under a flexible contract. These results do not depend on the presence of ex-ante competition.

Now consider the predictions in the “contracts as reference points” framework. In the rigid contract, the prices are fixed up-front, so the seller does not feel aggrieved ex-post, thus, no shading takes place. Therefore, we do not expect low quality to be chosen. In a flexible contract, by contrast, both parties have different entitlements, so shading by the seller does take place. Such shading is personally not very costly for the seller (a loss of 5 units in terms of the cost), however, it is costly for the buyer (a loss of 40 units in terms of the valuation in each state). The higher the price, the lower is the expected shading on the part of the seller. In response, buyers will attempt to increase the price to reduce shading by the seller. But increasing the price reduces the buyer's profits under flexible contracts, so they might end up preferring rigid contracts.

Recall that in a flexible contract, in the bad state, $p_f \in [95, 140]$. In an additional treatment, buyers are restricted to choose the date 1 price in the interval $[95, 95]$. Since there is no ex-post flexibility for the buyer in this case, we expect lower shading by the seller. Yet another treatment eliminates ex-ante seller-competition. In the absence of seller-competition, the sellers may increase their entitlements, so they may engage in more shading.

The empirical results confirm the predictions of the Hart–Moore model that are outlined above, but reject the model of purely self-regarding preferences. The results are summarized in Table 3.5.

The *competitive price* is simply the cost of the seller plus the outside option. Since trade never takes place in the bad state under rigid contracts, the competitive price in a rigid contract is 35 in the good state. Analogously, in a flexible contract that allows trade in the bad state, the competitive price is 95. From Table 3.5, the average prices in rigid and fixed contracts are 40.7 and 40.2, respectively, which are not statistically different.

Table 3.5 Summary of the results for rigid and flexible contracts.

Contract type	Rigid contract		Flexible contract	
	Good	Bad	Good	Bad
Average price	40.7	—	51.1	98.4
Relative frequency of normal quality	0.94	—	0.75	0.70
Average auction outcome	40.7		40.2	
Average profit buyer (per state)	96.8	10	78.9	29.7
Average profit seller (per state)	20.4	10	29.8	16.9
Average profit buyer (over both states)	77.9		68.9	
Average profit seller (over both states)	18.1		27.2	
Relative frequency of contract	0.50		0.50	

Source: Fehr et al. (2011).

Looking at the results in each state, although the lower bound for price is only 35 in the good state of a flexible contract, buyers use their residual control right over prices to offer a price of 51.1. This is significantly higher than the price of 40.7 in the good state of a rigid contract where no shading can occur. In a flexible contract, 75% of the prices are above the lower bound. The difference between actual prices and the lower bound is significant, stable, and persists over time. This is consistent with the prediction that in a flexible contract, buyers offer a higher price to take account of the possibility of shading by the sellers.

Sellers rarely shade in the rigid contract in the good state (no trade takes place, by design, in the bad state) and in 94% of the cases, sellers choose normal quality. However, the corresponding figure for flexible contracts is 75%. A non-parametric signed rank test confirms that the difference between the two figures is significant. This is consistent with the hypothesis that sellers are aggrieved when they receive a payoff below their entitlements.

As noted above, if all parties had self-regarding preferences, then buyers should always prefer the flexible contract. However, in the Hart–Moore framework, buyers may get a higher payoff from the rigid contracts on account of the shading behavior of sellers in a flexible contract. The data confirms this prediction. Average profit over both states in the rigid contract is 77.9, while it is 68.9 in the flexible contract. By contrast, the average profits of sellers are higher in the flexible contract relative to the rigid contract. It would seem that buyers may have conceded a higher price that more than offsets the higher costs of the sellers from shading in the flexible contracts.

Possibly in response to these differences in profits, on average, buyers choose the rigid contract in 50% of the cases. Over time, the share of rigid contracts increases from 38% in round 1 to 56% by the last round (round 15) and the time trend is statistically significant. These results provide insights about why some forms of rigid contracts such as the employment contract may be preferred in many cases.

3.11 Moral hazard, loss aversion, and optimal contracts

Consider a simple principal–agent model in which the principal is risk neutral and the agent has prospect theory preferences (de Meza and Webb, 2003, 2007). The outside option of the agent

is zero. The agent chooses an action $a \in \{0, 1\}$. Each action probabilistically determines one of three observable output levels, $0 \leq y_L < y_M < y_H$. Let P and Q be the distribution functions, respectively, under $a = 1$ and $a = 0$. Under action $a = 1$, the probability of outcome y_i is p_i , $i = L, M, H$, while under action $a = 0$, the corresponding probability is q_i . We make two assumptions about the probability distribution of outcomes.

- A1. The distribution of outcomes under $a = 1$ first order stochastically dominates the distribution under $a = 0$, so $P_i \leq Q_i$ for $i = L, M, H$.
- A2. The *monotone likelihood ratio property* (MLRP) holds, so

$$\frac{p_L}{q_L} < \frac{p_M}{q_M} < \frac{p_H}{q_H}.$$

Thus, the probability of a higher outcome under the higher action is higher. In other words, a higher output is more likely to have come from a higher effort level.

The only observable and verifiable signal of the agent's effort is the output level produced. Hence, the principal conditions the agent's reward scheme on the observed output level. The agent receives a wage w_i when the observed output level is y_i . The agent has limited liability ($w_j \geq 0$) and there is no renegotiation of contracts. For pedagogical ease, let the *reference wage* of the agent be w_M . Since the MLRP holds, the optimal wage is increasing in the observed output level, so w_M will also be the median wage. The utility function of the agent from the wage w_i and action a , is given by

$$v(w_i | w_M, a) = \begin{cases} u(w_i) - a & \text{if } w_i > w_M \\ u(w_i) - \lambda (w_M - w_i) - a & \text{if } w_i \leq w_M \end{cases},$$

where $u' > 0$, $u'' < 0$, and $\lambda > 0$ is the coefficient of loss aversion. Thus, when $w_i \leq w_M$ the decision maker suffers disutility from loss aversion that is linear in the distance $w_M - w_i$. This formulation is similar to the Köszegi–Rabin framework in which the individual derives consumption utility and gain–loss utility, except that here we only capture disutility from losses. In order to ensure an interior solution, we impose the Inada condition, $\lim_{x \rightarrow 0} u'(x) = -\infty$.

A contract is given by four parameters, (w_L, w_M, w_H, a) , where w_i is the contractible wage level if the observed output is y_i , and a is the *non-contractible, desired action*. We assume that the principal always wishes to induce the higher action, $a = 1$. Thus, we wish to focus attention on contracts of the form $(w_L, w_M, w_H, 1)$ in which the principal optimally chooses w_L, w_M, w_H . Let $u_i = u(w_i)$ and define the expected utility of the agent and expected cost of the principal under the distribution P by

$$\bar{u}(P) = p_H u_H + p_M u_M + p_L u_L, \bar{C}(P) = p_H w_H + p_M w_M + p_L w_L. \quad (3.64)$$

One can analogously define $\bar{u}(Q)$ and $\bar{C}(Q)$ under the distribution Q . Define

$$\Delta \bar{u}(P) = \bar{u}(P) - \bar{u}(Q), \Delta p_i = p_i - q_i, \Delta w_M = w_M - w_L; i = L, M, H.$$

With this definition, we can write the MLRP in Assumption A2 as

$$\frac{\Delta p_L}{p_L} < \frac{\Delta p_M}{p_M} < \frac{\Delta p_H}{p_H}. \quad (3.65)$$

The expected utility of the agent when he chooses the action $a = 1$ and induces the distribution P over outcomes can be written as

$$E_P v(w_i | w_M, 1) = \bar{u}(P) - p_L \lambda \Delta w_M - 1.$$

Similarly, the agent's expected utility when he takes action $a = 0$ and induces the distribution Q over outcomes is

$$E_Q v(w_i | w_M, 0) = \bar{u}(Q) - q_L \lambda \Delta w_M.$$

In order to implement the action $a = 1$, the principal chooses w_L, w_M, w_H to solve Problem 1:

$$\text{Min } \bar{C}(P)$$

subject to:

$$\begin{aligned} 0 &\leq w_L \leq w_M. & (\text{NDW}) \\ E_P v(w_i | w_M, 1) &\geq 0 \Leftrightarrow \bar{u}(P) - p_L \lambda \Delta w_M - 1 \geq 0. & (\text{IR}) \\ E_P v(w_i | w_M, 1) &\geq E_Q v(w_i | w_M, 0) \Leftrightarrow \Delta \bar{u}(P) - \Delta p_L \lambda \Delta w_M - 1 \geq 0. & (\text{IC}) \end{aligned}$$

The first constraint restricts attention to non-decreasing wage schemes (NDW); this is motivated by the MLRP assumption. The *individual rationality constraint* (IR) requires that the agent earns at least his outside option of zero when he takes the action $a = 1$. Since the reference point is w_M , the agent only experiences a loss when he receives a wage w_L . The *incentive compatibility constraint* (IC) requires that the agent gets at least as much expected utility from taking the action $a = 1$ as he does from the action $a = 0$. Setting $\lambda = 0$ recovers the no loss aversion case.

Letting $\phi \geq 0, \gamma \geq 0, \mu \geq 0$ be the Lagrangian multipliers on the constraints, NDW, IR, and IC, respectively, the Lagrangian for the minimization problem, Problem 1, can be written as:

$$\Gamma = \bar{C}(P) + \gamma [-\bar{u}(P) + p_L \lambda \Delta w_M + 1] + \mu [-\Delta \bar{u}(P) + \Delta p_L \lambda \Delta w_M + 1] - \phi \Delta w_M. \quad (3.66)$$

Our assumptions on preferences and constraints ensure that the Kuhn–Tucker conditions are necessary and sufficient for a minimum. The Kuhn–Tucker conditions are as follows.

$$\frac{\partial \Gamma}{\partial w_H} = p_H - \gamma p_H u'_H - \mu \Delta p_H u'_H \geq 0. \quad (3.67)$$

$$\frac{\partial \Gamma}{\partial w_M} = p_M - \gamma (p_M u'_M - p_L \lambda) - \mu (\Delta p_M u'_M - \Delta p_L \lambda) - \phi \geq 0. \quad (3.68)$$

$$\frac{\partial \Gamma}{\partial w_L} = p_L - \gamma (p_L u'_L + p_L \lambda) - \mu (\Delta p_L u'_L + \Delta p_L \lambda) + \phi \geq 0. \quad (3.69)$$

$$\frac{\partial \Gamma}{\partial \gamma} \leq 0; \quad \frac{\partial \Gamma}{\partial \mu} \leq 0; \quad \frac{\partial \Gamma}{\partial \phi} \leq 0. \quad (3.70)$$

$$w_i \frac{\partial \Gamma}{\partial w_i} = 0, i = L, M, H. \quad (3.71)$$

$$\gamma \frac{\partial \Gamma}{\partial \gamma} = 0; \mu \frac{\partial \Gamma}{\partial \mu} = 0; \phi \frac{\partial \Gamma}{\partial \phi} = 0. \quad (3.72)$$

The first order conditions with respect to the state-contingent wages are given in (3.67)–(3.69). The condition (3.70) recovers the three constraints. The two complementary slackness conditions are given in (3.71) and (3.72). Given the Inada condition, $\lim_{x \rightarrow 0} u'(x) = -\infty$, it must be the case that $w_i > 0$, $i = L, M, H$, otherwise, $\frac{\partial \Gamma}{\partial w_i} < 0$, which is not possible. But if $w_i > 0$, then we get from (3.71) that $\frac{\partial \Gamma}{\partial w_i} = 0$, $i = L, M, H$.

The decision maker feels loss averse whenever $\Delta w_M > 0$. On the one hand, as in classical agency theory, $\Delta w_M > 0$ helps the principal to counter the moral hazard problem by providing higher incentives for better outcomes. However, because the agent is loss averse, this also requires a higher wage compensation to the agent in order to fulfill his IR condition. The size of this compensation is directly proportional to the loss aversion parameter, λ . Indeed, if λ is high enough, the second effect is likely to dominate the first effect, and the firm may find it optimal to set $\Delta w_M = 0$. In this case, the agent does not face any downside risk. De Meza and Webb (2007) conjecture that this case applies to CEO salary compensation that takes the form of stock options, which protect the CEO from downside risk. The next proposition formalizes this idea.

Proposition 3.36 *There exists a critical value of the parameter of loss aversion, λ_c , such that $\lambda \geq \lambda_c$ is necessary and sufficient for $\Delta w_M = w_M - w_L = 0$.*

Proof: Here we prove the if part. Suppose that $\Delta w_M = w_M - w_L = 0$, so $u'_L = u'_M$. Recall that $\frac{\partial \Gamma}{\partial w_i} = 0$, $i = L, M, H$. Thus, from (3.67) we must have

$$\frac{1}{u'_H} = \gamma + \mu \frac{\Delta p_H}{p_H}. \quad (3.73)$$

Adding (3.68), (3.69), noting that $\Delta w_M = 0$ and $u'_L = u'_M$, and dividing by $p_L + p_M$ we have

$$\frac{1}{u'_M} = \gamma - \mu \frac{\Delta p_H}{1 - p_H}. \quad (3.74)$$

First we show that in this case, $\gamma > 0$ and $\mu > 0$. Suppose not. If $\gamma = 0$, then from (3.74) we get $u'_M < 0$, which is a contradiction. If $\mu = 0$, then, from (3.73), (3.74), we have $u'_M = u'_H$. Thus, in this case $w_L = w_M = w_H$. However, if the agent is paid a fixed wage, then this destroys incentives and he will choose the minimum action, $a = 0$.⁴⁴ This contradicts the desired implementation of $a = 1$. Thus, it follows that $\gamma > 0$ and $\mu > 0$.

From (3.69), $\frac{\partial \Gamma}{\partial w_L} = 0$ implies that

$$1 + \frac{\phi}{p_L} = (u'_L + \lambda) \left(\gamma + \mu \frac{\Delta p_L}{p_L} \right). \quad (3.75)$$

⁴⁴ The IC then gives $-1 \geq 0$, which is a contradiction.

When $\phi = 0$, (3.75) implies $\lambda = \lambda_c = \left(\gamma + \mu \frac{\Delta p_L}{p_L}\right)^{-1} - u'_L$. Thus, when $\lambda = \lambda_c$ it is optimal for the principal to choose $\Delta w_M = 0$ in the absence of the NDW constraint. But for any $\phi > 0$, (3.75) implies that $\lambda > \lambda_c$ to ensure that our maintained assumption, $\Delta w_M = 0$, holds. Thus, for all $\lambda \geq \lambda_c$ we have $\phi > 0$, so $\phi \frac{\partial \Gamma}{\partial \phi} = 0$ implies $\frac{\partial \Gamma}{\partial \phi} = -\Delta w_M = 0$, or $\Delta w_M = 0$. We leave the if and only if part to the reader. ■

Corollary 3.3 $\lambda < \lambda_c$ is necessary and sufficient for $\Delta w_M = w_M - w_L > 0$.

Proof: Proposition 3.36 shows that $\lambda \geq \lambda_c \Leftrightarrow \Delta w_M = 0$. Thus, if $\lambda < \lambda_c$, then $\Delta w_M \neq 0$. But $w_M < w_L$ would violate the MLRP when used in conjunction with (3.68), (3.69). So the only possibility is $\Delta w_M > 0$. ■

Classical moral hazard theory predicts *high-powered incentives*, i.e., wage schemes that are very sensitive to the observed and verifiable outcomes. In real life, one typically observes incentives that are *low-powered*, relative to these predictions.⁴⁵ High enough loss aversion, by leading to flat segments in the optimal incentive scheme ($\Delta w_M = 0$; see Proposition 3.36) provides one potential explanation.

In Proposition 3.36, the incentive scheme is flat at the bottom of the outcome range. However, when there are more than three states of the world, then it is possible that the incentive scheme may be flat in the intermediate ranges. Consider, for instance, four possible states of the world, in which the output levels are $0 \leq y_S < y_L < y_M < y_H$. In all other respects, the model is identical to the one with three states. In this case, suppose that the principal wishes to implement the high action, $a = 1$. Then the optimal contract specifies the wage levels (w_S, w_L, w_M, w_H) . The main result is summarized in the next proposition. The proof is similar to that of Proposition 3.36, so we leave it as an exercise.

Proposition 3.37 (de Meza and Webb, 2003): *There exist critical values of the parameter of loss aversion, $\lambda_1 \geq \lambda_2 \geq 0$, such that:*

- (i) $\lambda \geq \lambda_1$ is necessary and sufficient for $w_S = w_L = w_M$.
- (ii) $\lambda_1 \geq \lambda \geq \lambda_2$ is necessary and sufficient for $w_L = w_M$ but $w_S < w_L$.
- (iii) $\lambda_1 \geq \lambda_2 \geq \lambda$ is necessary and sufficient for $w_S < w_L < w_M < w_H$.

The main insight of this analysis is that, relative to the predictions of the classical agency model, loss aversion may lead to low-powered incentive schemes. This insight helps to explain several interesting economic phenomena. Herweg and Mierendorff (2013) show that if loss aversion of consumers is high enough, then a monopolist risk-neutral seller finds it optimal to offer a two-part tariff to the consumer. This involves a fixed charge to be paid by the consumer, and a variable charge that depends on the number of units consumed. The authors argue that their framework provides the rationale for a fixed payment in situations where the competing explanations are not very persuasive.

Heidhues and Köszegi (2005, 2008) consider a model of monopolist competition in which consumers are loss averse and have a reference price. If the price is higher than the reference price, they suffer a gain–loss disutility. Thus, at prices that exceed the reference point, firms compete more aggressively in prices. The reason is that lowering prices towards the reference price contributes greatly to an increase in consumer utility by avoiding expensive loss aversion. The reference point might then serve as a *focal point* around which firms compete. Loss aversion, thus, leads to non-responsiveness of prices to changes in economic circumstances, particularly when

⁴⁵ See Chapter 8 for the empirical evidence.

the desired price exceeds the focal reference price. Interestingly, a single focal price may exist even when firms are asymmetric. Spiegel (2012) considers a simplified version of this model. See Spiegel (2011, Chapter 9) for a further discussion of optimal contracts under loss aversion.

Herweg et al. (2010) consider another implication of loss aversion. They seek to explain why most real-world incentive schemes are of a binary nature, i.e., they have a fixed wage and a bonus payment that is conditional on the observed outcome exceeding a certain threshold. They consider a principal–agent framework with moral hazard on the part of the agent. The agent has Köszegi–Rabin preferences; see Section 2.8 above. As in de Meza and Webb (2007), Herweg et al. (2010) show that loss aversion leads to lower-powered incentive schemes. However, the main insights of their paper arise from the manner in which it is optimal for the principal to bundle different signals of the agent’s effort level, when providing a bonus. While it is not possible to give an exposition of their model without a fuller specification of the details, we consider below a simple example of the sorts of issues that can arise; the interested reader should pursue the much richer model in Herweg et al. (2010).

Suppose that the observable and verifiable output levels are $0 \leq y_1 \leq y_2 \leq \dots \leq y_n$. The principal is interested in eliciting a particular effort level from the agent that gives rise to a probability distribution P over outcomes. So the outcomes $y_1, y_2, \dots, y_n$ occur with the respective probabilities $p_1, p_2, \dots, p_n$. In order to explain the intuition behind the results, let us restrict attention to the case of increasing wage schemes, so $0 \leq w_1 \leq w_2 \leq \dots \leq w_n$. Assume also that the material utility is given by $m(x) = x$ and gain–loss utility, μ , is linear and is operative only for losses as in (2.77). Suppose that the realized state is i , with the associated wage w_i . Then, the Köszegi–Rabin utility of the agent (see Section 2.8) is

$$u(w_i | r) = w_i - \sum_{j=i+1}^n p_j \lambda (w_j - w_i).$$

Thus, given $m(x) = x$ and (2.77), the expected Köszegi–Rabin utility of the agent, prior to the realization of the state of the world is

$$Eu = \sum_{j=1}^n p_j u(w_j | r) = \sum_{j=1}^n p_j w_j - \lambda \sum_{i=1}^{n-1} p_i \sum_{j=i+1}^n p_j (w_j - w_i). \quad (3.76)$$

We now consider a sequence of hypothetical incentive schemes to illustrate the intuition behind the results. Suppose that we begin with an initial fixed wage contract, $w_j = w, j = 1, 2, \dots, n$; we call this the *baseline contract*. In the baseline contract, there is no gain–loss utility, so the Köszegi–Rabin utility, in (3.76), denoted by Eu_0 , is

$$Eu_0 = w.$$

The baseline contract illustrates the classic conflict between insurance and incentives. By putting zero risk on the agent, the contract eliminates incentives completely, so a self-regarding agent will put in the least amount of effort. In other words, this contract is not incentive compatible. Starting from the baseline contract, consider a $n - 1$ *perturbed contract* that offers the following incentive structure

$$w_1 = w \text{ and } w_j = w + b, j = 2, 3, \dots, n,$$

where $b > 0$ is a credible bonus payment. Recall that the distribution function of outcomes for the desired effort level is P . The expected Köszegi–Rabin utility, in (3.76), for the $n - 1$ *perturbed contract*, Eu_1 , can be written as

$$Eu_1 = w + b(1 - P_1)(1 - \lambda P_1),$$

where $P_1 = \text{Prob}(y \leq y_1)$. Now consider a $n - 2$ perturbed contract

$$w_1 = w_2 = w \text{ and } w_j = w + b, j = 3, 4, \dots, n.$$

The expected Köszegi–Rabin utility, (3.76), for this incentive scheme, Eu_2 , can be written as

$$Eu_2 = w + b(1 - P_2)(1 - \lambda P_2),$$

where $P_1 = \text{Prob}(y \leq y_2)$. These calculations reveal the pattern of recursion. In general, for any $n - k$ perturbed contract

$$w_1 = w_2 = \dots = w_k = w \text{ and } w_{k+1} = w_{k+2} = \dots = w_n = w + b,$$

we have that the expected Köszegi–Rabin utility is

$$Eu_k = w + b(1 - P_k)(1 - \lambda P_k), \quad (3.77)$$

where $P_k = \text{Prob}(y \leq y_k)$.

In the absence of loss aversion ($\lambda = 0$), bonus payments always increase the expected utility of the agent, hence, they serve to loosen the individual rationality constraint of the agent. The exact bonus payment may be calculated in such a manner that the individual rationality constraint always holds. Now consider the presence of loss aversion ($\lambda > 1$). If loss aversion is low enough, in the sense that $\lambda < \frac{1}{P_k}$ for all $1 < k \leq n$, then using (3.77), bonuses for performance that exceeds some threshold always improve expected utility.

The more interesting situation arises when $\frac{1}{P_{k+1}} < \lambda < \frac{1}{P_k}$. In this case, offering bonuses for equalling or exceeding the output threshold of y_{k+1} (a $n - (k + 1)$ perturbed contract) reduces the expected utility of the agent. This result is contrary to that obtained in the standard model. Thus, once this threshold is reached, the principal will have to increase the fixed wage, w , of the agent to implement the desired effort. On the other hand if the performance threshold is chosen to be below y_{k+1} , say at y_k then bonuses for performance that exceeds y_k increase the expected utility of the agent. Such an incentive scheme loosens the individual rationality constraint and makes implementation of the desired effort easier. On the other hand if k is very high, then $1 - P_k$ is very low, which reduces the marginal impact on the agent's utility from an increase in the bonus payment. For a characterization of the optimal contract under these sorts of considerations, see Herweg et al. (2010).

3.12 Renegotiation, long-term contracts, and loss aversion

In the traditional theory of the firm, once the state of the world is realized, parties that are bound by some initial contract use Coasian bargaining to renegotiate to an efficient outcome. Suppose that the terms of the initial contract serve to form a reference point for loss averse contracting parties, as in Hart and Moore (2008). However, parties write a contract with each other before the realization of state of nature. Once the state of nature is realized, the contracting parties can renegotiate the ex-ante reference contract to an efficient ex-post contract.




Figure 3.12 Sequence of moves in a bilateral trade model with renegotiation.

Source: Herweg, F. and Schmidt, K. M., "Loss aversion and inefficient renegotiation," *Review of Economic Studies* (2015) 82(1): 297–332, by permission of Oxford University Press.

An ex-post change in the terms of the contract through renegotiation may improve the payoff of one of the parties, but reduce the payoff of the other party. If the harmed party is loss averse, the extent of the harm is magnified. The party that benefits may be able to compensate the harm caused to the other party, but might not be able to compensate the scaled-up harm that arises on account of loss aversion. In this case, ex-post renegotiation may successfully alter the original contract only partly towards the efficient outcome.

Furthermore, this might lead to sticky ex-post renegotiation; such renegotiation may take place only for a restricted range of states of nature, while for all other states, it completely breaks down. Anticipating these ex-post problems, the parties may not wish to write long-term contracts, preferring instead to use the spot market, particularly if the spot market is fairly competitive. We consider below these issues in a model due to Herweg and Schmidt (2013).

Consider a buyer (B) and a seller (S) who agree at date 0 on two contractual features of one unit of a good/widget that they might trade in the future. These features are: a specification $x \in X$ of the good and a price $p \in P$, where X is a compact, possibly multidimensional, space and P is some non-negative closed real interval.

The sequence of moves is shown in Figure 3.12. At date 0, the two parties sign a contract $(\bar{x}, \bar{p})$. At date 1, a state of the world, $\theta \in \Theta$ is realized. This state determines which specification $x \in X$ of the traded good is most appropriate. At date 0, the two parties know that θ is distributed with the distribution function $F(\theta)$. They also know that ex-post, the state of the world is not verifiable to a third party, so they cannot write state-contingent contracts at date 0. Having observed the state at date 1, the parties can renegotiate to another feasible contract, $(\hat{x}, \hat{p})$, at date 2. If they fail to renegotiate, the default contract, $(\bar{x}, \bar{p})$ is implemented. The valuation of the buyer and the cost of the seller are, respectively, $v(x, \theta)$ and $c(x, \theta)$. Their respective payoffs are $v(x, \theta) - p$ and $p - c(x, \theta)$, which are realized at date 3.

Define the joint surplus of the two parties as

$$W(x, \theta) = v(x, \theta) - c(x, \theta), x \in X, \theta \in \Theta.$$

We assume that $W(x, \theta)$ is strictly concave, so the first order conditions are sufficient. In any given state $\theta \in \Theta$, denote the socially optimum specification found by maximizing W with respect to $x \in X$ by $x^*(\theta)$. This is given by the solution to the following equation.

$$\frac{\partial W}{\partial x} = \frac{\partial v(x^*(\theta), \theta)}{\partial x} - \frac{\partial c(x^*(\theta), \theta)}{\partial x} = 0. \quad (3.78)$$

Both parties have Köszegi–Rabin utilities. Define gain–loss utility by

$$\mu(z) = \begin{cases} 0 & \text{if } z > 0 \\ \lambda z & \text{if } z \leq 0 \end{cases}, \quad (3.79)$$

then the Kőszegi–Rabin utilities of the buyer and the seller from the contract $(\hat{x}, \hat{p})$, conditional on the state of the world $\theta \in \Theta$ are given, respectively, by

$$U_B(\hat{x}, \hat{p}; \theta) = v(\hat{x}, \theta) - \hat{p} + \mu(\bar{p} - \hat{p}) + \mu(v(\hat{x}, \theta) - v(\bar{x}, \theta)). \quad (3.80)$$

$$U_S(\hat{x}, \hat{p}; \theta) = \hat{p} - c(\hat{x}, \theta) + \mu(\hat{p} - \bar{p}) + \mu(c(\bar{x}, \theta) - c(\hat{x}, \theta)). \quad (3.81)$$

The original contract $(\bar{x}, \bar{p})$ establishes the relevant entitlements that determine the subsequent gains and losses of the two parties. We assume a common or universal gain–loss function, so that the form of μ is identical for each party and in each dimension (price, valuation, or cost). Moreover, we assume that gain–loss utility from price variation and value/cost variation is additively separable. The buyer feels a loss if the renegotiated price is higher than the date 0 price and the seller suffers a loss if it is lower. The buyer suffers a loss if the value of the renegotiated specification is lower than the date 0 specification and the seller suffers a loss if the cost of this specification is higher.

Clearly, if the original contract is not renegotiated, then there is no gain–loss utility, so

$$U_B(\bar{x}, \bar{p}; \theta) = v(\bar{x}, \theta) - \bar{p} \text{ and } U_S(\bar{x}, \bar{p}; \theta) = \bar{p} - c(\bar{x}, \theta). \quad (3.82)$$

Using (3.80)–(3.82) the buyer and the seller would wish to renegotiate if

$$B : v(\hat{x}, \theta) - v(\bar{x}, \theta) + \mu(v(\hat{x}, \theta) - v(\bar{x}, \theta)) \geq \hat{p} - \bar{p} - \mu(\bar{p} - \hat{p}). \quad (3.83)$$

$$S : c(\bar{x}, \theta) - c(\hat{x}, \theta) + \mu(c(\bar{x}, \theta) - c(\hat{x}, \theta)) \geq \bar{p} - \hat{p} - \mu(\hat{p} - \bar{p}). \quad (3.84)$$

Conditional on any initial contract, $(\bar{x}, \bar{p})$, the set of contracts $(\hat{x}, \hat{p})$ that satisfies (3.83), (3.84) is said to be *individually rational*. For a given initial contract, $(\bar{x}, \bar{p})$, let the set of all individually rational contracts be denoted by $I(\bar{x}, \bar{p})$. The *renegotiation set*, R , is simply the set of specifications of the good that the parties wish to voluntarily renegotiate to at some mutually agreed price, i.e.,

$$R(\bar{x}, \bar{p}) = \{\hat{x} \in X : \exists \hat{p} \in P \text{ and } (\hat{x}, \hat{p}) \in I(\bar{x}, \bar{p})\}.$$

The next proposition summarizes the renegotiation possibilities.

Proposition 3.38 Consider an initial contract $(\bar{x}, \bar{p})$ at date 0, and the realization of a state $\theta \in \Theta$ at date 1. Then the renegotiation set $R(\bar{x}, \bar{p})$ at date 2 is characterized as follows:

- (i) For some $\hat{x} \in X$, if $v(\bar{x}, \theta) \leq v(\hat{x}, \theta)$ and $c(\bar{x}, \theta) \geq c(\hat{x}, \theta)$, then $\hat{x} \in R(\bar{x}, \bar{p})$.
- (ii) For some $\hat{x} \in X$, if $v(\bar{x}, \theta) \leq v(\hat{x}, \theta)$ and $c(\bar{x}, \theta) \leq c(\hat{x}, \theta)$, then $\hat{x} \in R(\bar{x}, \bar{p})$ if and only if

$$v(\hat{x}, \theta) - v(\bar{x}, \theta) \geq (1 + \lambda)^2 [c(\hat{x}, \theta) - c(\bar{x}, \theta)]. \quad (3.85)$$

- (iii) For some $\hat{x} \in X$, if $v(\bar{x}, \theta) \geq v(\hat{x}, \theta)$ and $c(\bar{x}, \theta) \geq c(\hat{x}, \theta)$, then $\hat{x} \in R(\bar{x}, \bar{p})$ if and only if

$$c(\bar{x}, \theta) - c(\hat{x}, \theta) \geq (1 + \lambda)^2 [v(\bar{x}, \theta) - v(\hat{x}, \theta)]. \quad (3.86)$$

Proof: (i) If there is a specification such that at date 2, the valuation of the buyer is higher and the cost of the seller is lower, then each party benefits by agreeing to renegotiate to this specification even if they kept the initial price $\bar{p}$, unchanged.

(ii) Since renegotiation is voluntary, the buyer will have to offer $\hat{p} \geq \bar{p}$ in order to cover the higher costs of the seller. Thus, the buyer suffers a loss in the price domain and a gain in the valuation domain. The seller suffers a loss in the cost domain and a gain in the price domain. From (3.83), (3.84) we get that

$$v(\hat{x}, \theta) - v(\bar{x}, \theta) \geq (1 + \lambda)(\hat{p} - \bar{p}) \Leftrightarrow \hat{p} \leq \bar{p} + \frac{v(\hat{x}, \theta) - v(\bar{x}, \theta)}{1 + \lambda}. \quad (3.87)$$

$$(1 + \lambda)(c(\bar{x}, \theta) - c(\hat{x}, \theta)) \geq \bar{p} - \hat{p} \Leftrightarrow \hat{p} \geq \bar{p} - (1 + \lambda)(c(\bar{x}, \theta) - c(\hat{x}, \theta)). \quad (3.88)$$

Hence, we get

$$\bar{p} - (1 + \lambda)(c(\bar{x}, \theta) - c(\hat{x}, \theta)) \leq \bar{p} + \frac{v(\hat{x}, \theta) - v(\bar{x}, \theta)}{1 + \lambda}. \quad (3.89)$$

Simplifying (3.88), we get (3.85), as claimed. Thus, $\hat{x} \in R(\bar{x}, \bar{p})$ for any $\hat{x}$ satisfying (3.85).

(iii) In this case, the seller must offer a price reduction to the buyer in order to induce him to voluntarily renegotiate. Thus, the buyer faces a loss in the value dimension and the seller faces a loss in the price dimension. Using (3.83), (3.84) and proceeding as in the proof of part (ii), an easy calculation gives rise to the claimed result. ■

Proposition 3.38 illustrates at least two ideas:

1. In the absence of loss aversion ($\lambda = 0$) each of the renegotiation conditions (3.85) and (3.86) is easier to satisfy. Consider, for instance, (3.85). It might be the case that $v(\hat{x}, \theta) - v(\bar{x}, \theta) \geq c(\hat{x}, \theta) - c(\bar{x}, \theta)$ but $v(\hat{x}, \theta) - v(\bar{x}, \theta) < (1 + \lambda)^2 [c(\hat{x}, \theta) - c(\bar{x}, \theta)]$. Hence, the presence of loss aversion may create stickiness in renegotiation.
2. The results in Proposition 3.38 are independent of prices. This is because the utilities are quasilinear in prices and specifications.

Next, we need to specify how the two parties bargain at date 2. Herweg and Schmidt (2013) assume that the two parties bargain by using generalized Nash bargaining.⁴⁶ For expositional ease, we consider instead the two polar extremes in which either the buyer or the seller have all the bargaining power. In what follows, we consider the case where the buyer has all the bargaining power (as an exercise, the reader can work through the case where the seller has all the bargaining power).

To make further progress, let us consider the following example.

$$v(x, \theta) = \theta x, \quad c(x, \theta) = \frac{1}{2}x^2 + \frac{1}{3}\theta. \quad (3.90)$$

Using (3.78), we get the unique socially optimal (or first best) state-contingent solution

$$x^*(\theta) = \theta. \quad (3.91)$$

⁴⁶ For various bargaining solutions, and the empirical evidence, see Part 4 of the book.

Suppose that the two parties have agreed to a contract $(\bar{x}, \bar{p})$ at date 0. For simplicity, let us normalize $\bar{x} = 1$. Define $\bar{\theta}$ such that

$$\frac{\partial v(1, \bar{\theta})}{\partial x} = \frac{\partial c(1, \bar{\theta})}{\partial x},$$

so, using (3.91), $\bar{\theta} = 1$.

Suppose that at date 1, the realized state of the world is $\theta_H > 1$. How should the two parties renegotiate at date 2? There are two possibilities.

I. For any $\hat{x} > 1$, $v(1, \theta_H) < v(\hat{x}, \theta_H)$ and $c(1, \theta_H) < c(\hat{x}, \theta_H)$. The buyer experiences an increase in value and the seller experiences an increase in cost. This is the case described in Proposition 3.38(ii), hence, $\hat{x} \in R(1, \bar{p})$ if and only if (3.85) holds. Using (3.90), we can write (3.85) as

$$\theta_H(\hat{x} - 1) \geq (1 + \lambda)^2 \frac{1}{2} (\hat{x}^2 - 1) \Leftrightarrow \lambda \leq \sqrt{\frac{2\theta_H}{(\hat{x} + 1)}} - 1. \quad (3.92)$$

In the absence of loss aversion ($\lambda = 0$), the condition in (3.92) is more likely to be satisfied. However, if loss aversion is high enough, then no renegotiation can take place. Using (3.92), $\hat{x} \in R(\bar{x}, \bar{p})$ if

$$\theta_H \geq \frac{(1 + \lambda)^2}{2} (\hat{x} + 1). \quad (3.93)$$

Thus, when, say, $\lambda = 1$, θ_H must exceed the critical value of 4 before any renegotiation can occur. So, we have that at date 3, $x = \bar{x} = 1$ for $1 \leq \theta \leq 4$. However, for any θ in this range, the socially efficient outcome is $x^*(\theta) = \theta$. Thus, loss aversion leads to rigidity in contract renegotiation, in the sense that there is a set of states for which it is socially optimal to renegotiate but privately suboptimal.

Since the buyer is assumed to have all the bargaining power, he ensures that the individual rationality constraint of the seller in (3.88) holds with equality. Thus,

$$\hat{p}(\hat{x}) = \bar{p} + (1 + \lambda) \frac{1}{2} (\hat{x}^2 - 1). \quad (3.94)$$

II. For any $\hat{x} < 1$ we have $v(1, \theta_H) > v(\hat{x}, \theta_H)$ and $c(1, \theta_H) > c(\hat{x}, \theta_H)$. The buyer's value falls and the seller's cost also falls, so this is the case in Proposition 3.38(iii), hence, $\hat{x} \in R(1, \bar{p})$ if and only if (3.86) holds. Using (3.90), we can write (3.86) as

$$\frac{1}{2} (1 - \hat{x}^2) \geq (1 + \lambda)^2 \theta_H (1 - \hat{x}) \Leftrightarrow \lambda \leq \sqrt{\frac{(1 + \hat{x})}{2\theta_H}} - 1. \quad (3.95)$$

Since λ is non-negative, $\hat{x} < 1$ and $\theta_H > 1$, so (3.95) cannot hold. Thus, following a realization of any state $\theta_H > \bar{\theta} = 1$, the two contracting parties never reduce the specification of the good that they trade. The only relevant case, therefore, is Case I.

Substitute $\hat{p}(\hat{x})$ from (3.94) into (3.80) and use the assumed functional forms in (3.90) to get

$$U_B(\hat{x}, \hat{p}(\hat{x}); \theta_H) = \theta_H \hat{x} - \hat{p}(\hat{x}) - \lambda (\hat{p}(\hat{x}) - \bar{p}); 1 < \theta_H \in \Theta. \quad (3.96)$$

The buyer experiences a loss in the price dimension (last term in (3.96)) and a gain in the value dimension, but given (3.79), the gain is set to zero. Since the buyer has all the bargaining power, the ex-post solution $\hat{x} = \hat{x}^*$ maximizes $U_B(\hat{x}, \hat{p}(\hat{x}); \theta_H)$.

$$\frac{\partial U_B(\hat{x}, \hat{p}(\hat{x}); \theta_H)}{\partial \hat{x}} = \theta_H - (1 + \lambda)^2 \hat{x}.$$

It is immediate that U_B is strictly concave and the first order condition is sufficient. Thus,

$$\hat{x}^* = \frac{\theta_H}{(1 + \lambda)^2}. \quad (3.97)$$

There are two noteworthy features of the renegotiated solution under loss aversion. First, the socially optimal solution requires $x^*(\theta) = \theta$. But when renegotiation takes place, we have $\hat{x}^* < x^*$, i.e., the specification is lower than the socially optimal specification. Second, the socially optimal solution requires a one for one adjustment of the specification to θ (because $x^*(\theta) = \theta$), while the renegotiated solution in (3.97) exhibits a less than one for one adjustment to θ . Coupled with the fact that renegotiation takes place for only some values of θ (see (3.93)) it follows that in the presence of loss aversion, renegotiation is inefficient and sluggish.

Substituting $\hat{x}^*$ in (3.94) gives the optimal solution for the renegotiated price, $\hat{p}^* = \hat{p}(\hat{x}^*)$. Thus, the utility of the buyer from the long-term contract in any state of the world, $\theta_H > 1$, is given by $U_B(\hat{x}^*, \hat{p}^*; \theta_H)$. One can proceed analogously to find the utility of the buyer from a long-term contract for any $\theta_L < 1$. Integrating over all states $\theta \in \Theta$ gives the expected utility of the buyer from a long-term contract, U_B^{LT} .

The other possibility is that the parties may not sign a long-term contract, but following the realization of θ at date 1, choose to trade on the spot market. Should the buyer prefer spot markets? The answer depends on three factors.

1. The degree of competition in the spot market: the greater the competition, the lower is the spot price, so the higher is the utility of the buyer.
2. Reference point of the buyer in the spot market: In a long-term contract, the initial contract serves as an entitlement. However, adjustment to the realized state is easier in a spot contract because there is no initial contract. In a simple class of models this may be mathematically similar to having a lower value of loss aversion in spot market contracts (see Herweg and Schmidt, 2013).
3. The degree of loss aversion in long-term contracts. The greater the loss aversion, the higher is the inefficiency and sluggishness of renegotiation in long-term contracts, which reduces their attractiveness.

We refer the reader to Herweg and Schmidt (2013) for a more detailed analysis of these factors.⁴⁷

⁴⁷ The interested reader can also consult Herweg et al. (2014) for an application of the UPE concept to the incomplete contracts problem.

CHAPTER 4

Human Behavior under Ambiguity

4.1 Introduction

Ambiguity arises when individual choices depend on the *source* of the uncertain information, but within each source, the individual possesses reasonable subjective probabilities (non-negative probabilities that add to one). This is sometimes also referred to as *source-dependent uncertainty*. Ambiguity is a major topic, and deserves a book length treatment in its own right. However, the discussion in this chapter is relatively brief.¹ The reason is that the approach of most of the modern literature on ambiguity does not appear behavioral enough. A troubling feature of several of these models is that under pure risk they simply reduce to expected utility theory, which is not supported by the evidence.

For this reason, we have split this chapter into two parts. Section 4.3 outlines *neoclassical models of ambiguity*, and Section 4.4 considers *behavioral models of ambiguity*. There is, by no means, a neat or established separation between the two sets of models. However, the main difference is that all the models in Section 4.4 do not reduce to EU under pure risk, while many of the models in Section 4.3 do.² In order to maintain brevity of discussion, we focus almost exclusively on the *Ellsberg paradox* and consider its resolution in the different models.

In Section 4.2, we consider the classic Ellsberg paradox that illustrates the phenomenon of ambiguity in lab experiments. Suppose that one urn, the *known urn*, is known to have 50 red and 50 black balls. Another urn, the *unknown urn*, has 100 balls of only two colors, red and black, but their exact composition is unknown. As a first pass, one may appeal to the *principle of insufficient reason* to argue that the prior probability of a red and a black ball from the unknown urn is 0.5 each. Indeed, when individuals are asked to bet *separately* on the two urns, the evidence suggests that they assign equal probabilities to drawing either a red or a black ball from either urn. This is the sense in which they have reasonable subjective beliefs about each urn. However, individuals are unwilling to exchange a bet on any color on the known urn with the same bet on the unknown urn. This is the sense in which the subjective beliefs across the two urns are not reasonable. The decision maker is then said to have a preference for avoiding ambiguity (or source-dependent

¹ The reader can also consult chapter 17 in Gilboa (2009), chapter 11 in Wakker (2010), Etner et al. (2012), and Machina and Siniscalchi (2014) for more details.

² This is not true of several models. For instance, Choquet expected utility (CEU) that is considered in Section 4.3 does not reduce to EU for the case of risk.

uncertainty) associated with the unknown urn.³ Thus, SEU cannot be applied because no consistent subjective probabilities exist. The Ellsberg paradox is also a violation of *probabilistic sophistication* (Definition 1.13).

Section 4.3 gives a brief discussion of the neoclassical models of ambiguity and outlines their explanation for the Ellsberg paradox. Section 4.3.1 discusses *non-additive probability measures* (or *capacities* or *Choquet capacities*). Choquet capacities are different from probabilities in that they need not be additive.⁴ This framework gives rise to *Choquet expected utility* (CEU) and is able to explain the Ellsberg paradox.

Section 4.3.2 describes the *multiple priors model*. In this class of models, decision makers have probability distributions over their prior beliefs that reflect their uncertainty about the true model of the world (Gilboa and Schmeidler, 1989). We describe the *maximin model* and its two main axioms. The maximin model has a pessimistic attitude to uncertainty. Among the set of minimum outcomes achieved through all actions, the decision makers choose the action with the maximum outcome. None of the axioms used in deriving such a preference representation gives any explicit clue about this extreme degree of pessimism. We also describe another popular multiple priors model, the α -MEU model (Hurwicz, 1951; Ghirardato et al., 2004).⁵ This class of models can also explain the Ellsberg paradox.

Section 4.3.3 considers the *failure of the reduction of compound lotteries axiom* as an explanation of the Ellsberg paradox in a *two-stage recursive model* (Segal, 1987, 1990). Section 4.3.4 considers the *smooth ambiguity aversion model* (Klibanoff et al., 2005). Unlike some models of multiple priors, such as the *maximin model*, the objective function here is differentiable. The two-stage recursive models and the smooth ambiguity aversion models are both multiple priors models. However, the main difference between the two is that the smooth ambiguity aversion model uses a different utility function at each stage. Ambiguity is captured by the concavity of the second stage utility function and the burden of explaining risk and ambiguity attitudes in this model lies with the utility function.

Section 4.3.5 gives a brief account of other models that can account for the Ellsberg paradox. Section 4.3.6 shows that the empirical evidence for neoclassical models of ambiguity aversion is mixed (Hey et al., 2010; Conte and Hey, 2013). No single model can account for all the experimental results. What causes ambiguity aversion? Some experimental evidence suggests that *fear of negative evaluation by others* and one's competence in the problem could be important explanations (Curley et al., 1986; Trautmann et al. 2008).

Section 4.4 considers the small literature on behavioral models of ambiguity. None of these models reduce to EU under pure risk. In Section 4.4.1, we consider promising work that was conducted in the 1990s on *support theory* but was ignored by the subsequent literature in economics. This work is associated most with the names of Amos Tversky and his collaborators. In *support theory*, the support for one hypothesis over another depends on how the two hypotheses are framed and how events are subdivided. For instance, the probability that one assigns to death by homicide in an area might be lower than the sum of probabilities that one assigns to homicide by friends and homicide by strangers, although classically the two answers should be identical. Splitting the event into its components, brings more such instances to mind (as in the

³ It is possible to define ambiguity generally as a preference for information from one source relative to another; see Definition 11.1.1 in Wakker (2010).

⁴ Lack of additivity means that for two disjoint events, E_i, E_j , in some well-defined sample space, the subjective probability $\mu(E_i \cup E_j)$ need not equal $\mu(E_i) + \mu(E_j)$.

⁵ Luce and Raiffa (1957) already discuss multiple priors (maximin EU) and the α -maximin model (see their Ch. 13 on ambiguity, with §13.5 pp. 304–5). I am grateful to Peter Wakker for pointing this out.

availability heuristic), hence, increasing its perceived probability. In support theory, in the first stage, individuals form beliefs about events by using such considerations. Once these beliefs are formed, then one evaluates the uncertain situation using PT.

Section 4.4.2 considers the possibility that decision makers perceive information from different *sources* of uncertainty in a fundamentally different manner (Abdellaoui, Baillon et al., 2011). Applied to the Ellsberg paradox, in this view, each urn is a different source of uncertainty. This accounts for the unwillingness of decision makers to exchange bets on the known urn (first source) with bets on the unknown urn (second source). Empirical results presented by the authors confirm this view, although they also reveal a rich diversity in the attitudes of individuals to uncertainty. In particular, the differences in the two Ellsberg urns cannot be ascribed to differences in the utility functions but rather to the elicited probability weighting functions (or source functions) across them. In an important study that applies these concepts, PT outperforms all others in a 11-way horse race (Kothiyal et al., 2014). This strengthens the case for PT to apply to risk, uncertainty, and ambiguity, which is a central theme that runs through Wakker (2010).

The explanation of ambiguity in these behavioral models is different from neoclassical models, such as the smooth ambiguity aversion model. The neoclassical models locate the source of ambiguity aversion in the shape of the utility function. This is analogous to the feature of EU under risk and uncertainty, which places the entire burden of risk attitudes on the shape of the utility function; a finding that is rejected by the evidence. On the other hand, non-EU theories such RDU and PT rely jointly on the shapes of the utility function and the probability weighting function in capturing attitudes to risk; a finding that is consistent with the evidence.

4.2 A problem for SEU: The Ellsberg paradox

Under uncertainty, there is no problem in applying EU, RDU, or PT (see Sections 1.3, 2.3.2, and 2.4) provided that decision makers have well-defined subjective probabilities. However, evidence suggests that this assumption might not always hold. The most famous example of this is the Ellsberg (1961) paradox. We present two variants of the Ellsberg paradox below; the two-colors and the three-colors experiment.

Example 4.27 (*Two-colors experiment*): Consider the simplest demonstration of the Ellsberg paradox, the two-colors experiment. Urn 1 has 50 red and 50 black balls, while Urn 2 contains 100 balls that are an unknown mixture of red and black balls.

- (1) Suppose that subjects are asked to bet on a color (a red or a black ball) for each urn and they receive a monetary prize, z , if a ball of their chosen color is drawn. Subjects reveal, through their choices, that they are indifferent among the two colors for each urn.⁶ Thus, they appear to assign equal probabilities to the two colors in both urns.
- (2) Now suppose that subjects, in addition, are also asked which urn they would like to bet on. Irrespective of their chosen color, they typically choose Urn 1 over Urn 2. Let p_r and p_b represent the subjective probabilities of a red and a black ball, respectively, in Urn 2 such that $p_r + p_b = 1$. Then, the indifference in (1) implies that $p_r = 0.5$. However, if the chosen color is red, the preference for Urn 1 implies that $0.5u(z) > p_ru(z) \Leftrightarrow p_r < 0.5$. Analogously if the chosen color is black, the stated

⁶ The original Ellsberg experiments were not incentivized and are thought experiments. However, once these potential drawbacks are addressed by modern experimental methods (see below) the results still survive.

Table 4.1 Description of the four Savage acts in the three-colors experiment.

	r	b	y
f_1	100	0	0
f_2	0	100	0
f_3	100	0	100
f_4	0	100	100

preference implies $0.5u(z) > p_b u(z) \Leftrightarrow p_b < 0.5$. This contradicts $p_r = p_b = 0.5$ from (1). Hence, the subjects do not hold reasonable subjective beliefs.⁷ Furthermore, this violates probabilistic sophistication (see Definition 1.13).

Example 4.28 (Three-colors experiment) In the three-colors experiment, subjects are told that an urn contains 90 balls of 3 different colors: red, black, and yellow. There are 30 red balls and the remaining balls are either black or yellow. One ball is drawn at random from the urn. Let r, b, y be the events: red ball, black ball, and yellow ball. The minimum “reasonableness” that we might expect from subjective beliefs is that they should satisfy an additive probability measure, q , such that

$$q(r) = \frac{1}{3}, q(b) = 1 - p - \frac{1}{3}; q(y) = p; p \in [0, 2/3]. \quad (4.1)$$

The decision maker is offered a choice between the following two pairs of lotteries or Savage acts (see Section 1.3).

$$f_1 = (r, 100; b, 0; y, 0), f_2 = (r, 0; b, 100; y, 0), \quad (4.2)$$

$$f_3 = (r, 100; b, 0; y, 100), f_4 = (r, 0; b, 100; y, 100). \quad (4.3)$$

Thus, for instance, the Savage act f_2 is the lottery that pays \$100 if a black ball turns up but \$0 if any other color comes up. Analogously, the Savage act f_4 is the lottery that pays \$100 if either a black or a yellow ball is drawn and \$0 if a red ball comes up. This information is summarized in Table 4.1.

Subjects in experiments, typically express the preferences $f_1 \succ f_2$ and $f_4 \succ f_3$. Assuming SEU preferences (see Section 1.3), these preferences imply the following inequalities.

$$f_1 \succ f_2 \Leftrightarrow \frac{1}{3}u(100) > \left(1 - p - \frac{1}{3}\right)u(100) \Leftrightarrow p > \frac{1}{3}. \quad (4.4)$$

$$f_4 \succ f_3 \Leftrightarrow \frac{1}{3}u(100) + pu(100) < \left(1 - p - \frac{1}{3}\right)u(100) + pu(100) \Leftrightarrow p < \frac{1}{3}. \quad (4.5)$$

From (4.4) and (4.5), the decision maker does not assign consistent subjective probabilities. This is illustration of the Ellsberg paradox.

⁷ This problem was also discussed by Keynes (1921).

The examples suggest the following observations.

1. The preferences $f_1 \succ f_2$ and $f_4 \succ f_3$ suggest that the decision maker tries to search for greater certainty in the probability of the outcome. For instance, the preference $f_1 \succ f_2$ may suggest that the decision maker wishes to avoid the possibility that the probability of the black ball is very small or zero. Similarly, in the preference $f_4 \succ f_3$, the decision maker knows that the probability of getting a prize of 100 under f_4 is $2/3$, while under f_3 it might just be $1/3$. This is the sense in which the individual is *ambiguity averse*.
2. From Table 4.1, the outcomes under the two Savage acts f_1, f_2 are identical if a yellow ball comes up but $f_1 \succ f_2$. Similarly, the outcomes under the two Savage acts f_3, f_4 are identical if a yellow ball comes up. If Savage's *sure thing principle* (see Definition 1.11) is applied we should observe $f_3 \succ f_4$. However, we observe $f_4 \succ f_3$. Hence, the Ellsberg paradox is essentially a violation of the sure thing principle. For this reason, explanations of the Ellsberg paradox have typically tried to relax this principle.
3. The Ellsberg paradox violates probabilistic sophistication (Definition 1.13). For instance, in the two-colors experiment in Example 4.27 above, for each urn, separately, $p_r = 0.5$ but when the two urns are compared, the decision maker does not behave as if $p_r = 0.5$. It is as if the probabilities depend on the act itself.⁸ Suppose that an individual bets on the color red when presented with each urn separately. Let x_i denote the event that the a ball of color $x = \text{Red, Black}$ is drawn from urn $i = 1, 2$. Then for each urn, the expressed preference is $(\text{Red}1, z; \text{Black}1, 0) \sim (z, 0.5, 0, 0.5)$ and $(\text{Red}2, z; \text{Black}2, 0) \sim (z, 0.5, 0, 0.5)$, where $\sim$ is the indifference relation, so probabilistic sophistication holds for each urn separately. But $(\text{Red}1, z; \text{Black}1, 0) \succ (\text{Red}2, z; \text{Black}2, 0)$, so $(z, 0.5, 0, 0.5) \succ (z, 0.5, 0, 0.5)$, where $\succ$ is the strict preference relation; this is a contradiction of probabilistic sophistication.

We now consider alternative explanations of the Ellsberg paradox within two classes of models. The neoclassical approach and the behavioral approach.

4.3 Neoclassical models of ambiguity

In this section, we consider models of ambiguity that are based on the neoclassical tradition. There is no clear dividing line between models in the neoclassical and the behavioral traditions, so any such division is necessarily arbitrary. However, in the presence of pure risk, many of the models in this section collapse to expected utility, but clearly some such as Choquet expected utility do not. Given the refutations of EU (see Chapter 1), this is clearly a troublesome feature of many of these models (Wakker, 2010); and some may argue, a direct refutation of these models. By contrast, the behavioral models of ambiguity (see Section 4.4) do not typically reduce to EU under risk. Furthermore, the main difference between the neoclassical and the behavioral models of ambiguity is that, in the behavioral models, ambiguity is *source-dependent*.

The following two comments on this class of models illustrate the reservations that many behavioral economists have with the neoclassical approach to ambiguity, despite its immense popularity and widespread acceptance. Hey (2014, p. 838) writes: "We remain agnostic as to whether these can be really considered as models of ambiguity." Kothiyal et al. (2014, p. 3) write: "Since 1990, quantitative ambiguity models have become available (Gilboa 1987;

⁸ Probabilistic sophistication was originally introduced by Machina and Schmeidler (1992). For subsequent developments, see Grant (1995) and Chew and Sagi (2006). See also the exposition in Wakker (2010).

Gilboa & Schmeidler 1989; Schmeidler 1989), and many extensions have been proposed . . . These models were all normatively motivated . . . assuming expected utility for risk, backward induction in an Anscombe-Aumann (1963) model, and (with α -maximin excepted) universal ambiguity aversion. These assumptions, even if accepted normatively, fail descriptively.” For this reason, we provide a relatively brief treatment of these models, focusing mainly on their approach to the Ellsberg paradox.

4.3.1 Choquet expected utility

One resolution to the Ellsberg paradox is to employ non-additive probability measures. Consider a sample space, S and a sigma algebra Σ of the subsets of S .

Definition 4.64 A sigma algebra Σ corresponding to a sample space S is a set of subsets of S such that the following hold. (i) The empty set $\emptyset \in \Sigma$. (ii) If $A \in \Sigma$, then $A^c \in \Sigma$, where A^c is the complement of A . (iii) For countably many sets $A_1, A_2, \dots, A_n, \dots \in \Sigma$, the union $A_1 \cup A_2 \cup \dots \cup A_n \cup \dots \in \Sigma$.

The non-additive measure, $\varsigma : \Sigma \rightarrow [0, 1]$, also sometimes called a *capacity* or *Choquet capacity*, has the following properties.

1. $\varsigma(\emptyset) = 0$, where $\emptyset$ is the empty set.
2. If $A \subseteq B$, then $\varsigma(A) \leq \varsigma(B)$; $A, B \in \Sigma$.
3. $\varsigma(S) = 1$.

Thus, as with probability measures, the probability assigned to the empty set, $\emptyset$, is zero, while the probability assigned to the entire sample space, S , is one. However, the main difference lies in the weaker requirement in Property 2. For an additive probability measure, say, p , and for disjoint events A, B , $p(A \cup B) = p(A) + p(B)$. However, Property 2 does not impose any such condition for a non-additive probability measure. The triple (S, Σ, ς) is known as the *capacity space*.

Recall that for any partition of a sample space, S , into a set of *mutually exclusive* events $\{E_1, E_2, \dots, E_n\}$, a Savage act $f : S \rightarrow X$. Suppose that the outcomes or prizes corresponding to the events $E_1, E_2, \dots, E_n$ are, respectively, $z_1, z_2, \dots, z_n$. We assume that the events are arranged, so that $0 < z_1 < z_2 < \dots < z_n$ and $f = \{E_1, z_1; E_2, z_2; \dots; E_n, z_n\}$. Let F be the set of all Savage acts (or lotteries) under consideration. For non-additive measures, the Riemann integral is not adequate (Gilboa, 2009, pp. 148–9). The *Choquet integral* that we define next, turns out to be useful in this regard.

Definition 4.65 (*Choquet integral*): Consider the capacity space (S, Σ, ς) . Define a Savage act $f = \{E_1, z_1; E_2, z_2; \dots; E_n, z_n\}$, $0 < z_1 < z_2 < \dots < z_n$, over a sigma algebra, Σ , of the sample space S . The Choquet integral of f with respect to the Choquet capacity ς is given by:

$$\int_S f d\varsigma = \sum_{j=1}^n z_j \left[\varsigma \left(\bigcup_{k=j}^n E_k \right) - \varsigma \left(\bigcup_{k=j+1}^n E_k \right) \right]. \quad (4.6)$$

If ς were additive, then $\varsigma \left(\bigcup_{k=j}^n E_k \right) = \sum_{k=j}^n \varsigma(E_k)$, so (4.6) can be written as

$$\int_S f d\varsigma = \sum_{j=1}^n z_j \varsigma(E_j),$$

which is simply the Riemann integral.

Definition 4.66 (Comonotonicity): Consider two Savage acts $f = (s_1, z_1; \dots; s_n, z_n)$, and $g = (s_1, x_1; \dots; s_n, x_n)$. Then f and g are comonotonic if for all i, j

$$f(s_i) < f(s_j) \Rightarrow g(s_i) < g(s_j),$$

where $f(s_i) = z_i$ and $g(s_j) = x_j$ for all $i, j = 1, \dots, n$.

For our next definitions, we need to define mixtures over horse-roulette acts. Hence, we introduce the following terminology (Machina and Siniscalchi, 2014). Recall that we represented a lottery under objective probabilities in the *one-stage form* as $L = (x_1, p_1; \dots; x_n, p_n)$. Consider a partition $E_1, \dots, E_m$ of the sample space. A *horse-roulette act* or an *Anscombe–Aumann act* (or an *objective–subjective prospect*) takes the form, $(L_1, E_1; \dots; L_m, E_m)$ with the interpretation that the decision maker plays the objective lottery L_i if the event E_i takes place. Thus, horse-roulette acts are mappings from the sample space to the objective, single-stage, lotteries. These sorts of lotteries were used by Anscombe and Aumann (1963) in providing an axiomatic derivation of subjective expected utility (SEU) that is simpler than the axiomatization in Savage (1954).

Recall that mixtures over objective one-stage lotteries are defined as follows. If $L_1 = (x_1, p_1; \dots; x_n, p_n)$ and $L_2 = (y_1, q_1; \dots; y_m, q_m)$, then for $\alpha \in (0, 1)$ the $\alpha : (1 - \alpha)$ probability mixture of L_1 and L_2 , is denoted by

$$\alpha L_1 + (1 - \alpha) L_2 = (x_1, \alpha p_1; \dots; x_n, \alpha p_n; y_1, (1 - \alpha) q_1; \dots; y_m, (1 - \alpha) q_m).$$

If the reduction of compound lotteries axiom holds, then one can reduce the mixture lottery to an objective one-stage lottery such that for any common prize, $x_i = y_i$, in the two lotteries L_1 and L_2 , the corresponding probability is $\alpha p_i + (1 - \alpha) q_i$.

In most models of ambiguity aversion, preferences are defined over horse-roulette acts (or simply “acts”). Probability mixtures over such acts are constructed as follows. Suppose that we have a common partition $\{E_1, \dots, E_n\}$ of the sample space. For two acts $f = (L_1, E_1; \dots; L_n, E_n)$ and $g = (L'_1, E_1; \dots; L'_n, E_n)$, and $\alpha \in (0, 1)$, the $\alpha : (1 - \alpha)$ mixture of f and g , is defined as

$$\alpha f + (1 - \alpha) g = (\alpha L_1 + (1 - \alpha) L'_1, E_1; \dots; \alpha L_n + (1 - \alpha) L'_n, E_n).$$

Definition 4.67 (Comonotonic independence): Consider three acts f, g, h that satisfy pairwise comonotonicity and let $\succeq$ denote the decision maker’s preference relation. Then comonotonic independence implies that

$$f \succeq g \Leftrightarrow \alpha f + (1 - \alpha) h \succeq \alpha g + (1 - \alpha) h \text{ for all } \alpha \in (0, 1).$$

Schmeidler (1986, 1989) demonstrated axiomatically that if the relevant notion of independence is comonotonic independence, then,

$$f \succeq g \Leftrightarrow \int_S u(f) d\zeta \geq \int_S u(g) d\zeta,$$

where the relevant notion of integration is Choquet integration (Definition 4.65), ς is a capacity and u is a non-constant function such that $u : X \rightarrow \mathbb{R}$. $\int_S u(f) d\varsigma$ is termed as the Choquet expected utility (CEU) of the act, f .⁹

Allowing for subjective, non-additive, beliefs raises problems for the Pareto criterion in welfare economics. Consider, for instance, the following hypothetical problem posed by Gilboa et al. (2004). In an age of chivalry, two gentlemen, A and B , are considering a resolution of a nagging dispute by engaging in a duel. The objective probability that A wins is p and that B wins is $1 - p$. However the two gentlemen do not know the objective probabilities. Given their tastes, each prefers to engage in the duel if their respective probability of winning is greater than 0.8. Suppose that A and B assign respective subjective probabilities of their winning the duel at 0.9. It is, therefore, privately optimal for each to fight the duel. However, there is no pair of objectively shared probabilities $p, 1 - p$ such that it is optimal to fight. Hence, it may not seem justified for society to accept the individual preferences for a duel, knowing that at least one of the gentlemen has got his beliefs wrong.

Let us now give an explanation of the preferences observed in the three-colors experiment in Example 4.28. Under the four acts in Example 4.28, the sample space is

$$S = \{\{r\}, \{b\}, \{y\}, \{r, b\}, \{r, y\}, \{b, y\}\}.$$

Suppose that we assign the following capacities to these events

$$\varsigma(r) = \varsigma(rb) = \varsigma(ry) = \frac{1}{3}; \varsigma(b) = \varsigma(y) = 0; \varsigma(by) = \frac{2}{3},$$

which satisfy the conditions for a capacity. Let us now compute the CEU corresponding to the two preferences and let $u(0) = 0$.

$$\begin{aligned} \int_S u(f_1) d\varsigma - \int_S u(f_2) d\varsigma &= \frac{1}{3}u(100) - 0 > 0 \Leftrightarrow f_1 \succ f_2. \\ \int_S u(f_3) d\varsigma - \int_S u(f_4) d\varsigma &= \frac{1}{3}u(100) - \frac{2}{3}u(100) < 0 \Leftrightarrow f_4 \succ f_3. \end{aligned}$$

Thus, we are able to explain the observed preference patterns in the Ellsberg paradox if we allow for non-additive probability measures.

4.3.2 Models of multiple priors

Multiple priors models offer one way of capturing ambiguity. Instead of point estimates of the prior probability, decision makers might assign a probability distribution over the prior probabilities. This could arise because decision makers may not fully trust any particular prior or they may have many possible models of the world in mind, each model having some probability of being the true one.

⁹ For other axiomatizations of CEU, see Gilboa (1987), Sarin and Wakker (1992), Nakamura (1990), Ghirardato et al. (2003), and Lehrer (2009).

Gilboa and Schmeidler (1989) propose the following axiomatized model, which can explain the Ellsberg paradox.¹⁰ Consider some act $f \in F$, where F is the set of all acts under consideration. Suppose that the state space is $S = \{s_1, s_2, \dots, s_n\}$ and let Δ be the set of possible subjective probability distributions over S . Let Λ be some closed and convex subset of Δ and denote a typical element of Λ by $\rho = (\rho(s_1), \rho(s_2), \dots, \rho(s_n))$. The set Λ captures all priors that the decision maker has about the set of states; each prior may be taken to denote a different model of the world. Now consider two acts f, g . Then, under certain conditions (see below), Gilboa and Schmeidler (1989) showed that:

$$f \succeq g \Leftrightarrow \min_{\rho \in \Lambda} \int_S u(f(s)) d\rho(s) \geq \min_{\rho \in \Lambda} \int_S u(g(s)) d\rho(s), \quad (4.7)$$

where $u : X \rightarrow R$ is a non-constant function. In words, (4.7) says the following. For each possible act, compute the minimum expected utility for that act from the set of all possible priors under consideration. Then choose the act which gives the maximum utility from the set of minimum expected utilities obtained in the first step. This is also called as the *maximin expected utility model* (MEU).¹¹

In order to derive the representation in (4.7), Gilboa and Schmeidler (1989) weakened the independence axiom in the Anscombe–Aumann framework to *C-independence*, and they also required a notion of *uncertainty aversion*. We now define these concepts.

Definition 4.68 (*C-independence*): Suppose that f, g are two acts and h is some constant act that gives the same outcome in each state of the world. Then *C-independence* requires that

$$f \succeq g \Leftrightarrow \alpha f + (1 - \alpha)h \succeq \alpha g + (1 - \alpha)h \text{ for all } \alpha \in (0, 1).$$

Notice the difference between *C-independence* and comonotonic independence: the acts here are not required to be pairwise comonotonic and h is a constant act.

Definition 4.69 (*Uncertainty aversion*): Suppose that the decision maker is indifferent between two acts f, g , i.e., $f \sim g$, then

$$\delta f + (1 - \delta)g \succeq f \text{ for all } \delta \in (0, 1).$$

Roughly speaking, uncertainty aversion requires that the decision maker would like to hedge bets by preferring a combination of two equivalent acts to any one of the acts.

The representation in (4.7) exhibits extreme pessimism because the decision maker focuses on the worst possible expected utility under each prospect. However, the two conditions, *C-independence* and *uncertainty aversion*, do not directly suggest such pessimism.

Consider now the three-color Ellsberg experiment in Example 4.28. For the four lotteries in (4.2), (4.3), using the most pessimistic subjective probability distribution, we get that for $p \in [0, 2/3]$ (the joint probability of a black and yellow ball) the minimum expected utility for each of the four acts, f_1, f_2, f_3, f_4 is:

¹⁰ There are several axiomatizations of the multiple priors model. See also, Bewley (2002), Klibanoff et al. (2005), Maccheroni et al. (2006a,b).

¹¹ For a generalization of the maximin criteria, see Chateauneuf and Faro (2009). These authors restrict the set of allowable priors to only those on which the decision maker can put a certain minimum degree of confidence. One undesirable feature of MEU models is that they violate monotonicity (see exercise 11.10.1 in Wakker, 2010).

$$\begin{cases} f_1 : & \frac{1}{3}u(100) \\ f_2 : & (1 - p - \frac{1}{3})u(100) = 0 \text{ for } p = 2/3 \\ f_3 : & [\frac{1}{3}u(100) + pu(100)] = \frac{1}{3}u(100) \text{ for } p = 0 \\ f_4 : & [(1 - p - \frac{1}{3})u(100) + pu(100)] = \frac{2}{3}u(100) \text{ for } p = 0 \text{ or } p = 2/3 \end{cases}$$

It is clear from these calculations that

$$\min_{\rho \in \Lambda} \int_S u(f_1(s))d\rho(s) \geq \min_{\rho \in \Lambda} \int_S u(f_2(s))d\rho(s), \text{ so } f_1 \succ f_2,$$

$$\min_{\rho \in \Lambda} \int_S u(f_4(s))d\rho(s) \geq \min_{\rho \in \Lambda} \int_S u(f_3(s))d\rho(s), \text{ so } f_4 \succ f_3,$$

which is in agreement with the observed choices: f_1 over f_2 and f_4 over f_3 .

The models of multiple priors and the MEU approach has also been applied to macroeconomics in the work of Hansen and Sargent (2001a,b). Roughly speaking, the idea is as follows. Suppose that a policymaker has uncertainty about which model best describes the real world. For instance, the policymaker could be uncertain about the relevant coefficients of the expectations augmented Phillips curve. The policymaker must make a choice between $n > 1$ alternative policies. Then the policymaker computes, first, the maximum loss that could arise by using each of the policies. At the second stage, the policymaker chooses the policy that gives the least loss from among the set of losses computed in the first step. For more on the precise link between *robust control methods* of this sort on the one hand and the maximin model, see Maccheroni et al. (2006b).

Some of the assumptions behind the MEU model have been relaxed. For instance, Maccheroni et al. (2006a) relax C-independence. By relaxing the assumption of uncertainty aversion (Definition 4.69 above) in the MEU model, Ghirardato et al. (2004) obtained the α -MEU model. In this model, any act f is evaluated in the following manner:

$$V(f) = \alpha(f) \min_{\rho \in \Lambda} \int_S u(f(s))d\rho(s) + [1 - \alpha(f)] \max_{\rho \in \Lambda} \int_S u(f(s))d\rho(s) \quad (4.8)$$

where Λ is some closed and convex subset of the set of possible probability distributions, Δ , over the set of states. The parameter $\alpha(f) \in [0, 1]$ depends on the act, f , and it captures the decision maker's *ambiguity aversion* for the act. For $\alpha(f) = 1$, this is simply the MEU model. For $\alpha(f) = 0$, this representation turns out to be the same as that for a special case of incomplete preferences as in Bewley (2002), which we examine in Section 4.3.5. In this case, it is sometimes known as the maximax EU model. It has recently been shown that the preference foundations for the α -MEU model are quite restricted. Eichberger et al. (2011) show that the proposal by Ghirardato et al. (2004) holds only for $\alpha = 0$ and $\alpha = 1$. Siniscalchi (2006) shows that α cannot be interpreted as a parameter of ambiguity aversion.

We note that (4.8) is reminiscent of the *Arrow–Hurwicz criterion*, which preceded the axiomatic developments. Arrow and Hurwicz (1972) generalized in turn, the Wald criterion. The latter advocated looking at the worst possible outcome of an act. The generalization of Arrow–Hurwicz was to advocate a weighted average of the worst and best outcome under each act.

4.3.3 Two-stage recursive models and failure of compounding

Segal (1987, 1990) provided an explanation for the Ellsberg paradox within the class of two-stage recursive models where subjects violate the reduction of compound lotteries axiom.¹² Consider the following simplified Ellsberg setup.

Example 4.29 Suppose that there are two urns. Urn 1 is known to contain 5 red and 5 black balls. Urn 2 has 10 balls that are either red or black, but their composition is not known. In this sense, as in the original setup of the Ellsberg paradox, there is ambiguity associated with Urn 2, while Urn 1 involves known risks.

The decision maker is offered a monetary prize, z , if a red ball comes up in a random draw. No prize is given if a black ball comes up. Suppose that the decision maker uses RDU. Then, the utility from betting on a red ball from Urn 1, denoted by RDU_1 , is

$$RDU_1 = u(z)w(0.5), \quad (4.9)$$

where $w(p)$ is a probability weighting function. Suppose that the decision maker has the following mental models of the color composition of Urn 2.

Definition 4.70 With probability $\rho > 0$, Urn 2 has 10 red balls (Model 1). With probability ρ , Urn 2 has 0 red balls (Model 2). Finally, with the complementary probability $1 - 2\rho$, Urn 2 has 5 red and 5 black balls (Model 3). We denote Model i by M_i , $i = 1, 2, 3$.

The probability of a red ball in the three models is, respectively, ρ , 0, $(1 - 2\rho)/2$. A decision maker who respects the reduction axiom, computes the total probability of a red ball to be $\rho + 0 + (1 - 2\rho)/2 = 0.5$, and the total probability of a black ball to be 0.5. Thus, such a decision maker evaluates Urn 2, based on his mental model in Definition 4.70, as the lottery

$$L = (0, 0.5; z, 0.5). \quad (4.10)$$

If a decision maker who follows RDU did respect compounding, then he would evaluate the lottery in (4.10) drawn from Urn 2, and based on the mental models in Definition 4.70, as $RDU(L) = u(z)w(0.5)$.

Suppose now that the decision maker does not respect the reduction axiom, so is unable to reduce compound lotteries in the usual manner. Further, the decision maker follows RDU. Segal (1987, 1990) provided the following formalization of the departure from compounding, in the context of RDU. Assume that $w(p)$ is convex.¹³ The certainty equivalent of the lottery $L = (0, 1 - p; z, p)$ is a prize, c , earned with probability 1, such that

$$u(c) = u(z)w(p). \quad (4.11)$$

$$\Rightarrow c = u^{-1}(u(z)w(p)). \quad (4.12)$$

¹² For a list of two-stage models that preceeded the work of Segal and suggest giving up compounding so as to model ambiguity, see the annotated bibliography on Peter Wakker's homepage under the heading "second-order probabilities to model ambiguity."

¹³ Convexity of the weighting function has certain desirable properties under RDU. For instance, it is a necessary condition for risk aversion (see Chew et al., 1987). For other desirable properties, see Segal (1987).

The decision maker essentially faces a two-stage lottery when considering the draw of a ball from Urn 2. In the first stage, there is risk on account of three possible mental models described in Definition 4.70 (these occur with respective probabilities $\rho, 0, 1 - 2\rho$). For each mental model, in the second stage, he faces another risky situation because the prize z is won with some probability (only in the first mental model is it won with certainty).

The decision maker violates the reduction axiom in the following manner. For each mental model, he replaces the second stage lottery by its certainty equivalent, which violates compounding. Denote by $c_i, i = 1, 2, 3$, the certainty equivalent in model M_i , then

$$\begin{cases} c_1 = u^{-1}(u(z)w(1)) = u^{-1}(u(z)) = z, \\ c_2 = u^{-1}(u(z)w(0)) = u^{-1}(0) = 0, \\ c_3 = u^{-1}(u(0)[1 - w(0.5)] + u(z)w(0.5)) = u^{-1}(u(z)w(0.5)). \end{cases}$$

By replacing the second stage lottery by its certainty equivalent, and rolling back the decision tree, the individual faces the following lottery. The outcomes are arranged in an increasing order, as required by RDU (w is convex, so $u^{-1}(u(z)w(0.5)) < u^{-1}(u(z)) = z$).

$$(0, \rho; u^{-1}(u(z)w(0.5)), 1 - 2\rho; z, \rho). \quad (4.13)$$

We can now compute the utility of a RDU decision maker, when faced with the lottery in (4.13). This is the utility of the decision maker from betting on a red ball from Urn 2 (the ambiguous urn in the classical terminology), when he employs the mental models in Definition 4.70, and reduces lotteries in the manner described above. Under RDU, this utility, denoted by RDU_2 , is:

$$\begin{aligned} RDU_2 &= u(u^{-1}(u(z)w(0.5))) [w(1 - \rho) - w(\rho)] + u(z)w(\rho) \\ &= u(z) [w(0.5)w(1 - \rho) + w(\rho)(1 - w(0.5))]. \end{aligned} \quad (4.14)$$

Using the assumed convexity of the weighting function we get that

$$\begin{aligned} w(0.5)w(1 - \rho) + w(\rho)(1 - w(0.5)) &< w(w(0.5)(1 - \rho) + (1 - w(0.5))\rho) \\ &= w(w(0.5)(1 - 2\rho) + \rho) \\ &< w(w(0.5)(1 - 2\rho)) \quad (\because w \text{ is increasing}) \\ &< w(w(0.5)) \quad (\because (1 - 2\rho) \in (0, 1)) \\ &< w(0.5) \quad (\because w(0.5) < 0.5). \end{aligned} \quad (4.15)$$

From (4.14), (4.15), and using (4.9), we get that

$$RDU_2 < u(z)w(0.5) = RDU_1. \quad (4.16)$$

From (4.16), the decision maker prefers to bet on a red ball from Urn 1 rather than Urn 2. Hence, using the classical terminology, one can explain a preference for the risky urn relative to the ambiguous urn, solely by a failure of a decision maker, using RDU, to understand or invoke compounding, when the decision maker has more than one model of the world as in Definition 4.70.

Halevy (2007) found evidence consistent with the predictions of models that allow for a failure of compounding. These models include those of Segal (1987, 1990), Klibanoff et al. (2005), and

Halevy and Feltkamp (2005). Halevy found that most subjects in experiments who were able to reduce compound lotteries were ambiguity neutral. Conversely, most subjects who were ambiguity neutral were able to reduce compound lotteries. Over the entire population of subjects in experiments, none of the existing theories can explain all the experimental findings. However, the results do strongly suggest that the reduction of compound lotteries must be an important part of the explanation of the Ellsberg paradox.

4.3.4 The smooth ambiguity aversion model

The *smooth ambiguity aversion model* proposed by Klibanoff et al. (2005) belongs to a similar class of models as Segal (1987, 1990). The main differences are the nature of lotteries in the two frameworks (objective versus subjective lotteries) and the nature of preferences in the two stages (identical versus different preferences). Consider the three-colors Ellsberg experiment in Example 4.28. The decision maker has a set of priors, or initial belief, or *first order beliefs*, given in (4.1). These are given by the distribution function Q^1 with the associated probabilities: $q^1(r) = \frac{1}{3}$, $q^1(b) = \frac{2}{3} - p$; $q^1(y) = p$; $p \in [0, 2/3]$. There are an infinite number of such priors because $p \in [0, 2/3]$. Suppose that the decision maker has *second order beliefs* captured by the distribution Q^2 about the set of possible priors $p \in [0, 2/3]$.

Preferences over the initial or first order beliefs are represented by a von Neumann–Morgenstern utility function, u , and $u(0) = 0$. Preferences over second order beliefs are represented by an increasing transformation, ϕ , of u , i.e., $\phi(u) : \mathbb{R} \rightarrow \mathbb{R}$. Then, the observed preferences in the Ellsberg paradox (i.e., $f_1 \succ f_2$ and $f_4 \succ f_3$) are represented by:

$$f_1 \succ f_2 \Leftrightarrow \phi\left(\frac{1}{3}u(100)\right) > \int_{-\infty}^{\infty} \phi\left(\left(\frac{2}{3} - p\right)u(100)\right) dQ^2, \quad (4.17)$$

$$f_4 \succ f_3 \Leftrightarrow \phi\left(\frac{2}{3}u(100)\right) > \int_{-\infty}^{\infty} \phi\left(\frac{1}{3}u(100) + pu(100)\right) dQ^2. \quad (4.18)$$

In this framework, u captures the individual's attitudes to risk. However, ϕ measures attitudes to ambiguity, i.e., the uncertainty faced by the decision maker about the set of priors. Ambiguity aversion is associated with concavity of ϕ , i.e., $\phi' > 0$, $\phi'' < 0$. The preferences expressed in (4.17), (4.18) are reasonable; this is shown in the next example.

Example 4.30 (Machina and Siniscalchi, 2014): Suppose that $u(100) = 1$, $u(0) = 0$, $\phi(u) = \sqrt{u}$. The distribution Q^2 assigns a probability 0.5 each to one of two possible Q^1 distributions. In the first distribution, $p = \frac{2}{3}$, and the second, $p = 0$. We can now rewrite (4.17), (4.18) as follows.

$$\begin{aligned} f_1 \succ f_2 &\Leftrightarrow \phi\left(\frac{1}{3}\right) > \frac{1}{2}\phi\left(\frac{2}{3}\right) \Leftrightarrow 0.5774 > 0.4083. \\ f_4 \succ f_3 &\Leftrightarrow \phi\left(\frac{2}{3}\right) > \frac{1}{2}\phi(1) + \frac{1}{2}\phi\left(\frac{1}{3}\right) \Leftrightarrow 0.8165 > 0.7887. \end{aligned}$$

Hence, there is no paradox in this case.

More generally, let the initial priors over the set of states in the sample space, S , be captured by the distribution, Q^1 . An ambiguity averse decision maker's preferences over two acts f, g in the smooth ambiguity aversion model are represented by:

$$f \succ g \Leftrightarrow \int_{\Lambda} \phi \left(\int_S u(f(s)) dQ^1 \right) dQ^2 > \int_{\Lambda} \phi \left(\int_S u(g(s)) dQ^1 \right) dQ^2, \quad (4.19)$$

where u is a von Neumann–Morgenstern utility function, $\phi''(u) < 0$, the distribution Q^1 defined over S , captures prior beliefs, and the distribution Q^2 defined over Λ , captures uncertainty about prior beliefs. In a sense, the decision maker applies subjective expected utility twice; once over the prior distribution with von Neumann–Morgenstern preferences given by u and then over the posterior distribution with preferences given by $\phi(u)$.

In contrast to the MEU model of Gilboa and Schmeidler (1989), where the objective function is not differentiable, here it is differentiable. For this reason, when ϕ is smooth, the Klibanoff et al. (2005) model is sometimes known as the *smooth ambiguity aversion model*. Insofar as expected utility is applied at each of the two stages, within each stage, the model is susceptible to the criticisms that have already been made of EU. Furthermore, the burden of explaining risk and ambiguity attitudes in this model lies on the utility function; in contrast the evidence indicates that the shape of the probability weighting function plays an important role in risk attitudes. Finally, the theory is not easy to test because second order acts are difficult to observe. Wakker (2010) gives a more detailed discussion of these criticisms; see also Baillon et al. (2012) for a further critical evaluation.

Related models can be found in Nau (2006) and Ergin and Gul (2009). The idea in Nau (2006) is that the state space can be partitioned in such a manner that each partition has a similar degree of uncertainty. In a similar framework, Ergin and Gul (2009) obtain a representation that is similar to (4.19). Epstein (2009) has criticized the smooth ambiguity aversion model of Klibanoff et al. (2005), raising concerns about its very foundations and interpretations. He also contests the claim that the model provides a separation between attitudes to ambiguity and to what one would reasonably agree to be an ambiguous situation.¹⁴

4.3.5 A brief note on the related literature

Bewley (2002) allowed for the possibility that one might not have complete preferences over acts. Thus, in the Ellsberg paradox, some of the acts may not be comparable. In this model, let Λ be the set of acts over which the preferences are complete. Then, using the same terminology that was used in (4.7),

$$f \succeq g \Leftrightarrow \int_{\Lambda} u(f(s)) d\rho(s) \geq \int_{\Lambda} u(g(s)) d\rho(s) \text{ for all } \rho \in \Lambda,$$

otherwise (for $\rho \notin \Lambda$), the acts f, g are simply non-comparable. In any non-comparable situation, it is assumed that the status quo is chosen. This framework can also be used to explain the Ellsberg paradox.

A range of other proposals can explain the Ellsberg paradox. We briefly note some of these. Siniscalchi (2009) proposes a model of *vector expected utility*. In this model, for any act, the individual computes his expected utility using a baseline probability distribution. However, the individual also makes an adjustment to utility which depends on a comparison of the act with similar

¹⁴ The criticisms also apply to the related representation by Ergin and Gul (2009). For a reply, see Klibanoff et al. (2009) who argue that Epstein's thought experiments either strengthen their smooth ambiguity aversion model, or point shortcomings that apply to a much wider range of ambiguity models.

acts, using a measure of statistical similarity. Maccheroni et al. (2006b) propose a representation that they call the *variational representation of preferences*. This is based on a weakening of C-independence (Definition 4.68). Their main axiom, *weak certainty independence*, requires that preferences over acts be invariant to changes in location, but not scale of utilities. Cerreia-Vioglio et al. (2011) propose a general model that provides a unified account of MEU and smooth ambiguity aversion models. Gul and Pesendorfer (2013) propose an axiomatization of their model, *expected uncertain utility*. Chew and Sagi (2006, 2008) propose a theory of “small worlds.” The idea here is to divide the domain of possible events into groups of events. The decision maker has a good idea of probabilities within a group. However, he/she reveals ambiguity when asked to compare across the groups.

We have already considered some definitions of ambiguity aversion. In Definition 4.69, we considered Schmeidler’s definition. In Section 4.3.3 we considered the definition of ambiguity aversion given by Klibanoff et al. (2005). In the α -MEU model of Ghirardato et al. (2004) considered above, the magnitude of α is used to capture the degree of ambiguity (although we have also noted the problems with this approach). These definitions depend on the shape of the utility function. However, it is also possible to give definitions of ambiguity aversion that are independent of the shape of the utility function.¹⁵

Yet another definition of ambiguity aversion is due to Chateauneuf and Tallon (2002). Suppose that there are n acts $f_1, f_2, \dots, f_n$ and the decision maker is indifferent between them, i.e., $f_1 \sim f_2 \sim \dots \sim f_n$. Suppose also that there exist n real numbers $\alpha_1, \alpha_2, \dots, \alpha_n$ with the property that $\sum_{i=1}^n \alpha_i = 1$ and $\sum_{i=1}^n \alpha_i f_i = x \in X$, the set of outcomes. Notice that x could be considered to be a hedged bet relative to any of the acts, f_i . Then, ambiguity aversion requires that $x \succ f_i$, $i = 1, 2, \dots, n$. This definition is similar in idea to Schmeidler’s definition. It too relies on a taste for diversification; the decision maker prefers a hedged bet to an unhedged bet.

Some other definitions of ambiguity aversion include those by Epstein (1999) and Ghirardato and Marinacci (2002). An issue with these definitions is that they do not fully separate attitudes to ambiguity from attitudes to risk. For one such separation of ambiguity and risk preferences, see Gajdos et al. (2008).¹⁶

4.3.6 Some empirical evidence on ambiguity aversion

The experimental results in Section 4.4.1 question some of the underlying premises of the modern ambiguity aversion literature; for a survey, see Trautmann et al. (2016). For instance, ambiguity disappears in Ellsberg style experiments when individuals bet on one urn at a time rather than across both urns. Knowledge/expertise in the sorts of questions asked could even create ambiguity-seeking behavior that is problematic for the modern theoretical models of ambiguity aversion. Furthermore, in the presence of losses or lower probabilities with which events occur, we may observe ambiguity neutrality or even ambiguity-seeking behavior.

Hey et al. (2010) use an experimental setup that modifies the one in Halevy (2007). The authors argue that the original Ellsberg paradox is subject to one important criticism. Namely, that the urn with the unknown composition of balls invites *suspicion*. Subjects in experiments could be suspicious that the experimenter wishes to reduce their total monetary payoff and could somehow

¹⁵ See, for instance, Epstein and Zhang (2001), Nehring (2001), Ghirardato et al. (2004), and Ghirardato and Siniscalchi (2012). Tversky and Wakker (1995) define ambiguity aversion as source preference for risk, which does not involve any utility.

¹⁶ Klibanoff et al. (2005) also claim a separation between attitudes to risk and to ambiguity. This claim is contested in Epstein (2009).

have rigged the urns in his favor. Hence, the “ambiguous urn” could turn into a “suspicious urn”. Subjects might even use the heuristic that they are playing a game against a *malevolent nature*. The very fact that someone offers you a bet, could make you suspicious. For instance, you might wonder if the other party has access to some information that you do not have.

The solution proposed by Hey et al. (2010) for the suspicious urn problem is to use a bingo-blower. This is a rectangular device with transparent glass walls that has a collection of balls being spun around in a continuous motion. In order to draw a ball from the box, the box is tilted and a ball is ejected randomly, and transparently, through a transparent tube. In this case, the balls can be seen, but their proportions cannot be calculated on account of the continuous motion of the balls. The authors argue that this enables them to term the urn as “ambiguous” rather than “suspicious.”

Hey et al. (2010) examine 12 alternative specifications. These include expected utility, the decision field theory (DFT) of Busemeyer and Townsend (1993), MEU and α -MEU models, the Hurwicz criterion, Choquet expected utility, a form of prospect theory that is not PT,¹⁷ and minimum regret. Theories with second order beliefs (see, for instance, Sections 4.3.3 and 4.3.4) are not included because the experiments do not have the multiple stages required to test them. DFT and OPT emerge as the two best theories out of the 12. There is support for Gilboa and Schmeidler’s (1989) MEU model and Ghirardato et al.’s (2004) α -MEU models.

Ahn et al. (2014) find greater support for multiple priors models and rank dependent models relative to smooth ambiguity aversion models. Hayashi and Wada (2010) reject multiple priors models on the grounds that priors that lead to intermediate values of expected utility also play a critical role. By contrast, from (4.8), only those priors that lead to maximum and minimum expected utility should matter.

Conte and Hey (2013) employ a new experimental setup that allows for the testing of theories that employ second order beliefs. They use a mixture model that considers five different theories. They find that the choices of 22% of the subjects are consistent with EU, 53% are consistent with the smooth ambiguity aversion model, 22% are consistent with the rank dependent model and 3% with the α -MEU model. There are vexed econometric issues in interpreting the results from any of these studies; see Hey (2014) for a discussion.

Recall that the Ellsberg paradox poses problems for Savage’s SEU, because it rules out well-defined subjective probabilities in the sense of SEU. Machina (2009) proposes two-urn problems that give rise to a similar problem for models based on Choquet expected utility and axiomatized, for instance, by Gilboa (1987) and Schmeidler (1989).¹⁸ Baillon et al. (2011) show in fact that Machina’s results falsify four main models: maximin expected utility, variational preferences, α -maximin and the smooth model of ambiguity aversion.¹⁹

Gilboa et al. (2009) illustrate a shortcoming of the Bayesian approach by asking the following hypothetical question. Suppose that you were to hear the title of a seminar in a University: “Are all cydophines, also abordites?” What probability do you assign to the event? As the authors say: “For all you know these could be enzymes, grammatical structures in an ancient language, or Abelian groups.” So you cannot justifiably assign any probability. Using the principle of insufficient reason

¹⁷ It is an extension of OPT but Kothiyal et al. (2014) express reservations about using this variant in the domain of lotteries that they consider.

¹⁸ Recall that models based on Choquet expected utility violate *event separability* because the Choquet capacities involve other events on account of the cumulative treatment of events. A particular kind of separability, *tail separability*, still remains in such models. Machina (2009) exploits the notion of tail separability for his results.

¹⁹ The vector expected utility model of Siniscalchi (2009) is, however, not subjected to the anomalies pointed out by Machina (2009).

is not fully satisfactory either. The authors conclude: “The Bayesian approach is lacking because it is not rich enough to describe one’s degree of confidence in one’s assessments. For any probability question it requires a single probability number as an answer, excluding the possibility of replies such as ‘I don’t know’ or ‘I’m not so sure.’” In a similar vein, Gilboa et al. (2009) ask readers, what is the probability that the US President, six years from now, will be a Democrat?

These arguments rule out other possible solutions, particularly expert advice and the option value of waiting for uncertainty to resolve. For instance, suppose that instead a non-economics scholar at a university were to see a seminar titled “Are Hicksian cross partial derivatives negative?” The scholar could ask a friend in the economics department for advice. In many situations of interest in real life, if it were terribly important for you to know the relevant probabilities, you would often ask someone who knows better. For instance, you would ask an expert for financial/insurance information. Or perhaps you might rely on some institutional solution; for instance you need not figure out the probability that the saturated fat content of a jar of powdered milk drink is high or low, you simply read information about the contents, which must be provided by law in most countries.

But in some situations, nobody might currently know the relevant probabilities, not even experts. The modern ambiguity aversion literature requires decision makers to form an immediate probability assessment of an event. In many relevant situations of interest to economics, there is substantial *option value of waiting*. The reason is that ambiguity can resolve as time passes. If you are not sure that you have met the right person for marriage, or you are not sure if this is the right time for having a child, or divorcing your partner, you can rationally wait for more information to unravel. Similarly, if you are not sure that this is the right time to make a particular investment, you can wait to see how events unfold. Countries might put off having a common currency, e.g., the euro, or defer on the death penalty, if they are not sure about the probabilistic merit of the case.

The findings of Halevy (2007), Hey et al. (2010), and Kothiyal et al. (2014) cast sufficient empirical doubt about some of the modern ambiguity aversion theories. The findings of Amos Tversky and colleagues that we outline in Section 4.4.1, question some of the central assumptions in all neoclassical theories of ambiguity.

The empirical literature on ambiguity aversion till about the mid 1990s is considered in Camerer (1995). We can summarize his stylized facts on ambiguity aversion as follows. First, subjects’ aversion to more ambiguous choices does not vanish if one gives them written arguments to illustrate their paradoxical choices. Second, ambiguity and attitudes to risk are not correlated. So, for instance, someone who is more risk averse is not necessarily more ambiguity averse. Third, compound lotteries are considered to be more ambiguous, i.e., subjects are more keen to bet on lotteries with some unknown probability rather than a probability distribution over an unknown probability. Fourth, as one widens the domain of the probability distribution of unknown events (for instance, the range of probabilities p in Example 4.28), choices are considered by subjects to be more ambiguous. Fifth, ambiguity aversion is not caused by a belief on the subjects’ part that the odds are stacked against them in the ambiguous choice.

There is also some evidence that ambiguity aversion might be more relevant in the domain of gains rather than in the domain of losses (Cohen et al., 1987). Similar findings can be found in Du and Budescu (2005) who find decision makers to be nearly ambiguity-neutral in the loss domain. Furthermore, while there appears to be some correlation between risk and ambiguity attitudes in the domain of gains, there is none in the domain of losses (Chakravarty and Roy, 2009). Gender differences in ambiguity aversion have also been documented. Women are relatively more ambiguity averse in situations where the degree of ambiguity is high (Borghans et al., 2009). Using

data from trainee truck drivers, Burks et al. (2008) find that attitudes to risk and ambiguity can be separated in their data. Furthermore, indices of ambiguity aversion are found to be negatively correlated with cognitive skills (measured through, say, non-verbal IQ tests).

Several field experiments illustrate the effect of ambiguity aversion in real-world choice. Investors have a home country bias, investing far more in home country securities relative to diversifying by using international securities (French and Poterba, 1991). Referees are more willing to accept papers when refereeing is single-blind rather than double-blind (Blank, 1991).²⁰ Insurance contracts can charge an ambiguity premium (Hogarth and Kunreuther, 1992). Furthermore, even professionals such as actuaries and insurance underwriters exhibit ambiguity aversion (Hogarth and Kunreuther, 1989; Kunreuther et al., 1995).

What causes ambiguity aversion? This remains an intriguing question. One possible explanation is that people take account of the negative evaluation of others, of their decision (Curley et al., 1986). Hence, if they are being directly observed by others, they are found to be more ambiguity averse. This is sometimes referred to as the *fear of negative evaluation*, FNE. When the perception is that others are more competent, then FNE (and, hence, ambiguity aversion) might be enhanced even more.²¹ Trautmann et al. (2008) introduce an experimental design which is able to control for FNE. They find that this leads to the disappearance of ambiguity aversion. The authors conjecture that FNE might be a necessary condition for the presence of ambiguity aversion.

4.4 Behavioral models of ambiguity

In this section, we consider behavioral models of ambiguity. Under pure risk, this class of models does not reduce to EU. We consider two classes of models. The common theme running through these models is that the individual's attitudes towards ambiguity are influenced by the *source of information*.

4.4.1 *Support theory*

In a series of papers, Amos Tversky and several collaborators examine the nature of ambiguity within and outside the Ellsberg context. We now examine this line of work which is influential in psychology, and its decision theoretic foundations have recently been articulated by Ahn and Ergin (2010).²²

Fox and Tversky (1998) begin with the observation that subjective expected utility (SEU) determines beliefs and tastes in one fell swoop. SEU is agnostic on what the decision maker forms first: beliefs or preferences. However, beliefs may precede preferences. For instance, a punter might bet on team A rather than on team B because he believes the former is more likely to win. Punters do not, however, infer this belief from a observation that they prefer betting on team A.

In the first step, individuals form beliefs about events using the *support theory* of Tversky and Koehler (1994). The starting point of support theory is the empirical observation that *procedure*

²⁰ These experiments used data on referees from the *American Economic Review*. In single blind refereeing, the identity of the author is known to the referee (but not vice versa). In double blind refereeing, the identity of each is private knowledge (so this is a more ambiguous situation).

²¹ See for instance, Heath and Tversky (1991), Fox and Tversky (1998), and Fox and Weber (2002).

²² These authors derive a partition dependent expected utility representation of preferences. The same sample space can be finely or coarsely partitioned. The decision maker then has non-additive weights over events in any partition. Also see Cohen and Jaffray (1980) for an earlier attempt along this direction.

invariance is violated when one elicits beliefs using different methods. Suppose that $E_1, E_2, \dots, E_n$ are disjoint events, then we typically observe the following form of subadditivity (Fischhoff et al., 1978; Teigen, 1974).

$$P(E_1 \cup E_2 \cup \dots \cup E_n) < P(E_1) + P(E_2) + \dots + P(E_n).$$

Motivated by this finding, support theory assumes that individuals do not attach probability to events, but rather to the “description of events” or “hypotheses.” So depending on how the same event is described or framed, the individual could assign different probabilities to it.

If only two possible hypotheses, A and B , hold, then the subjective probability that A rather than B holds, denoted by $P(A, B)$, is given by:

$$P(A, B) = \frac{s(A)}{s(A) + s(B)} \in [0, 1], \quad (4.20)$$

where, for any hypothesis, A , $s(A) \geq 0$ is the “support value” of the hypothesis. This reflects the strength of evidence for this hypothesis. It is obvious from (4.20) that $P(A, B) + P(B, A) = 1$; this is known as *binary complementarity*. For $A = A_1 \cup A_2$, the support value satisfies:

$$s(A) \leq s(A_1 \text{ or } A_2) \leq s(A_1) + s(A_2). \quad (4.21)$$

The first inequality in (4.21) implies that by unpacking an event into its constituent parts increases the support for the event. The second inequality implies that this, in turn, is dominated by the sum of support for the individual constituents. The idea behind this kind of *subadditivity* is that by breaking an event into its sub-parts, one increases the salience of the event; quite possibly, one might be able to recall a greater number of instances of the event. For instance, it is likely that the support value for homicide (event A) is lower than the support value for the two constituent events, homicide by strangers (A_1) and homicide by acquaintances (A_2), simply because the latter two might bring to mind more instances of such events and thereby increase their salience. This intuition is borne out by evidence, which includes field studies of the behavior of experienced physicians.²³

Once beliefs consistent with (4.20), (4.21) are allowed for, one transforms an uncertain situation (what is the probability of events A and B that lack an objective probability?) into a risky one (well-defined probabilities $P(A, B)$, $P(B, A)$ can be formed). Such probabilities are said to be consistent with support theory.

For n possible events $E_1, E_2, \dots, E_n$ that represent a partition of the sample space S , denote the associated subjective probabilities that are consistent with support theory by $p_1, p_2, \dots, p_n$. One can then move on to the second stage. In this stage, the decision maker simply uses prospect theory to evaluate what has become a risky prospect.

Tversky and Wakker (1995) extend the weighting function under risk, $w(p)$, to uncertainty, $W(E)$, as follows.

1. (Weight on extreme events) $W(\phi) = 0$, $W(S) = 1$.
2. (Bounded subadditivity) Denote by E_i, E_j , two events in the sample space such that $i, j = 1, 2, \dots, n$, $i \neq j$ and $E_i \cap E_j = \phi$, then:²⁴

²³ For the evidence, see, for instance, Tversky and Koehler (1994) and Rottenstreich and Tversky (1997). For the specific case of experienced physicians exhibiting the ideas in (4.21), see Redelmeier et al. (1995).

²⁴ Support for the bounded subadditivity hypothesis can be found in Tversky and Fox (1995).

- 2a. (Possibility effect or lower subadditivity) $W(E_i) \geq W(E_i \cup E_j) - W(E_j)$ whenever $W(E_i \cup E_j) \leq W(S - E)$ for some event E . The possibility effect captures the idea that increasing the probability of winning a prize from 0 to p has a larger effect than increasing the probability from $q > 0$ to $q + p \leq 1$. This is seen by considering successively the cases $E_j = \phi$ (where the inequality binds) and $E_j \neq \phi$ (where it need not bind).
- 2b. (Certainty effect or upper subadditivity) $W(S) - W(S - E_i) \geq W(E_i \cup E_j) - W(E_j)$ whenever $W(E_j) \geq W(E')$ for some E' . The certainty effect shows that the impact of an event E_i is greater when it is subtracted from a certain event (S) rather than an uncertain event.

Consider simple prospects of the form (E_i, x) that offer a monetary prize x if the event E_i occurs. Denote the reference point as $\bar{x}$. Then the value of the prospect (E_i, x) is given by

$$V = v(x - \bar{x})W(E_i) = v(x - \bar{x})w(p_i),$$

where p_i is the subjective probability associated with the event E_i and it is consistent with support theory. The function $v(x - \bar{x})$ is as defined in prospect theory.

Heath and Tversky (1991) introduce the *competence hypothesis*. The competence hypothesis requires that an individual's propensity to bet on a lottery depends on how knowledgeable or competent the individual feels in that context. Several cognitive explanations for this hypothesis are suggested. For instance, people learn from experience that they perform better in challenges where they have better knowledge, skill, or competence. Another possibility is that in addition to the monetary consequences from betting, there are also psychic consequences, e.g., credit or blame for an outcome, lowered or improved self esteem, and approval or condemnation from others. In pure games of chance, one does not expect these psychic payoffs to be very important. However, in games involving skill, judgment, or knowledge, they are likely to be far more important. Competent people can take credit when they succeed in games of skill and can sometimes be insulated from blame ("the expert's prediction was incorrect but the expert is very competent, so it must be a very noisy environment, making predictions too difficult"). On the other hand, if a non-expert gets it correct he gets no credit ("succeeded by chance") and if he fails he might get the blame ("he had no expertise; failure was expected"). Hence, we would expect people with the relevant knowledge and competence to be more likely to bet on games that involve knowledge/competence.

In the context of the Ellsberg paradox, people might dislike betting on the "vague" urn for fear of being shown up to be less knowledgeable or competent. Indeed, Curley et al. (1986) show that aversion to ambiguity increases if others can observe one's decision and the actual outcome ("fear of negative evaluation"). The Ellsberg paradox is not, however, a particularly good application of the competence hypothesis because it is not clear what knowledge or competence is involved.

Fox and Tversky (1995) made the important observation that the literature following the Ellsberg paradox typically invoked a *within-subjects design* of experiments. In these experiments, subjects always compared clear and vague bets; for instance, a comparison of two urns, one with a known composition of balls and another with an unknown composition. They propose, instead, a *between-subjects design* in which subjects compare either one or the other urn at a time. Their main experimental finding is that ambiguity aversion no longer exists in a between-subjects design.

Consider, for instance, the following experiment reported in Fox and Tversky (1995). Bag A has 50 red chips and 50 black chips. Bag B has 100 red and black chips, but the exact proportion is unknown. A "ticket" to a bag entitles the holder to play the following game. First, choose a color,

Table 4.2 Results from the Fox and Tversky (1995) study.
Standard errors in parantheses.

	Clear bet	Vague bet
Comparative condition	\$24.34 (2.21)	\$14.58 (1.80)
Non-comparative condition	\$17.94 (2.50)	\$18.82 (2.87)

then draw randomly a chip from the bag. If the chip is of the same color as the chosen color, then the ticket holder wins \$100, otherwise he/she wins nothing. Experimental subjects are asked how much they would pay for a “ticket” to each bag?

Two conditions were identified and the subjects were split roughly evenly between the two conditions. In the *comparative condition* (67 subjects) the two bags were simultaneously presented to the subjects and they had to state a price they were willing to pay for a ticket for each bag. The other half of the subjects (74 subjects) played the *non-comparative condition* and they were split into two groups. One group (35 subjects) had to state the price they were willing to pay for the ticket to Bag A only and the remaining (39 subjects) had to state the price they were willing to pay for Bag B only.

The results are show in Table 4.2 for the case of the *clear bet* (betting on bag A) and the *vague bet* (betting on bag B). The results are in line with the conjecture made by the authors. Ambiguity aversion, as traditionally understood, arises in the comparative condition; subjects are willing to pay significantly more for a ticket to Bag A. However, in the non-comparative condition, such ambiguity aversion completely disappears. This result is unaltered when the authors introduce real monetary stakes.

Subjects do seem aware that relevant information is missing in the case of Bag B. However, unless there is a direct comparison between the clear bet and the vague bet, ambiguity aversion does not arise, so missing information is not a sufficient condition for ambiguity aversion. These results are problematic for neoclassical theories of ambiguity.

4.4.2 Prospect theory, source-dependent uncertainty, and ambiguity

We have briefly alluded to the work of Chew and Sagi (2008) above. They partition the domain of events into groups such that there is homogeneity in uncertainty within each group, yet heterogeneity between the groups. Section 4.4.1 also alluded to the *source of information* in influencing the beliefs of individuals about uncertain events (Tversky and Fox, 1995; Fox and Tversky, 1995).

In an important breakthrough, Abdellaoui, Baillon et al. (2011) formalize some of these ideas and also present evidence for their insights. The central idea is that the decision maker distinguishes between information from difference *sources* of uncertainty. The authors (p. 699) give an informal definition of a *source*: “A source of uncertainty concerns a group of events that is generated by a common mechanism of uncertainty.” Consider the two-colors example of the Ellsberg paradox (Example 4.27) in which there is a known urn (K) and the unknown urn (U). Decision makers are indifferent between the two colors in each urn when the urns are presented separately. Yet, they are not willing to exchange a bet on a color in urn K with a bet on the same color in urn U . Thus, it appears that each urn constitutes a separate source of uncertainty.

We have outlined PT under uncertainty in Section 2.4.1. We now employ similar machinery to develop the ideas here. Let $f = (E_1, x_1; E_2, x_2; \dots; E_n, x_n)$ be a Savage act, as in (1.6), that gives

a prize x_i if event E_i occurs and $\cup_{i=1}^n E_i = E$, where E is the universal event. The decision maker does not know which of the n events will occur. Let $P(E_i)$ be the subjective probability of the event E_i , which can be elicited from the choices made by the subjects. Suppose that the decision maker has complete and transitive preferences $\succeq$ over F , the set of all Savage acts. Under SEU (see Section 1.3) the decision maker evaluates the act f by the functional

$$SEU(f) = \sum_{i=1}^n P(E_i)u(x_i), \quad (4.22)$$

where u is a continuous and strictly increasing utility function.

In the empirical tests undertaken by Abdellaoui, Baillon et al. (2011), $n = 2$. Thus, they have prospects of the form $(E_1, y; E_2, x)$, where $y \leq x$. The reader should have no problems with the next two definitions.

Definition 4.71 (*Weighting function for uncertainty*): A weighting function under uncertainty is a mapping $\tilde{w} : S \rightarrow [0, 1]$ such that:

- (i) $\tilde{w}(\emptyset) = 0$, where $\emptyset$ is the empty set.
- (ii) $\tilde{w}(E) = 1$, where E is the universal event.
- (iii) If $E_i \subseteq E_j$, then $\tilde{w}(E_i) \leq \tilde{w}(E_j)$.

Definition 4.72 (*Binary rank dependent utility*): The utility representation of the prospect $(E_1, y; E_2, x)$ under rank dependent utility is given by

$$RDU(E_1, y; E_2, x) = [1 - \tilde{w}(E_2)] u(y) + \tilde{w}(E_2)u(x). \quad (4.23)$$

We now wish to define the sense in which two events are perceived to be equally likely.

Definition 4.73 (*Exchangeable events*):

- (i) Two disjoint events E_1 and E_2 are exchangeable if permuting their respective prizes has no effect on preferences, i.e.,

$$(E_1, x_1; E_2, x_2; \dots; E_n, x_n) \sim (E_1, x_2; E_2, x_1; \dots; E_n, x_n).$$

- (ii) An entire partition $(E_1, E_2, \dots, E_n)$ of the universal event is exchangeable if all pairs of events in it are exchangeable.

Definition 4.74 (*Source*; Abdellaoui, Baillon et al., 2011, p. 699): A source, S is an algebra. It contains the universal event (certain to happen), the vacuous event (certain not to happen), the complement of each element in the set, and the union of each pair of elements in the set.²⁵

We distinguish between the various sources of uncertainty (such as the single urn and the two urns in the Ellsberg paradox) by giving the following definition of a uniform source.

Definition 4.75 (*Uniform sources*): A source S is said to be uniform if for any partition of events in S , $(E_1, E_2, \dots, E_n)$, probabilistic sophistication (Definition 1.13) holds. In other words, choices among prospects of the form $(E_1, x_1; E_2, x_2; \dots; E_n, x_n)$ can equivalently be expressed as choices over prospects of the form $(x_1, P(E_1); x_2, P(E_2); \dots; x_n, P(E_n))$ where $P(E_i)$ is the subjective probability of the event E_i . The decision maker is indifferent between two event-contingent prospects “from the same source” that lead to the same probability-contingent prospect. This property need not hold if we were to change the source.

²⁵ Thus sources also contain every finite union and intersection of their elements.

If any finite partition of a source is exchangeable (Definition 4.73), then the source is a uniform source.

Example 4.31 (Source Preference; Wakker, 2010, p. 318) Suppose that A and B are two sources and we have

$$A \succeq B \text{ and } A^c \succ B^c \text{ for all } A \in \mathcal{A} \text{ and } B \in \mathcal{B}$$

but it is never the case that $B \succeq A$ and $B^c \succ A^c$, then we say that the decision maker prefers the source A to B .

Definition 4.76 (Source functions): Since probabilistic sophistication holds in any uniform source, S , so for any $E_i \subseteq S$ we have

$$\tilde{w}(E_i) = w_s(P(E_i)),$$

where w_s is known as a source function. In this case we can write (4.23) as

$$RDU(E_1, y; E_2, x) = [1 - w_s(P(E_2))] u(y) + w_s(P(E_2)) u(x).$$

Example 4.32 Consider the prospect $(E_1, x_1; E_2, x_2; \dots; E_n, x_n)$ in which the probabilities of all events are “objective probabilities” and $\cup_{i=1}^n E_i = S$. Then S is a uniform source. In this case we have the standard situation of choice under risk in which the source function w_s is written as a standard probability weighting function w without a subscript.

To check if the behavior of decision makers differs when they receive information from non-uniform sources, Abdellaoui, Baillon et al. (2011) propose the following modification of the Ellsberg experiments. They consider urns K and U as in the Ellsberg paradox. Urn K is known to contain 8 balls of different colors, while urn U is also known to have 8 balls; the color composition in urn U is unknown, but it contains no new colors relative to urn K . So urn U may contain one ball each of the 8 colors, or all 8 balls of the same color, or any combinations of the 8 colors that are present in urn K .

In order to test for uniformity, events $E_2, E'_2, E''_2, \dots$ are specified that have the same number of balls (3 balls in the experiments), for each of the two urns. The decision maker is then asked to state certainty equivalents for the prospects $(E_1, 0; E_2, 25)$, $(E_1, 0; E'_2, 25)$, $(E_1, 0; E''_2, 25)$, $\dots$ from each of the urns. Since each urn is hypothesized to be a uniform source, one would expect the events in each urn to be exchangeable, hence, the certainty equivalents for each prospect in each urn should be identical. However, if the K and U urns are different sources, then a comparison of the certainty equivalents for a pair of prospects across the two urns should reveal a difference.

Utility elicitation was carried out by using the method of Abdellaoui et al. (2008) (see Section 2.5). Certainty equivalents, CE , were elicited for prospects of the form $(E_1, 0; E_2, x)$, where x varied from 0 euros to 25 euros, and the event E_2 specified balls of 4 colors, red, blue, yellow, cyan, $E_2 = \{R_s, B_s, Y_s, C_s\}$; the sources being $s = K, U$. Since the event E_2 specifies 4 out of 8 colors, so we expect the subjective probability of the event to be 0.5 for each urn, i.e., $P(E_2) = 0.5$. A power form of utility was assumed, so $u(x) = \left(\frac{x}{25}\right)^{\rho_s}$, where $\rho_s > 0$ is a parameter. Using the definition of certainty equivalent, this implies that $\left(\frac{CE_s}{25}\right)^{\rho_s} = w_s(0.5) \left(\frac{x}{25}\right)^{\rho_s}$, where CE_s is the certainty equivalent when $s = K, U$. Non-linear least squares is then used to estimate the utility function and the source function at the probability 0.5.

For probabilities different from 0.5, we obtain the source function by varying the number of colors j in the event E_2 , $j = 1, 2, 3, 5, 6, 7, 8$. Once we obtain the 8 values of the source function

$(w_s(\frac{1}{8}), w_s(\frac{2}{8}), \dots, w_s(\frac{8}{8}))$, a two-parameter Prelec function, $w(p) = e^{-\beta(-\ln p)^\alpha}$ (Definition 2.22), is fitted through them.

The results are as follows.

1. Uniformity is not rejected for each of the two urns separately. This result holds at the aggregate and the individual-level data.
2. The estimates of ρ_s , $s = K, U$ are indistinguishable. Hence, potential explanations for the Ellsberg paradox should not rely on the differences in the utility functions across the two sources.
3. The most telling differences between the two urns come from the fitted source functions, w_s , $s = K, U$. The Prelec parameters for urn K are $\alpha = 0.85$, $\beta = 0.93$, and for urn U these are $\alpha = 0.64$, $\beta = 1.03$; Figure 4.1 plots the two source functions. The diagonal line corresponds to EU, the dotted curve to the source function for urn K and the heavy curve for urn U . Both parameter estimates for the source functions are significantly different from 1, hence, EU cannot explain the data (recall that for EU, $\alpha = \beta = 1$). A plausible interpretation is that each urn is viewed as a different source of uncertainty by the decision maker. The difference between the source functions for the two urns may be taken as a measure of the *ambiguity attitudes* of the decision maker. A measure of pessimism, that is developed by the authors, is significantly higher for urn U relative to urn K .
4. There is significant heterogeneity among the decision makers, as shown by the source functions for three subjects in Figure 4.2. Subject 2 is most pessimistic; subject 44 is most likely to be insensitive (flattest middle segment of the source function for urn U); subject 66 is *ambiguity seeking* because the source function for urn K lies outside the source function for urn U . Thus, subjects are characterized by a rich range of behaviors in the domain of uncertainty.

We have noted earlier that PT can be applied to risk, uncertainty, and ambiguity. An important empirical test of PT in this regard is provided by Kothiyal et al. (2014). These authors offer a more satisfactory treatment of PT relative to that reported in the tests of Hey et al. (2010). These tests involve a combination of two major advances: the source method employed in Abdellaoui,




Figure 4.1 Source-dependent probability weighting functions. Plots of the diagonal line (solid thin line) and the source functions for urn K (dashed curve) and urn U (solid line).

Source: Abdellaoui et al. (2008).




Figure 4.2 Subject heterogeneity for three subjects expressed in terms of their source functions.

Source: Abdellaoui, Baillon et al. (2011).

Baillon et al. (2011) and the experimental implementation of the method for testing multiple prior models in Hey et al. (2010).

The device of a three-color (pink, blue, and yellow) bingo blower (see discussion above in Section 4.3.6) is used to ensure that the *ambiguous urn* does not turn into a *suspicious urn*. The total number of balls was varied to be 10, 20, and 40. Three-event prospects of the form $(E_1, 100; E_2, 10; E_3, -10)$ are presented. An example is: $E_1 = \{\text{pink}\}$, $E_2 = \{\text{yellow, blue}\}$ and $E_3 = \emptyset$ (an impossible event that never occurs). In most alternatives to EU, the decision maker maximizes an objective function of the form

$$V = \pi_1 u(100) + \pi_2 u(10) + \pi_3 u(-10),$$

where π_i is the decision weight associated with event E_i , $i = 1, 2, 3$ and u is a utility function. Under SEU, the decision maker evaluates utility as shown in (4.22). *Subjective expected value* (EV) is similar except that $u(x) = x$. Under *Choquet expected utility* (CEU) described in Section 4.3.1, the weighting function, $\hat{w}$, has the properties: (i) $\hat{w}(\emptyset) = 0$, where $\emptyset$ is the empty set; (ii) if $E_i \subseteq E_j$, then $\hat{w}(E_i) \leq \hat{w}(E_j)$; (iii) $\hat{w}(\cup_i E_i) = 1$. The decision weights in this case are given as in cumulative transformations of probabilities under PT, i.e.,

$$\pi_1 = \hat{w}(E_1), \pi_2 = \hat{w}(E_1 \cup E_2) - \hat{w}(E_1), \pi_3 = 1 - \hat{w}(E_1 \cup E_2).$$

The decision weights under PT are similar except that we allow for different weighting functions in the domains of gains and losses; in this example, only event E_3 pushes the decision maker into the domain of losses. Thus, in PT, unlike CEU, the decision weights need not add up to 1. For estimation purposes, the parameters of the models need to be reduced. We leave the reader to pursue the details on this issue from the original paper.²⁶

A tournament is then run among 11 competing models (as in Hey et al. 2010); this list does not include models with second order beliefs. On most yardsticks, PT beats every other theory.

²⁶ For instance, the weighting functions for gains and losses under PT are treated to be identical, and the Prelec weighting function is used. The authors also examine the implications of relaxing some of these assumptions.

These yardsticks include log likelihoods and the number of subjects conforming to the theory. Among the other rank dependent theories, α -MEU with $\alpha = 1$ does best. The authors conclude (p. 16): “This facilitates the study of ambiguity (the difference between uncertainty and risk), where prospect theory, unlike most other ambiguity theories today, need not commit to the descriptively failing expected utility for risk.”

Exercises for Part 1

1. Using the reduction axiom, show that Proposition 1.1 holds for the case of any finite number of outcomes.¹
2. Using Definition 1.5, show that risk neutrality, risk aversion, risk seeking are equivalent, respectively, to a linear, concave, convex utility function.
3. (St. Petersburg paradox) Suppose that a decision maker were to accept the following gamble. A fair coin is tossed repeatedly in independent trials until a head comes up. If a head comes up on the n th toss, then the payoff from this gamble is 2^n . Show that the expected value of this gamble is infinite. People are typically willing to pay only about \$20 for this gamble. How much would a decision maker who followed the expected value criterion (EV) pay for this lottery? Show that a sufficient condition to reconcile the EV criterion with the empirical findings is to have a sufficiently concave utility function, $u(x)$, where x is an outcome.
4. (Affine transformations) Let $F : \mathcal{L} \rightarrow \mathbb{R}$ and $G : \mathcal{L} \rightarrow \mathbb{R}$, where $\mathcal{L}$ is the set of all simple lotteries. We say that G is a positive affine transformation of F if, and only if, there are real numbers, α, β , with $\beta > 0$, such that $G = \alpha + \beta F$. Show that $G : \mathcal{L} \rightarrow \mathbb{R}$ is a positive affine transformation of $F : \mathcal{L} \rightarrow \mathbb{R}$ if, and only if, F is a positive affine transformation of G .
5. Given two functions, $F : \mathcal{L} \rightarrow \mathbb{R}$ and $G : \mathcal{L} \rightarrow \mathbb{R}$, let the respective induced preferences on $\mathcal{L}$, be denoted by $\preceq_F$ and $\preceq_G$. Let the two functions, $F : \mathcal{L} \rightarrow \mathbb{R}$ and $G : \mathcal{L} \rightarrow \mathbb{R}$, satisfy, for all $L_1, L_2, L \in \mathcal{L}$, and all $p \in [0, 1]$,
 - (i) $F(x_1) \leq F(L) \leq F(x_n)$ and $F(x_1) < F(x_n)$,
 - (ii) $G(x_1) \leq G(L) \leq G(x_n)$ and $G(x_1) < G(x_n)$.
 - (iii) $F(L_1, p; L_2, 1-p) = pF(L_1) + (1-p)F(L_2)$,
 - (iv) $G(L_1, p; L_2, 1-p) = pG(L_1) + (1-p)G(L_2)$.
 Show that the induced orders, $\preceq_F$ and $\preceq_G$, are identical if, and only if, G is a positive affine transformation of F .
6. Show that Example 1.3 can be recast as a common consequence violation, as defined in Definition 1.15.
7. Prove Propositions 1.6 and 1.7.
8. Draw a probability triangle to explain the common consequence violation corresponding to Example 1.3 and show that the indifference curves in the triangle must *fan out* for consistency with the evidence.
9. (Probabilistic insurance; Wakker et al., 1997) Suppose that a decision maker has wealth, $W > 0$. With some probability $0 < p < 1$, the decision maker incurs a monetary loss, $L < 0$. In effect, the decision maker faces the lottery $(W - L, p; W, 1 - p)$. The utility function of the risk averse decision maker is given by $u(z)$, which is increasing and concave. Since utility is measured on an ordinal scale, we can, for pedagogical simplicity, choose the following normalization $u(W) = 1, u(W - L) = 0$.
 Suppose that the decision maker is offered a hypothetical insurance contract that charges a premium, $P > 0$, so that he just prefers to buy the insurance contract. The insurance contract offers the following lottery, $(W - P, p; W - P, 1 - p) \equiv (W - P, 1)$. Now, suppose that the decision maker is offered *probabilistic insurance*. Under this form of insurance contract, the decision maker will pay a premium of θP , $0 \leq \theta \leq 1$. Should a loss occur, then (1) with probability $1 - \theta$, the loss will not be covered, and the premium θP will be

¹ I am grateful to Ali al-Nowaihi for contributing several problems in this part.

returned, and (2) with probability θ , the loss will be fully covered, but the decision maker will have to pay an extra premium of $(1 - \theta)P$ (so in effect, a total premium of P in that state). All this is summarized in the following lottery under probabilistic insurance.

$$(W - L, (1 - \theta)p; W - P, \theta p; W - \theta P, 1 - p). \quad (4.24)$$

Show that conditional on just accepting the insurance contract at a premium P , an EU decision maker will prefer probabilistic insurance.

10. Prove Proposition 2.8.
11. Prove Proposition 2.9.
12. Consider the probability weighting function $w(p) = p^\gamma$, where $\gamma > 1$. Is this probability weighting function subproportional? Can it capture the certainty effect?
13. Check that the weighting functions in (2.5), (2.6), and (2.7) satisfy the conditions of a probability weighting function. Also check the conditions under which these functions are subproportional.
14. (al-Nowaihi and Dhami, 2011): Let $w(p)$ be the Prelec function. Suppose $\alpha \neq 1$. Then prove the following:
 - (a) $w''(\tilde{p}) = 0$ for some $\tilde{p} \in (0, 1)$ and, for any such $\tilde{p}$:
 - (i) For $\alpha < 1$: $p < \tilde{p} \Rightarrow w''(p) < 0$, $p > \tilde{p} \Rightarrow w''(p) > 0$.
 - (ii) For $\alpha > 1$: $p < \tilde{p} \Rightarrow w''(p) > 0$, $p > \tilde{p} \Rightarrow w''(p) < 0$
 - (b) The Prelec function has a unique inflexion point, $\tilde{p} \in (0, 1)$, and is characterized by $\alpha\beta(-\ln \tilde{p})^\alpha + \ln \tilde{p} + 1 - \alpha = 0$.
 - (c) $\beta = 1 \Rightarrow \tilde{p} = e^{-1}$.
 - (d) $\frac{\partial \tilde{p}}{\partial \beta} = \frac{\alpha \tilde{p}(-\ln \tilde{p})^{1+\alpha}}{(\alpha-1)(\alpha-\ln \tilde{p})}$.
 - (e) $\frac{\partial [w(\tilde{p}) - \tilde{p}]}{\partial \beta} = \frac{\tilde{p}(-\ln \tilde{p})^{1+\alpha}}{(\alpha-1)(\alpha-\ln \tilde{p})} \left(e^{\frac{1-\alpha}{\alpha}} (\tilde{p})^{\frac{1-\alpha}{\alpha}} - \alpha \right)$.
 - (f) $\tilde{p} \leq w(\tilde{p}) \Leftrightarrow \beta \leq 1$.
15. Prove Propositions 2.10, 2.11 and Corollary 2.1.
16. Prove Proposition 2.12.
17. Prove Proposition 2.13.
18. Consider the lottery $L = (15, 0.8; 0, 0.2)$ that captures the risky returns associated with a financial decision. The utility function of the decision maker satisfies $u(0) = 0$. The individual uses RDU to evaluate lotteries and has the probability weighting function $w(p) = p^2$. Relative to a decision maker who follows expected utility (EU), is this decision maker optimistic or pessimistic?
19. Consider the lottery $L = (9, \frac{1}{3}; 16, \frac{1}{3}; 25, \frac{1}{3})$ that captures the discounted risky returns associated with the purchase of a financial security. The utility function of the decision maker is $u(x) = \sqrt{x}$. [Restrict calculations to four decimal places.]
 - (a) Suppose that the decision maker follows EU. Compute the expected utility of the lottery. Is the decision maker risk averse, risk neutral, or risk loving?
 - (b) Suppose that the decision maker uses RDU to evaluate lotteries and has the probability weighting function $w(p) = \sqrt{p}$. Is this decision maker risk averse, risk neutral, or risk loving?
 - (c) If there is a difference in the risk attitudes found in parts (a) and (b), what accounts for the difference?
20. (RDU can explain the Allais paradox) Suppose that the utility function is $u(x) = x^{0.6} - x_0^{0.6}$, where x is the (final) level of wealth and x_0 is the initial level of wealth. The constant

$-x_0^{0.6}$ is added, so that $u(x_0) = 0$ (as was assumed in Example 1.2). Take the Prelec (1998) probability weighting function with $\alpha = 0.15$, $\beta = 0.5$. Show that RDU is able to predict the “common ratio violation” in Example 1.2.

21. Let a decision maker have initial wealth, x_0 , and utility function, u . Now consider problems 1 and 2 from Kahneman and Tversky (1979) that were given to 72 subjects. All lotteries are in units of Israeli pounds.

Problem 1: (72 subjects) Choose between

$$L_1 = (x_0, 0; x_0 + 2400, 1) \text{ and } L_2 = (x_0, 0.01; x_0 + 2400, 0.66; x_0 + 2500, 0.33).$$

Problem 2: (72 subjects) Choose between

$$L_3 = (x_0, 0.66; x_0 + 2400, 0.34) \text{ and } L_4 = (x_0, 0.67; x_0 + 2500, 0.33).$$

The data shows that 82% of the subjects chose L_1 in problem 1 and 83% of subjects chose L_4 in problem 2. Each of these preferences was significant at the 1% level.

- (a) Is the choice data of the majority of subjects consistent with expected utility?
 (b) Assume that $u(x_0) = 0$, and the weighting function is very steep in the neighborhood of $p = 1$, and flat in the middle, which, for instance, is the case with the Prelec function. Is the choice data of the majority of subjects consistent with rank dependent utility?
22. Prove Proposition 2.14.
23. Construct an example, analogous to Example 2.12 to illustrate Remark 2.6.
24. (Huber et al., 1982; Ok et al., 2015) Consider three alternatives, x, y, z . Each alternative has two dimensions (say, price and quality). Alternative x dominates y in one dimension and is dominated in the other. Alternative z (i) is dominated by x in both dimensions, and (ii) dominates y in one dimension, while being dominated in the other dimension. When given a choice from the set $\{x, y\}$ the decision maker chooses y and when given a choice from the set $\{x, y, z\}$ he chooses x . This is known as the *asymmetric dominance effect* or the *attraction effect*.
 (a) Within the neoclassical framework, does such an individual have a well-defined utility function?
 (b) How can you explain this effect using prospect theory?
25. Suppose that a decision maker who uses prospect theory faces the following lottery with two outcomes 0, x such that $x < 0$.

$$L = (0, 1 - p; x, p) \text{ where } x < 0.$$

Suppose that the utility function for losses is given by $u(x) = -2.5\sqrt{-x}$ and we treat the reference point r as zero wealth (so $u(0) = 0$). Let us use the Prelec weighting function, $w(p) = e^{-\beta(-\ln p)^\alpha}$ with $\alpha = 0.5$, $\beta = 1$. Show that depending on the value of p , such a decision maker can be both risk loving and risk averse. In particular, show that there is a cutoff point $p^* = 0.0183$ such that $w(p^*)^2 = p^*$, and (i) for $p^* < 0.0183$, the decision maker is risk averse, while (ii) for $p^* > 0.0183$, the decision maker is risk loving.

26. (Loomes et al., 1992) This exercise highlights various features of regret theory.
 (a) (Violation of transitivity under regret theory) Suppose that we have three states of the world, s_1, s_2 , and s_3 ; each occurs with probability $\frac{1}{3}$. There are three actions, A_1, A_2 , and A_3 ;

outcomes corresponding to these actions in the various states ($x_1 < x_2 < x_3$) are given in the following table.

	$s_1(1/3)$	$s_2(1/3)$	$s_3(1/3)$
A_1	x_1	x_2	x_3
A_2	x_3	x_1	x_2
A_3	x_2	x_3	x_1

Actions A_1, A_2, A_3 are said to be *stochastically equivalent* because the three states are equiprobable. Suppose that $A_1 \prec A_2, A_2 \prec A_3$ and $A_3 \prec A_1$, which violates transitivity. Can regret theory account for this pattern of intransitivity?

(b) (Violation of monotonicity): Suppose that we have only two actions, A_1, A_2 . The outcomes are shown in the following table:

	$s_1(\frac{1}{3})$	$s_2(\frac{1}{3})$	$s_3(\frac{1}{3})$
A_1	x_1	$x_2 + \epsilon$	x_3
A_2	x_3	x_1	x_2

where $\epsilon > 0$ is a sufficiently small number, so that $x_2 + \epsilon < x_3$ and

$$\frac{1}{3}R(x_3, x_1) + \frac{1}{3}R(x_1, x_2 + \epsilon) + \frac{1}{3}R(x_2, x_3) > 0. \quad (4.25)$$

Show that $A_1 \prec A_2$ but A_1 first order stochastically dominates A_2 .

27. (al-Nowaihi and Dhami, 2010a,b) Suppose that a decision maker can suffer the loss, $L > 0$, with probability $p \in (0, 1)$. He/She can buy coverage, $C \in [0, L]$, at the cost rC , where $r \in (0, 1)$ is the *premium rate* (which is *actuarially fair* if $r = p$). Hence, with probability, $1 - p$, the decision maker's wealth is $W - rC$, and with probability p , her wealth is $W - rC - L + C \leq W - rC$. Suppose that the reference point of the individual is the status quo, i.e., $y_0 = W$, the utility function is the power form in Equation (2.20) with $\gamma^+ \simeq \gamma^- = \gamma$, and the objective function is given in Definition 2.34. Assume that we have actuarially fair premium, $r = p$.
 - (a) Show that the optimal purchase is either zero or full coverage, i.e., $C^* = 0$ or $C^* = L$.
 - (b) What is the predicted insurance purchase under PT?
 - (c) What is the predicted insurance purchase under composite prospect theory (CPT)?
 - (d) Two kinds of insurance behavior are observed for low probability losses. Kunreuther et al. (1978) report little or no take-up of insurance for many individuals. However, Sydnor (2010) reports over-insurance for modest risks for many individuals. Can CPT account for these conflicting findings?
28. (al-Nowaihi and Dhami, 2010a) Consider the same setup as in Problem 27:
 - (a) Suppose that a decision maker faces a loss, L , with probability $p = 0.001$. Recall the formal definition of CPT in Section 2.11.3. Motivated by the insurance data from Kunreuther et al. (1978), suppose that the CPF is given by

$$w(p) = \begin{cases} e^{-0.61266(-\ln p)^2}, & \text{i.e., } \alpha = 2, \beta = 0.61266, & \text{if } 0 \leq p < 0.25, \\ e^{-(-\ln p)^{\frac{1}{2}}}, & \text{i.e., } \alpha = 0.5, \beta = 1, & \text{if } 0.25 \leq p \leq 0.75, \\ e^{-6.4808(-\ln p)^2}, & \text{i.e., } \alpha = 2, \beta = 6.4808, & \text{if } 0.75 < p \leq 1. \end{cases} \quad (4.26)$$

In this case (i.e., $p = 0.001$) the findings of Kunreuther et al. (1978) show that only 20% of the decision makers insure. Suppose that the utility function of the decision maker is given in (2.20) with $\gamma^+ \simeq \gamma^- = \gamma = 0.88$.

(a1) For decision makers who use PT, the probability weighting function is the standard Prelec function with $\beta = 1$ and $\alpha = 0.50$, i.e., $w(p) = e^{-(\ln p)^{0.50}}$. Show that such a decision maker will fully insure.

(a2) Consider decision makers who use CPT. Decision makers who belong to the fraction $1 - \mu$ use PT and so, as in (a), will fully insure. Now consider those who belong to the fraction μ . They use the composite Prelec function given in (4.26). Show that these individuals will not insure.

(b) Let the probability of the loss be $p = 0.25$ (instead of $p = 0.001$). Kunreuther's (1978) data shows that 80% of subjects in experiments took up insurance in this case. Show that the predictions of both PT and CPT are in close (but not perfect) conformity with the evidence.

(c) What do we learn from comparing the (correct) answers to parts (a) and (b)?

29. (Dhami and al-Nowaihi, 2013) Suppose that income from legal activity is $y_0 \geq 0$ and income from an illegal activity is $y_1 \geq y_0$. Hence, the benefit, b , from the illegal activity is $b = y_1 - y_0 \geq 0$. If engaged in the illegal activity, the individual is caught with some probability p , $0 \leq p \leq 1$. If caught, the individual is asked to pay a fine F . Assume that we have the *hyperbolic punishment function* $F = \varphi(p) = b/p$.²

(a) Consider a decision maker with continuously differentiable and strictly increasing utility of income, u . Show that under EU, if the decision maker is risk neutral, or risk averse, then the hyperbolic punishment function $\varphi(p) = \frac{b}{p}$ will deter crime. It follows that given any probability of detection and conviction, $p > 0$, no matter how small, crime can be deterred by a sufficiently large punishment. This is the *Becker proposition* alluded to in the text (see Proposition 22).

(b) Now suppose that the decision maker uses PT or CPT and that the reference incomes from the two activities, crime and no-crime, are, respectively, y_R and y_r . If a decision maker commits a crime, and is caught, then assume that the outcome ($y_1 - F$) is in the domain of losses (i.e., $y_1 - F - y_R < 0$). If he commits a crime and is not caught, then assume that the outcome, y_1 , is in the domain of gains (i.e., $y_1 - y_R > 0$). The recent literature has sometimes made the case for the reference point to be the rational expectation of income and also the state-dependent reference point.³ In the spirit of these suggestions, suppose that the expected income from each activity is the reference point. Thus, $y_r = y_0$ and $y_R = y_1 - p\varphi(p)$. Assume the power form of utility in (2.20) with $\gamma^+ = \gamma^- = \gamma$.

(b-1) Show that as the probability of detection approaches zero, a decision maker using PT does not engage in the criminal activity.

(b-2) Show that as the probability of detection approaches zero, under CPT, the fraction μ of individuals is not deterred from the criminal activity (which is consistent with the empirical evidence reviewed in the text). Why does CPT give a better account of criminal behavior?

² al-Nowaihi and Dhami (2012) show that the hyperbolic punishment function is optimal for a wide class of cost of deterrence and damage from crime functions. Furthermore, it provides an upper bound on punishments for a large and sensible class of cost and damage functions. Thus, if the hyperbolic punishment function cannot deter crime, then neither can the optimal punishment function, if different from the hyperbolic.

³ See, for instance, Koszegi and Rabin (2006, 2009). The act dependence of the reference point that we use is also consistent with the third generation prospect theory of Schmidt et al. (2008).

30. (al-Nowaihi and Dhami, 2010a) Rieger and Wang (2006) show that the St. Petersburg paradox re-emerges under PT. How can the St. Petersburg paradox be resolved under CPT?
31. Prove Proposition 3.33.
32. (Loss aversion as an explanation of the St. Petersburg paradox; Camerer, 2005) Consider the gamble set out in the St. Petersburg paradox in Exercise 3. Bernoulli explained this paradox by using a concave utility function with EU.⁴ Camerer (2005) showed that the paradox could also be explained by using loss aversion and OPT/PT. Suppose that a casino, which offers the gamble has a budget of $\$b$, which is the maximum possible payout. Let m be the number of trials such that if a head first came out on the m th trial, the maximum possible payout would be just about exhausted. Suppose that $b = \$1\text{billion}$ (check this implies $m = 30$). Thus, if a head came up on any trial beyond the 30th, then the casino will not be able to honor the terms of the gamble. Suppose that the gamble is purchased for a price equal to P and the decision maker has the piecewise linear utility function,

$$u(x) = \begin{cases} (2^j - P) & \text{if } 2^j \geq P \\ \lambda(2^j - P) & \text{if } 2^j < P \end{cases} \quad (4.27)$$

For some n , let $2^n < P < 2^{n+1}$ and use the identity weighting function $w(p) = p$.

Show that the imposition of a concave utility function is not a necessary condition for resolving the St. Petersburg paradox. Loss aversion also provides a sufficient condition for the resolution of the paradox.

33. (Birnbbaum, 2008) Consider the Birnbbaum critique of PT in Section 2.9.1. Suppose that decision makers obeyed the three principles—transitivity, coalescing, restricted branch independence (RBI). Then show that there can be no Allais paradox.
34. Prove Proposition 2.16(b).
35. Prove Proposition 3.35.
36. Prove Proposition 3.37.
37. Consider the Bayesian approach to uncertainty (SEU). In SEU, beliefs are required to be updated using Bayes' rule. However, Bayesian analysis is agnostic as to the source of beliefs. Evaluate the implications for SEU of the following two scenarios.
- (a) Consider two alternative ways of forming subjective beliefs, $\mu(E_i)$. In one case, the decision maker uses an actual accurate historical record to construct $\mu(E_i)$. In another, the decision maker is ignorant of any data and merely uses a hunch to construct $\mu(E_i)$. For instance, the decision maker can use the *principle of insufficient reason* and simply assign each of the n outcomes a subjective probability equal to $1/n$.⁵ Schmeidler (1989) argued that the two probabilities feel very different.
- (b) Beliefs can sometimes be updated without any apparently new information being available as, for instance, in the *fact-free learning theory* of Aragonés et al. (2005).

⁴ Bernoulli used the logarithmic function.

⁵ The principle of insufficient reason may be more reasonable to apply in problems where the alternatives appear to be relatively more symmetric.

REFERENCES FOR PART 1

- Abdellaoui, M. A. (2000). Parameter-free elicitation of utility and probability weighting functions. *Management Science* 46(11): 1497–512.
- Abdellaoui, M. A. (2002). A genuine rank-dependent generalization of the Von Neumann-Morgenstern expected utility theorem. *Econometrica* 70(2): 717–736.
- Abdellaoui, M. A., Baillon, A., Placido, L., and Wakker, P. P. (2011). The rich domain of uncertainty: source functions and their experimental implementation. *American Economic Review* 101(2): 695–723.
- Abdellaoui, M. A. and Bleichrodt, H. (2007). Eliciting Gul's theory of disappointment aversion by the tradeoff method. *Journal of Economic Psychology* 28(6): 631–45.
- Abdellaoui, M. A., Bleichrodt, H., and Karmoun, H. (2013). Do financial professionals behave according to prospect theory? An experimental study. *Theory and Decision* 74(3): 411–29.
- Abdellaoui, M. A., Bleichrodt, H., and L'Haridon, O. (2008). A tractable method to measure utility and loss aversion in prospect theory. *Journal of Risk and Uncertainty* 36(3): 245–66.
- Abdellaoui, M. A., Bleichrodt, H., and Paraschiv, C. (2007). Loss aversion under prospect theory: a parameter-free measurement. *Management Science* 53(10): 1659–74.
- Abdellaoui, M. A., Diecidue, E., and Öncüler, A. (2011). Risk preferences at different time periods: an experimental investigation. *Management Science* 57(5): 975–87.
- Abdellaoui, M. A., Klibanoff, P., and Placido, L. (2014). Experiments on compound risk in relation to simple risk and to ambiguity. Forthcoming in *Management Science*.
- Abdellaoui, M. A., Vossman, F., and Weber, M. (2005). Choice-based elicitation and decomposition of decision weights for gains and losses under uncertainty. *Management Science* 51(9): 1384–99.
- Abdellaoui, M. A. and Wakker, P. P. (2005). The likelihood method for decision under uncertainty. *Theory and Decision* 58(1): 3–76.
- Abeler, J., Altmann, S., Kube, S., and Wibrall, M. (2010). Gift exchange and workers' fairness concerns: when equality is unfair. *Journal of the European Economic Association* 8(6): 1299–324.
- Abeler, J., Falk, A., Goette, L. and Huffman, D. (2011). Reference points and effort provision. *American Economic Review* 101(2): 470–92.
- Abellan-Perpiñan, J. M., Bleichrodt, H., and Pinto-Prades, J. L. (2009). The predictive validity of prospect theory versus expected utility in health utility measurement. *Journal of Health Economics* 28(6): 1039–47.
- Ahlbrecht, M. and Weber, M. (1997a). An empirical study on intertemporal decision making under risk. *Management Science* 43(6): 813–26.
- Ahlbrecht, M. and Weber, M. (1997b). Preference for gradual resolution of uncertainty. *Theory and Decision* 43: 167–84.
- Ahn, D., Choi, S., Gale, D., and Kariv, S. (2014). Estimating ambiguity aversion in a portfolio choice experiment. *Quantitative Economics* 5(2): 195–223.
- Ahn, D. and Ergin, H. (2010). Framing contingencies. *Econometrica* 78(2): 655–95.
- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'Ecole Americaine. *Econometrica* 21(4): 503–46.
- Allais, M., and Hagen, O. (eds.) (1979). *Expected Utility Hypotheses and the Allais Paradox*. Dordrecht: D. Reidel.
- Allingham, M. G. and Sandmo, A. (1972). Income tax evasion: a theoretical analysis. *Journal of Public Economics* 1(3–4): 323–38.
- al-Nowaihi, A., Bradley, I., and Dhami, S. (2008). A note on the utility function under prospect theory. *Economics Letters* 99(2): 337–9.
- al-Nowaihi, A. and Dhami, S. (2006). A simple derivation of Prelec's probability weighting function. *Journal of Mathematical Psychology* 50(6): 521–4.
- al-Nowaihi, A. and Dhami, S. (2010a). Composite prospect theory: a proposal to combine prospect theory and cumulative prospect theory. University of Leicester. Discussion Paper 10/11.
- al-Nowaihi, A. and Dhami, S. (2010b). The behavioral economics of the demand for insurance. University of Leicester. Discussion Paper.
- al-Nowaihi, A. and Dhami, S. (2011). Probability weighting functions. In *Wiley Encyclopaedia of*

- Operations Research and Management Science*. New Jersey: John Wiley and Sons.
- al-Nowaihi, A. and Dhami, S. (2012). Hyperbolic punishment functions. *Review of Law and Economics* 8(3): 759–87.
- Andersen, S., Cox, J. C., Harrison, G. W., et al. (2012). Asset integration and attitudes to risk: theory and evidence. Georgia State University, Experimental Economics Center. Working Paper 2012–12.
- Andreoni, J., Erard, B., and Feinstein, J. (1998). Tax compliance. *Journal of Economic Literature* 36(2): 818–60.
- Andreoni, J. and Sprenger, C. (2012a). Estimating time preferences from convex budgets. *American Economic Review* 102(7): 3333–56.
- Andreoni, J. and Sprenger, C. (2012b). Risk preferences are not time preferences. *American Economic Review* 102(7): 3357–76.
- Ang, A., Hodrick, R. J., Xing, Y., and Zhang, X. (2006). The cross-section of volatility and expected returns. *Journal of Finance* 61(1): 259–99.
- Anscombe, F. J., and Aumann, R. J. (1963). A definition of subjective probability. *Annals of Mathematical Statistics* 34(1): 199–205.
- Aragones, E., Gilboa, I., Postlewaite, A., and Schmeidler, D. (2005). Fact-free learning. *American Economic Review* 95(5): 1355–68.
- Ariely, D. and Simonson, I. (2003). Buying, bidding, playing, or competing? Value assessment and decision dynamics in online auctions. *Journal of Consumer Psychology* 13(1–2): 113–123.
- Arrow, K. and Hurwicz, L. (1972). An optimality criterion for decision making under ignorance. In C. Carter and J. Ford (eds.), *Uncertainty and Expectations in Economics*. Oxford: Blackwell, pp. 1–11.
- Babcock, L. and Loewenstein, G. (1997). Explaining bargaining impasse: the role of self-serving biases. *Journal of Economic Perspectives* 11(1): 109–26.
- Baillon, A., Driesen, B., and Wakker, P. P. (2012). Relative concave utility for risk and ambiguity. *Games and Economic Behavior* 75(2): 481–9.
- Baillon, A., L'Haridon, O., and Placido, L. (2011). Ambiguity models and the Machina paradoxes. *American Economic Review* 101(4): 1547–60.
- Bar-Ilan, A. (2000). The response to large and small penalties in a natural experiment. University of Haifa, Department of Economics. Working Paper.
- Bar-Ilan, A. and Sacerdote, B. (2001). The response to fines and probability of detection in a series of experiments. National Bureau of Economic Research. Working Paper 8638.
- Bar-Ilan, A., and Sacerdote, B. (2004). The response of criminals and noncriminals to fines. *Journal of Law and Economics* 47(1): 1–17.
- Barberis, N. C. (2013). Thirty years of prospect theory in economics: a review and assessment. *Journal of Economic Perspectives* 27(1): 173–96.
- Barberis, N. and Huang, M. (2008). Stocks as lotteries: the implications of probability weighting for security prices. *American Economic Review* 98(5): 2066–100.
- Barberis, N., Santos, T., and Huang, M. (2001). Prospect theory and asset prices. *Quarterly Journal of Economics* 116(1): 1–53.
- Barberis, N. and Thaler, R. H. (2003). A survey of behavioral finance. In G. M. Constantinides, M. Harris, and R. M. Stulz (eds.), *Handbook of the Economics of Finance, Volume 1*. 1st edition. Amsterdam: Elsevier, ch. 18, pp. 1053–128.
- Barron, G. and Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making* 16(3): 215–33.
- Barseghyan, L., Molinari, F., O'Donoghue, T., and Teitelbaum, J. C. (2013). The nature of risk preferences: evidence from insurance choices. *American Economic Review* 103(6): 2499–529.
- Bateman, I., Munro, A., Rhodes, B., Starmer, C., and Sugden, R. (1997). A test of the theory of reference-dependent preferences. *Quarterly Journal of Economics* 112(2): 479–505.
- Battalio, R. C., Kagel, J. H., and Jiranyakul, K. (1990). Testing between alternative models of choice under uncertainty: some initial results. *Journal of Risk and Uncertainty* 3(1): 25–50.
- Battalio, R. C., Kagel, J. H., and MacDonald, D. N. (1985). Animals' choices over uncertain outcomes: some initial experimental results. *American Economic Review* 75(4): 597–613.
- Baucells, M. and Heukamp, F. H. (2012). Probability and time trade-off. *Management Science* 58(4): 831–42.
- Beattie, J. and Loomes, G. (1997). The impact of incentives upon risky choice experiments. *Journal of Risk and Uncertainty* 14(2): 155–68.

- Becker, G. S. (1968). Crime and punishment: an economic approach. *Journal of Political Economy* 76(2): 169–217.
- Becker, G. S., DeGroot, M. H., and Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral Science* 9(3): 226–32.
- Bell, D. (1982). Regret in decision making under uncertainty. *Operations Research* 30(5): 961–81.
- Bell, D. (1985). Disappointment in decision making under uncertainty. *Operations Research* 33(1): 1–27.
- Bellemare, C., Krause, M., Kroger, S., and Zhang, C. (2005). Myopic loss aversion: information feedback vs. investment flexibility. *Economics Letters* 87(3): 319–24.
- Benartzi, S., and Thaler, R. H. (1995). Myopic loss-aversion and the equity premium puzzle. *Quarterly Journal of Economics* 110(1): 73–92.
- Bernasconi, M. (1998). Tax evasion and orders of risk aversion. *Journal of Public Economics* 67(1): 123–34.
- Bernasconi, M. and Loomes, G. (1992). Failures of the reduction principle in an Ellsberg-type problem. *Theory and Decision* 32(1): 77–100.
- Berns, G. S., Capra, C. M., Chappelow, J., Moore, S., and Noussair, C. (2008). Nonlinear neurobiological probability weighting functions for aversive outcomes. *Neuroimage* 39(4): 2047–7.
- Bewley, T. F. (2002). Knightian decision theory. Part I. *Decisions in Economics and Finance* 25(2): 79–110.
- Billot, A., Gilboa, I., Samet, D., and Schmeidler, D. (2005). Probabilities as similarity-weighted frequencies. *Econometrica* 73(4): 1125–36.
- Binswanger, H. P. (1981). Attitudes toward risk: theoretical implications of an experiment in rural India. *Economic Journal* 91(364): 867–90.
- Birnbaum, M. H. (1997). Violations of monotonicity in judgment and decision making. In A. A. J. Marley (ed.), *Choice, Decision, and Measurement: Essays in Honor of R. Duncan Luce*. Mahwah: Erlbaum, pp. 73–100.
- Birnbaum, M. H. (1999). Testing critical properties of decision making on the Internet. *Psychological Science* 10(5): 399–407.
- Birnbaum, M. H. (2001). A Web-based program of research on decision making. In U.-D. Reips and M. Bosnjak (eds.), *Dimensions of Internet Science*. Lengerich: Pabst Science Publishers, pp. 23–55.
- Birnbaum, M. H. (2004a). Causes of Allais common consequence paradoxes: an experimental dissection. *Journal of Mathematical Psychology* 48(2): 87–106.
- Birnbaum, M. H. (2004b). Tests of rank-dependent utility and cumulative prospect theory in gambles represented by natural frequencies: effects of format, event framing, and branch splitting. *Organizational Behavior and Human Decision Processes* 95(1): 40–65.
- Birnbaum, M. H. (2005). Three new tests of independence that differentiate models of risky decision making. *Management Science* 51(9): 1346–58.
- Birnbaum, M. H. (2007). Tests of branch splitting and branch-splitting independence in Allais paradoxes with positive and mixed consequences. *Organizational Behavior and Human Decision Processes* 102(2): 154–73.
- Birnbaum, M. H. (2008). New paradoxes of risky decision making. *Psychological Review* 115(2): 463–501.
- Birnbaum, M. H. and Chavez, A. (1997). Tests of theories of decision making: violations of branch independence and distribution independence. *Organizational Behavior and Human Decision Processes* 71(2): 161–94.
- Birnbaum, M. H. and Navarrete, J. B. (1998). Testing descriptive utility theories: violations of stochastic dominance and cumulative independence. *Journal of Risk and Uncertainty* 17(1): 49–79.
- Birnbaum, M. H., Patton, J. N., and Lott, M. K. (1999). Evidence against rank-dependent utility theories: violations of cumulative independence, interval independence, stochastic dominance, and transitivity. *Organizational Behavior and Human Decision Processes* 77(1): 44–83.
- Birnbaum, M. H. and Stegner, S. E. (1979). Source credibility in social judgment: bias, expertise, and the judge's point of view. *Journal of Personality and Social Psychology* 37(1): 48–74.
- Blanchard, O. J. and Fischer, S. (1989). *Lectures on Macroeconomics* Cambridge, MA: MIT Press.
- Blank, R. M. (1991). The effects of double-blind versus single-blind reviewing: experimental evidence from The American Economic Review. *American Economic Review* 81(5): 1041–67.
- Blavatsky, P. R. (2007). Stochastic expected utility theory. *Journal of Risk and Uncertainty* 34(3): 259–86.

- Blavatskyy, P. R. (2010). Reverse common ratio effect. *Journal of Risk and Uncertainty* 40(3): 219–41.
- Blavatskyy, P. R. (2011). Loss aversion. *Economic Theory* 46(1): 127–48.
- Blavatskyy, P. and Pogrebna, G. (2009). Myopic loss aversion revisited. *Economics Letters* 104(1): 43–5.
- Bleichrodt, H., Abellan-Perpiñan, J. M., Pinto-Prades, J. L., and Mendez-Martinez, I. (2007). Resolving inconsistencies in utility measurement under risk: tests of generalizations of expected utility. *Management Science* 53(3): 469–82.
- Bleichrodt, H., Cillo, A., and Diecidue, E. (2010). A quantitative measurement of regret theory. *Management Science* 56(1): 161–75.
- Bleichrodt, H., Filko, M., Kothiyal, A., and Wakker, P. P. (2013). Making case based decision theory directly observable. Mimeo, Erasmus University, Rotterdam.
- Bleichrodt, H. and Pinto, J. L. (2000). A parameter-free elicitation of the probability weighting function in medical decision analysis. *Management Science* 46(11): 1485–96.
- Bleichrodt, H. and Pinto, J. L. (2002). Loss aversion and scale compatibility in two-attribute trade-offs. *Journal of Mathematical Psychology* 46(3): 315–37.
- Bleichrodt, H., Pinto, J. L., and Wakker, P. P. (2001). Making descriptive use of prospect theory to improve the prescriptive use of expected utility. *Management Science* 47(11): 1498–514.
- Bleichrodt, H. and Wakker, P. P. (2014). Regret theory: a bold alternative to the alternatives. *Economic Journal* in press.
- Bombardini, M., and Trebbi, F. (2012). Risk aversion and expected utility theory: an experiment with large and small stakes. *Journal of the European Economic Association* 10(6): 1348–99.
- Bommier, A. (2007). Risk aversion, intertemporal elasticity of substitution and correlation aversion. *Economics Bulletin* 4(29): 1–8.
- Booij, A. S. and van de Kuilen, G. (2009). A parameter-free analysis of the utility of money for the general population under prospect theory. *Journal of Economic Psychology* 30(4): 651–66.
- Bordalo, P., Gennaioli, N., and Shleifer, A. (2012). Salience theory of choice under risk. *Quarterly Journal of Economics* 127(3): 1243–85.
- Bordalo, P., Gennaioli, N., and Shleifer, A. (2013a). Salience and asset prices. *American Economic Review* 103(3): 623–8.
- Bordalo, P., Gennaioli, N., and Shleifer, A. (2013b). Salience and consumer choice. *Journal of Political Economy* 121(5): 803–43.
- Borghans, L., Golsteyn, B. H. H., Heckman, J. J., and Huub, M. (2009). Gender differences in risk aversion and ambiguity aversion. *Journal of the European Economic Association* 7(2–3): 649–58.
- Bosch-Domènech, A. and Silvestre, J. (1999). Does risk aversion or attraction depend on income? An experiment. *Economics Letters* 65(3): 265–73.
- Bowman, D., Minehart, D., and Rabin, M. (1999). Loss aversion in a consumption-savings model. *Journal of Economic Behavior and Organization* 38(2): 155–78.
- Boyce, R. R., Brown, T. C., McClelland, G. H., Peterson, G. L., and Schulze, W. D. (1992). An experimental examination of intrinsic environmental values as a source of the WTA-WTP disparity. *American Economic Review* 82(5): 1366–73.
- Boyer, B., Mitton, T., and Vorkink, K. (2010). Expected idiosyncratic skewness. *Review of Financial Studies* 23(1): 169–202.
- Brookshire, D. S., and Coursey, D. L. (1987). Measuring the value of a public good: an empirical comparison of elicitation procedures. *American Economic Review* 77(4): 554–66.
- Brosnan, S. F., Jones, O. D., Lambeth, S. P., et al. (2007). Endowment effects in chimpanzees. *Current Biology* 17(19): 1704–7.
- Brown, T. C. and Gregory, R. (1999). Why the WTA-WTP disparity matters. *Ecological Economics* 28(3): 323–35.
- Bruhin, A., Fehr-Duda, H., and Epper, T. (2010). Risk and rationality: uncovering heterogeneity in probability distortion. *Econometrica* 78(4): 1375–412.
- Buffon, G. L. L. (1777). *Essai d'Arithmétique Morale*, Supplément à l'Histoire naturelle, vol. IV.
- Burks, S., Carpenter, J., Götte, L., Monaco, K., et al. (2008). Using behavioral economic field experiments at a large motor carrier: the context and design of the truckers and turnover project. In S. Bender, J. Lane, K. Shaw, F. Anderson, and T. von Wachter (eds.), *The Analysis of Firms and Employees: Quantitative and*

- Qualitative Approaches*. Chicago: University of Chicago Press.
- Busemeyer, J. R. and Townsend, J. T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review* 100(3): 432–59.
- Busemeyer, J. R., Weg, E., Barkan, R., Li, X., and Ma, Z. (2000). Dynamic and consequential consistency of choices between paths of decision trees. *Journal of Experimental Psychology General* 129(4): 530–45.
- Camerer, C. F. (1989). An experimental test of several generalized utility theories. *Journal of Risk and Uncertainty* 2(1): 61–104.
- Camerer, C. F. (1992). Recent tests of generalizations of expected utility theory. In W. Edwards (ed.), *Utility: Theories, Measurement, and Applications*. Norwell, MA: Kluwer, pp. 207–51.
- Camerer, C. F. (1995). Individual decision making. In J. Kagel and A. E. Roth (eds.), *The Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press.
- Camerer, C. F. (2000). Prospect theory in the wild: evidence from the field. In D. Kahneman and A. Tversky (eds.), *Choices, Values and Frames*. Cambridge: Cambridge University Press, pp. 288–300.
- Camerer, C. F. (2005). Three cheers—psychological, theoretical, empirical—for loss aversion. *Journal of Marketing Research* 42(2): 129–33.
- Camerer, C. F., Babcock, L., Loewenstein, G., and Thaler, R. H. (1997). Labor supply of New York City cabdrivers: one day at a time. *Quarterly Journal of Economics* 112(2): 407–41.
- Camerer, C. F. and Ho, T.-H. (1994). Violations of the betweenness axiom and nonlinearity in probability. *Journal of Risk and Uncertainty* 8(2): 167–96.
- Camerer, C. F. and Hogarth, R. M. (1999). The effects of financial incentives in experiments: a review and capital-labor-production framework. *Journal of Risk and Uncertainty* 19(1–3): 7–42.
- Carbone, E. and Hey, J. D. (1994). Estimation of expected utility and non-expected utility preference functionals using complete ranking data. In B. Munier and M. J. Machina (eds.), *Models and Experiments on Risk and Rationality*. Amsterdam: Kluwer Academic Publishers, pp. 119–39.
- Carmon, Z. and Ariely, D. (2000). Focusing on the forgone: how value can appear so different to buyers and sellers. *Journal of Consumer Research* 27(3): 360–70.
- Cerreia-Vioglio, S., Ghirardato, P., Maccheroni, F., Marinacci, M., and Siniscalchi, M. (2011). Rational preferences under ambiguity. *Economic Theory* 48(2–3): 341–75.
- Chakravarty, S. and Roy, J. (2009). Recursive expected utility and the separation of attitudes towards risk and ambiguity: an experimental study. *Theory and Decision* 66(3): 199–228.
- Chapman, G. (1998). Similarity and reluctance to trade. *Journal of Behavioral Decision Making* 11(1): 47–58.
- Charness, G. and Gneezy, U. (2012). Strong evidence for gender differences in risk taking. *Journal of Economic Behavior and Organization* 83(1): 50–58.
- Chateauneuf, A. and Cohen, M. (1994). Risk seeking with diminishing marginal utility in a non-expected utility model. *Journal of Risk and Uncertainty* 9(1): 77–91.
- Chateauneuf, A. and Faro, J. H. (2009). Ambiguity through confidence functions. *Journal of Mathematical Economics* 45(9–10): 535–58.
- Chateauneuf, A. and Tallon, J.-M. (2002). Diversification, convex preferences and non-empty core in the Choquet expected utility model. *Economic Theory* 19(3): 509–23.
- Chateauneuf, A. and Wakker, P. P. (1993). From local to global additive representation. *Journal of Mathematical Economics* 22(6): 523–45.
- Chateauneuf, A. and Wakker, P. P. (1999). An axiomatization of cumulative prospect theory for decision under risk. *Journal of Risk and Uncertainty* 18(2): 137–45.
- Chen, M. K., Lakshminarayanan, V., and Santos, L. (2006). How basic are behavioral biases? Evidence from capuchin monkey trading behavior. *Journal of Political Economy*. 114(3): 517–32.
- Cheung, S. (2014). On the elicitation of time preference under conditions of risk. Forthcoming in *American Economic Review*.
- Chew, S. H. (1983). A generalization of the quasi-linear mean with applications to the measurement of income inequality and decision theory resolving the Allais paradox. *Econometrica* 51(4): 1065–92.

- Chew, S. H. (1989). Axiomatic utility theories with the betweenness property. *Annals of Operations Research* 19(1): 273–98.
- Chew, S. H. and Epstein, L. G. (1989a). A unifying approach to axiomatic non-expected utility theories. *Journal of Economic Theory* 49(2): 207–40.
- Chew, S. H. and Epstein, L. G. (1989b). The structure of preferences and attitudes towards the timing of the resolution of uncertainty. *International Economic Review* 30(1): 103–17.
- Chew, S. H. and Epstein, G. L. (1990). Nonexpected utility preferences in a temporal framework with an application to consumption-savings behaviour. *Journal of Economic Theory* 50(1): 54–81.
- Chew, S. H., Epstein, L. G., and Segal, U. (1991). Mixture symmetry and quadratic utility. *Econometrica* 59(1): 139–63.
- Chew, S. H., Karni, E., and Safra, Z. (1987). Risk aversion in the theory of expected utility with rank-dependent preferences. *Journal of Economic Theory* 42(2): 370–81.
- Chew, S. H. and MacCrimmon, K. R. (1979). Alpha-nu choice theory: a generalization of expected utility theory. University of British Columbia. Working Paper 669.
- Chew, S. H. and Sagi, J. S. (2006). Event exchangeability: small worlds probabilistic sophistication without continuity or monotonicity. *Econometrica* 74(3): 771–86.
- Chew, S. H. and Sagi, J. (2008). Small worlds: modeling attitudes toward sources of uncertainty. *Journal of Economic Theory* 139(1): 1–24.
- Chew, S. H. and Waller, W. S. (1986). Empirical tests of weighted utility theory. *Journal of Mathematical Psychology* 30(1): 55–72.
- Coombs, C. H., Bezbinder, T. G., and Goode, F. M. (1967). Testing expectation theories of decision making without measuring utility or subjective probability. *Journal of Mathematical Psychology* 4(1): 72–103.
- Coble, K. H. and Lusk, J. L. (2010). At the nexus of risk and time preferences: an experimental investigation. *Journal of Risk and Uncertainty* 41(1): 67–79.
- Cohen, M. and Jaffray, J.-Y. (1980). Rational behavior under complete ignorance. *Econometrica* 48(5): 1281–99.
- Cohen, M., Jaffray, J.-Y., and Said, T. (1987). Experimental comparison of individual behavior under risk and under uncertainty for gains and for losses. *Organizational Behavior and Human Decision Processes* 39(1): 1–22.
- Conlisk, J. (1989). Three variants on the Allais example. *American Economic Review* 79(3): 392–407.
- Conrad, J., Dittmar, R. F., and Ghysels, E. (2013). Ex ante skewness and expected stock returns. *Journal of Finance* 68(1): 85–124.
- Conte, A. and Hey, J. D. (2013). Assessing multiple prior models of behaviour under ambiguity. *Journal of Risk and Uncertainty* 46(2): 113–32.
- Conte, A., Hey, J. D., and Moffatt, P. G. (2011). Mixture models of choice under risk. *Journal of Econometrics* 162(1): 79–88.
- Coursey, D. L., Hovis, J. L., and Schulze, W. D. (1987). The disparity between willingness to accept and willingness to pay measures of value. *Quarterly Journal of Economics* 102(3): 679–90.
- Cox, J. C. and Sadiraj, V. (2006). Small- and large-stakes risk aversion: implications of concavity calibration for decision theory. *Games and Economic Behavior* 56(1): 45–60.
- Crawford, V. P. and Meng, J. (2011). New York City cab drivers' labor supply revisited: reference-dependent preferences with rational-expectations targets for hours and income. *American Economic Review* 101(5): 1912–32.
- Croson, R. and Gneezy, U. (2009). Gender differences in preferences. *Journal of Economic Literature* 47(2): 448–74.
- Cubitt, R. P., Starmer, C., and Sugden, R., (1998). Dynamic choice and the common ratio effect: and experimental investigation. *Economic Journal* 108(450): 1362–80.
- Cummings, R. G., Brookshire, D. S., and Schulze, W. D. eds. (1986). *Valuing Environmental Goods: An Assessment of the Contingent Valuation Method*. Totowa: Rowman and Allanheld.
- Curley, S. P., Yates, J. F., and Abrams, R. A. (1986). Psychological sources of ambiguity avoidance. *Organizational Behavior and Human Decision Processes* 38(2): 230–56.
- de Finetti, B. (1937). La prevision: ses lois logiques, ses sources subjectives. *Annales de l'Institut Henri Poincare* 7(1): 1–68.
- De Giorgi, E. G., and Legg, S. (2012). Dynamic portfolio choice and asset pricing with narrow framing and probability weighting. *Journal of Economic Dynamics and Control* 36(7): 951–72.

- de Meza, D. and Webb, D. C. (2003). Principal agent problems under loss aversion: an application to executive stock options. Mimeo, London School of Economics.
- de Meza, D. and Webb, D. C. (2007). Incentive design under loss aversion. *Journal of the European Economic Association* 5(1): 66–92.
- Dekel, E. (1986). An axiomatic characterization of preferences under uncertainty: weakening the independence axiom. *Journal of Economic Theory* 40(2): 304–18.
- Dhami, S. and al-Nowaihi, A. (2007). Why do people pay taxes: expected utility versus prospect theory. *Journal of Economic Behavior and Organization* 64(1): 171–92.
- Dhami, S. and al-Nowaihi, A. (2010a). The behavioral economics of crime and punishment. University of Leicester. Discussion Paper 1014.
- Dhami, S. and al-Nowaihi, A. (2010b). Optimal income tax taxation in the presence of tax evasion: expected utility versus prospect theory. *Journal of Economic Behavior and Organization* 75(2): 313–37.
- Dhami, S. and al-Nowaihi, A. (2013). An extension of the Becker Proposition to non-expected utility theory. *Mathematical Social Sciences* 65(1): 10–20.
- Dhami, S. and al-Nowaihi, A. (2014). Prospect theory, crime and punishment. Forthcoming in K. Zeiler and J. C. Teitelbaum (eds.), *Research Handbook on Behavioral Law and Economics*. Cheltenham: Edward Elgar.
- Dhar, R. and Wertenbroch, K. (2000). Consumer choice between hedonic and utilitarian goods. *Journal of Marketing Research* 37(1): 60–71.
- Diecidue, E., Ulrich, S., and Horst, Z. (2009). Parametric weighting functions. *Journal of Economic Theory* 144(3): 1102–18.
- Dimmock, S. G. and Kouwenberg, R. (2010). Loss-aversion and household portfolio choice. *Journal of Empirical Finance* 17(3): 441–59.
- Dohmen, T., Falk, A., Huffman, D., and Sunde, U. (2010). Are risk aversion and impatience related to cognitive ability? *American Economic Review* 100(3): 1238–60.
- Dohmen, T., Falk, A., Huffman, D., and Sunde, U. (2012). The intergenerational transmission of risk and trust attitudes. *Review of Economic Studies* 79(2): 645–77.
- Dohmen, T., Falk, A., Huffman, D., et al. (2011). Individual risk attitudes: measurement, determinants and behavioral consequences. *Journal of the European Economic Association* 9(3): 522–50.
- Du, N. and Budescu, D. (2005). The effects of imprecise probabilities and outcomes in evaluating investment options. *Management Science* 51(12): 1791–803.
- Eckel, C. C. and Grossman, P. J. (2008). Men, women and risk aversion: experimental evidence. In C. R. Plott and V. L. Smith (eds.), *Handbook of Experimental Economics Results, Volume 1*. Amsterdam: Elsevier, ch. 113, pp. 1061–73.
- Edwards, W. (1954). The theory of decision making. *Psychological Bulletin* 51(4): 380–417.
- Edwards, W. (1962). Subjective probabilities inferred from decisions. *Psychological Review* 69(2): 109–35.
- Eichberger, J., Grant, S., Kelsey, D., and Koshevoy, G. A. (2011). The α -MEU model: a comment. *Journal of Economic Theory* 146(4): 1684–98.
- Eliasz, K. and Schotter, A. (2007). Experimental testing of intrinsic preferences for NonInstrumental information. *American Economic Review* 97(2): 166–9.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics* 75(4): 643–69.
- Engelmann, D. and Hollard, G. (2008). A shock therapy against the endowment effect. Mimeo, Royal Holloway, University of London.
- Epper, T. and Fehr-Duda, H. (2014a). Balancing on a budget line: comment on Andreoni and Sprenger (2012)’s “Risk preferences are not time preferences.” Forthcoming in *American Economic Review*.
- Epper, T. and Fehr-Duda, H. (2014b). The missing link: unifying risk taking and time discounting. Mimeo, University of Zurich.
- Epstein, L. G. (1999). A definition of uncertainty aversion. *Review of Economic Studies* 66(3): 579–608.
- Epstein, L. G. (2009). Three paradoxes for the smooth ambiguity model of preference. Mimeo, Boston University.
- Epstein, L. G. and Le Breton, M. (1993). Dynamically consistent beliefs must be Bayesian. *Journal of Economic Theory* 61(1): 1–22.
- Epstein, L. G. and Zhang, J. (2001). Subjective probabilities on subjectively unambiguous events. *Econometrica* 69(2): 265–306.

- Epstein, L. G. and Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. *Econometrica* 57(4): 937–69.
- Erev, I., Ert, E., and Roth, A. E. (2010). A choice prediction competition: choices from experience and from description. *Journal of Behavioral Decision Making* 23: 15–47.
- Ergin, A. and Gul, F. (2009). A theory of subjective compound lotteries. *Journal of Economic Theory* 144(3): 899–929.
- Ericson, K. M. M. and Fuster, A. (2011). Expectations as endowments: evidence on reference-dependent preferences from exchange and valuation experiments. *Quarterly Journal of Economics* 126(4): 1879–907.
- Etchart-Vincent, N. (2004). Is probability weighting sensitive to the magnitude of consequences? An experimental investigation on losses. *Journal of Risk and Uncertainty* 28(3): 217–35.
- Etner, J., Jeleva, M., and Tallon, J.-M. (2012). Decision theory under ambiguity. *Journal of Economic Surveys* 26(2): 234–70.
- Fafchamps, M., Kebede, B., and Zizzo, D. J. (2014). Keep up with the winners: experimental evidence on risk taking, asset integration, and peer effects. University of East Anglia, CBESS. Discussion Paper 14–03.
- Falk, A. and Knell, M. (2004). Choosing the Joneses: endogenous goals and reference standards. *Scandinavian Journal of Economics* 106(3): 417–35.
- Farber, H. S. (2005). Is tomorrow another day? The labor supply of New York City cabdrivers. *Journal of Political Economy* 113(1): 46–82.
- Farber, H. S. (2008). Reference-dependent preferences and labor supply: the case of New York City taxi drivers. *American Economic Review* 98(3): 1069–82.
- Fehr, D., Hakimova, R., and Kübler, D. (2015). The willingness to pay–willingness to accept gap: A Failed Replication of Plott and Zeiler. Mimeo, WZB Berlin Social Science Centre.
- Fehr, E. and Goette, L. (2007). Do workers work more if wages are higher? Evidence from a randomized field experiment. *American Economic Review* 97(1): 298–317.
- Fehr, E., Hart, O., and Zehnder, C. (2011). Contracts as reference points—experimental evidence. *American Economic Review* 101(2): 493–525.
- Fehr, E., Zehnder, C., and Hart, O. (2009). Contracts, reference points, and competition—behavioral effects of the fundamental transformation. *Journal of the European Economic Association* 7(2–3): 561–72.
- Fehr-Duda, H., Bruhin, A., Epper, T., and Schubert, R. (2010). Rationality on the rise: why relative risk aversion increases with stake size. *Journal of Risk and Uncertainty* 40(2): 147–80.
- Fehr-Duda, H. and Epper, T. (2012). Probability and risk: foundations and economic implications of probability-dependent risk preferences. *Annual Review of Economics* 4: 567–93.
- Fehr-Duda, H., Epper, T., Bruhin, A., and Schubert, R. (2011). Risk and rationality: the effects of mood and decision rules on probability weighting. *Journal of Economic Behavior and Organization* 78(1–2): 14–24.
- Fellner, G. and Sutter, M. (2009). Causes, consequences, and cures of myopic loss aversion: an experimental investigation. *Economic Journal* 119(537): 900–16.
- Fennema, H. and van Assen, M. (1998). Measuring the utility of losses by means of the trade-off method. *Journal of Risk and Uncertainty* 17(3): 277–96.
- Fischhoff, B., Slovic, P., Lichtenstein, S., Read, S., and Combs, B. (1978). How safe is safe enough? A psychometric study of attitudes towards technological risks and benefits. *Policy Sciences* 9(2): 127–52.
- Fishburn, P. C. (1978). On Handa's “New theory of cardinal utility” and the maximization of expected return. *Journal of Political Economy* 86(2): 321–4.
- Fishburn, P. C. (1982a). Nontransitive measurable utility. *Journal of Mathematical Psychology* 26(1): 31–67.
- Fishburn, P. C. (1982b). *The Foundations of Expected Utility*. Dordrecht: D. Reidel Publishing Company.
- Fishburn, P. C. (1983). Transitive measurable utility. *Journal of Economic Theory* 31(2): 293–317.
- Fishburn, P. C. and Kochenberger, G. A. (1979). Two-piece von Neumann Morgenstern utility functions. *Decision Sciences* 10(4): 503–18.
- Fox, C. R. and Tversky, A. (1995). Ambiguity aversion and comparative ignorance. *Quarterly Journal of Economics* 110(3): 585–603.

- Fox, C. R. and Tversky, A. (1998). A belief-based account of decision under uncertainty. *Management Science* 44(7): 879–95.
- Fox, C. R. and Weber, M. (2002). Ambiguity aversion, comparative ignorance, and decision context. *Organizational Behavior and Human Decision Processes* 88(1): 476–98.
- Fragaszy, D. M., Visalberghi, E., and Fedigan, L. M. (2004). *The Complete Capuchin: The Biology of the Genus Cebus*. Cambridge: Cambridge University Press.
- French, K. R. and Poterba, J. M. (1991). Investor diversification and international equity markets. *American Economic Review* 81(2): 222–6.
- Friedman, M. and Savage, L. (1948). The utility analysis of choices involving risk. *Journal of Political Economy* 56(4): 279–304.
- Fryer, R. G., Levitt, S. D., List, J. A., and Sadoff, S. (2012). Enhancing the efficacy of teacher incentives through loss aversion: a field experiment. NBER. Working Paper 8237.
- Gächter, S., Johnson, E. J., and Herrmann, A. (2010). Individual-level loss aversion in riskless and risky choices. University of Nottingham, CeDex. Discussion Paper 1749–3293.
- Gajdos, T., Hayashi, T., Tallon, J.-M., and Vergnaud, J.-C. (2008). Attitude toward imprecise information. *Journal of Economic Theory* 140(1): 27–65.
- Genesove, D. and Mayer, C. (2001). Loss aversion and seller behavior: evidence from the housing market. *Quarterly Journal of Economics* 116(4): 1233–60.
- Gennaioli, N. and Shleifer, A. (2010). What comes to mind. *Quarterly Journal of Economics* 125(4): 1399–433.
- Ghirardato, P., Maccheroni, F., and Marinacci, M. (2004). Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory* 118(2): 133–73.
- Ghirardato, P., Maccheroni, F., Marinacci, M., and Siniscalchi, M. (2003). A subjective spin on roulette wheels. *Econometrica* 71(6): 1897–908.
- Ghirardato, P. and Marinacci, M. (2002). Ambiguity made precise: a comparative foundation. *Journal of Economic Theory* 102(2): 251–89.
- Ghirardato, P. and Siniscalchi, M. (2012). Ambiguity in the small and in the large. *Econometrica* 80(6): 2827–47.
- Gilboa, I. (1987). Expected utility with purely subjective non-additive probabilities. *Journal of Mathematical Economics* 16(1): 65–88.
- Gilboa, I. (2009). *Theory of Decision under Uncertainty*. Cambridge: Cambridge University Press.
- Gilboa, I., Lieberman, O., and Schmeidler, D. (2006). Empirical similarity. *Review of Economics and Statistics* 88(3): 433–44.
- Gilboa, I., Postlewaite, A., and Schmeidler, D. (2009). Is it always rational to satisfy Savage's axioms? *Economics and Philosophy* 25(3): 285–96.
- Gilboa, I., Samet, D., and Schmeidler, D. (2004). Utilitarian aggregation of beliefs and tastes. *Journal of Political Economy* 112(4): 932–8.
- Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics* 18(2): 141–53.
- Gilboa, I. and Schmeidler, D. (1995). Case-based decision theory. *Quarterly Journal of Economics* 110(3): 605–39.
- Gilboa, I. and Schmeidler, D. (1996). Act similarity in case-based decision theory. *Economic Theory* 9: 47–62.
- Gilboa, I. and Schmeidler, D. (1997). Act-similarity in case-based decision theory. *Economic Theory* 9(1): 47–61.
- Gilboa, I. and Schmeidler, D. (2001). *A Theory of Case-Based Decisions*. New York: Cambridge University Press.
- Gill, D. and Prowse, V. (2012). A structural analysis of disappointment aversion in a real effort competition. *American Economic Review* 102(1): 469–503.
- Gintis, H. (2009). *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences*. Princeton, NJ: Princeton University Press.
- Gneezy, U., Kapteyn, A., and Potters, J. J. M. (2003). Evaluation periods and asset prices in a market experiment. *Journal of Finance* 58(2): 821–38.
- Gneezy, U. and Potters, J. (1997). An experiment on risk taking and evaluation periods. *Quarterly Journal of Economics* 112(2): 631–45.
- Goette, L. and Huffman, D. (2003). Reference dependent preferences and the allocation of effort over time: evidence from natural experiments. Mimeo, Institute for Empirical Research in Economics, University of Zurich.
- Goette, L., Huffman, D., and Fehr, E. (2004). Loss aversion and labor supply. *Journal of the European Economic Association* 2(2–3): 216–28.

- Goldstein, W. M. and Einhorn, H. J. (1987). Expression theory and the preference reversal phenomena. *Psychological Review* 94(2): 236–54.
- Gonzalez, R. and Wu, G. (1999). On the shape of the probability weighting function. *Cognitive Psychology* 38(1): 129–66.
- Grant, S. (1995). Subjective probability without monotonicity: or how Machina's mom may also be probabilistically sophisticated. *Econometrica* 63(1): 159–89.
- Grant, S., Kajii, A., and Polak, B. (2000). Temporal resolution of uncertainty and recursive non-expected utility models. *Econometrica* 68(2): 425–34.
- Grant, S., Özsoy, H., and Polak, B. (2008). Probabilistic sophistication and stochastic monotonicity in the Savage framework. *Mathematical Social Sciences* 55(3): 371–80.
- Green, J. R. and Jullien, B. (1988). Ordinal independence in nonlinear utility theory. *Journal of Risk and Uncertainty* 1(4): 355–87.
- Green, T. and Hwang, B.-H. (2012). Initial public offerings as lotteries: expected skewness and first-day returns. *Management Science* 58(2): 432–44.
- Grossman, S. J. and Hart, O. D. (1986). The costs and benefits of ownership: a theory of vertical and lateral integration. *Journal of Political Economy* 94(4): 691–719.
- Guerdjikova, A. (2008). Case-based learning with different similarity functions. *Games and Economic Behavior* 63(1): 107–32.
- Gul, F. (1991). A theory of disappointment aversion. *Econometrica* 59(3): 667–86.
- Gul, F. and Pesendorfer, W. (2013). Expected uncertainty utility theory. Mimeo, Princeton University.
- Gurevich, G., Kliger, D., and Levy, O. (2009). Decision-making under uncertainty: a field study of cumulative prospect theory. *Journal of Banking and Finance* 33(7): 1221–9.
- Haigh, M. S. and List, J. A. (2005). Do professional traders exhibit myopic loss aversion? An experimental analysis. *Journal of Finance* 60(1): 523–34.
- Halevy, Y. (2007). Ellsberg revisited: an experimental study. *Econometrica* 75(2): 503–36.
- Halevy, Y. (2008). Strotz meets Allais: diminishing impatience and the certainty effect. *American Economic Review* 98(3): 1145–62.
- Halevy, Y. and Feltkamp, V. (2005). A Bayesian approach to uncertainty aversion. *Review of Economic Studies* 72(2): 449–66.
- Halonon-Akatwijuka, M. and Hart, O. (2013). More is less: why parties may deliberately write incomplete contracts. Mimeo, Harvard University.
- Handa, J. (1977). Risk, probability, and a new theory of cardinal utility. *Journal of Political Economy* 85(1): 97–122.
- Hansen, L. P. and Sargent, T. J. (2001a). Acknowledging misspecification in macroeconomic theory. *Review of Economic Dynamics* 4(3): 519–35.
- Hansen, L. P. and Sargent, T. J. (2001b). Robust control and model uncertainty. *American Economic Review* 91(2): 60–6.
- Harbaugh, W. T., Krause, K., and Vesterlund, L. (2002). Risk attitudes of children and adults: choices over small and large probability gains and losses. *Experimental Economics* 5(1): 53–84.
- Harbaugh, W. T., Krause, K., and Vesterlund, L. (2010). The fourfold pattern of risk attitudes in choice and pricing tasks. *Economic Journal* 120(45): 595–611.
- Harless, D. W. (1992). Predictions about indifference curves inside the unit triangle: a test of variants of expected utility. *Journal of Economic Behavior and Organization* 18(3): 391–414.
- Harrison, G. W., Humphrey, S. J., and Verschoor, A. (2010). Choice under uncertainty: evidence from Ethiopia, India and Uganda. *Economic Journal* 120(543): 80–104.
- Harrison, G. W., Martínez-Correa, J., and Swarthout, J. T. (2012). Reduction of compound lotteries with objective probabilities: theory and evidence. Andrew Young School of Policy Studies, Experimental Economics Center. Working Paper 2012–04.
- Harrison, G. W. and Rutström, E. E. (2009). Expected utility theory and prospect theory: one wedding and a decent funeral. *Experimental Economics* 12(2): 133–58.
- Hart, O. (2009). Hold-up, asset ownership, and reference points. *Quarterly Journal of Economics* 124(1): 267–300.
- Hart, O. and Holmstrom, B. (2010). A theory of firm scope. *Quarterly Journal of Economics* 125(2): 483–513.
- Hart, O. and Moore, J. (1990). Property rights and the nature of the firm. *Journal of Political Economy* 98(6): 1119–58.

- Hart, O. and Moore, J. (2008). Contracts as reference points. *Quarterly Journal of Economics* 123(1): 1–48.
- Hayashi, T. and Wada, R. (2010). Choice with imprecise information: an experimental approach. *Theory and Decision* 69(3): 355–73.
- Heath, C., Larrick, R. P., and Wu, G. (1999). Goals as reference points. *Cognitive Psychology* 38(1): 79–109.
- Heath, C. and Tversky, A. (1991). Preference and belief: ambiguity and competence in choice and uncertainty. *Journal of Risk and Uncertainty* 4(1): 5–28.
- Heffetz, O. and List, J. A. (2014). Is the endowment effect an expectations effect? *Journal of the European Economic Association* 12(5): 1396–422.
- Heidhues, P. and Köszegi, B. (2005). The impact of consumer loss aversion on pricing. CEPR. Discussion Paper 4849.
- Heidhues, P. and Köszegi, B. (2008). Competition and price variation when consumers are loss averse. *American Economic Review* 98(4): 1245–68.
- Heiner, R. A. (1983). The origin of predictable behavior. *American Economic Review* 73(4): 560–95.
- Helson, H. (1964). *Adaptation Level Theory: An Experimental and Systematic Approach to Behavior*. New York: Harper and Row.
- Hershey, J. C. and Schoemaker, P. J. H. (1985). Probability vs. certainty equivalence methods in utility measurement: are they equivalent? *Management Science* 31(10): 1213–31.
- Herweg, F. (2013). The expectation-based loss-averse newsvendor. *Economics Letters* 120(3): 429–32.
- Herweg, F., Karle, H., and Müller, D. (2014). Incomplete contracting, renegotiation, and expectation-based loss aversion. CEPR. Discussion Paper 9874.
- Herweg, F. and Mierendorff, K. (2013). Uncertain demand, consumer loss aversion, and fiat rate tariffs. *Journal of the European Economic Association* 11(2): 399–432.
- Herweg, F., Müller, D., and Weinschenk, P. (2010). Binary payment schemes: moral hazard and loss aversion. *American Economic Review* 100(5): 2451–77.
- Herweg, F. and Schmidt, K. M. (2013). Loss aversion and inefficient renegotiation. Mimeo, University of Munich.
- Hey, J. D. (2014). Choice under uncertainty: empirical methods and experimental results. In M. J. Machina and W. K. Viscusi (eds.), *Handbook of the Economics of Risk and Uncertainty, Volume 1*. Oxford: North Holland, pp. 809–50.
- Hey, J. D., Lotito, G., and Maffioletti, A. (2010). The descriptive and predictive adequacy of theories of decision making under uncertainty/ambiguity. *Journal of Risk and Uncertainty* 41(2): 81–111.
- Hey, J. D. and Orme, C. (1994). Investigating generalizations of expected utility theory using experimental data. *Econometrica* 62(6): 1291–326.
- Hogarth, R. M. and Einhorn, H. J. (1990). Venture theory: a model of decision weights. *Management Science* 36(7): 780–803.
- Hogarth, R. M. and Kunreuther, H. (1989). Risk, ambiguity and insurance. *Journal of Risk and Uncertainty* 2(1): 5–35.
- Hogarth, R. M. and Kunreuther, H. (1992). Pricing insurance and warranties: ambiguity and correlated risks. *The Geneva Papers on Risk and Insurance Theory* 17(1): 35–60.
- Holt, C. A. (1986). Preference reversals and the independence axiom. *American Economic Review* 76(3): 508–15.
- Holt, C. A. and Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review* 92(5): 1644–55.
- Holt, C. A. and Laury, S. K. (2005). Risk aversion and incentive effects: new data without order effects. *American Economic Review* 95(3): 902–12.
- Horowitz, J. K. and McConnell, K. E. (2002). A review of WTA/WTP studies. *Journal of Environmental Economics and Management* 44(3): 426–47.
- Hu, W.-Y. and Scott, J. S. (2007). Behavioral obstacles in the annuity market. *Financial Analysts Journal* 63(6): 71–82.
- Huber, J., Payne, J. W., and Puto, C. (1982). Adding asymmetrically dominated alternatives: violations of regularity and the similarity hypothesis. *Journal of Consumer Research* 9(1): 90–8.
- Humphrey, S. J. (1995). Regret aversion or event-splitting effects? More evidence under risk and uncertainty. *Journal of Risk and Uncertainty* 11(3): 263–74.
- Humphrey, S. J. (1998). More mixed results on boundary effects. *Economics Letters* 61(1): 79–84.

- Humphrey, S. J. (2000). The common consequence effect: testing a unified explanation of recent mixed evidence. *Journal of Economic Behavior and Organization* 41(3): 239–62.
- Hurwicz, L. (1951). Some specification problems and applications to econometric models. *Econometrica* 19: 343–4.
- Irwin, J. R. (1994). Buying/selling price preference reversals: preference for environmental changes in buying versus selling modes. *Organizational Behavior and Human Decision Processes* 60(3): 431–57.
- Isoni, A., Loomes, G., and Sugden, G. (2011). The willingness to pay–willingness to accept gap, the “endowment effect,” subject misconceptions, and experimental procedures for eliciting valuations: comment. *American Economic Review* 101(2): 991–1011.
- Johnson, E. J., Gächter, S., and Herrmann, A. (2006). Exploring the nature of loss aversion. CeDEx. Discussion Paper 2006–02.
- Johnson, E. J., Haeubel, G., and Anat, K. (2007). Aspects of endowment: a query theory of loss aversion. *Journal of Experimental Psychology: Learning Memory and Cognition* 33(3): 461–74.
- Johnson, E. J. and Schkade, D. A. (1989). Bias in utility assessments: further evidence and explanations. *Management Science* 35(4): 406–24.
- Jullien, B. and Salanié, B. (2000). Estimating preferences under risk: the case of racetrack bettors. *Journal of Political Economy* 108(3): 503–30.
- Kachelmeier, S. J. and Shehata, M. (1992). Examining risk preferences under high monetary incentives: experimental evidence from the People’s Republic of China. *American Economic Review* 82(5): 1120–41.
- Kagel, J. H., MacDonald, D. N., and Battalio, R. C. (1990). Tests of “fanning out” of indifference curves: results from animal and human experiments. *American Economic Review* 80(4): 912–21.
- Kahneman, D. (2003). Maps of bounded rationality: psychology for behavioral economics. *American Economic Review* 93(5): 1449–75.
- Kahneman, D. (2012). *Thinking Fast and Slow*. London: Penguin Group.
- Kahneman, D., Knetsch, J. L., and Thaler, R. H. (1986). Fairness as a constraint on profit seeking: entitlements in the market. *American Economic Review* 76(4): 728–41.
- Kahneman, D., Knetsch, J. L., and Thaler, R. H. (1990). Experimental tests of the endowment effect and the Coase Theorem. *Journal of Political Economy* 98(6): 1325–48.
- Kahneman, D., Knetsch, J. L., and Thaler, R. H. (1991). Anomalies: the endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives* 5(1): 193–206.
- Kahneman, D. and Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica* 47(2): 263–91.
- Kahneman, D. and Tversky, A. (2000). *Choices, Values and Frames*. Cambridge: Cambridge University Press.
- Karmarkar, U. S. (1979). Subjectively weighted utility and the Allais paradox. *Organizational Behavior and Human Performance* 24(1): 67–72.
- Karni, E. (2003). On the representation of beliefs by probabilities. *Journal of Risk and Uncertainty* 26(1): 17–38.
- Karni, E. and Safra, Z. (1987). “Preference reversal” and the observability of preferences by experimental methods. *Econometrica* 55(3): 675–85.
- Keynes, J. M. (1921). *A Treatise on Probability*. London: Macmillan.
- Kilka, M. and Weber, M. (2001). What determines the shape of the probability weighting function under uncertainty? *Management Science* 47(12): 1712–26.
- Klibanoff, P., Marinacci, M., and Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica* 73(6): 1849–1892.
- Klibanoff, P., Marinacci, M., and Mukerji, S. (2009). On the smooth ambiguity aversion model: a reply. Oxford University. Discussion Papers in Economics 449.
- Kliger, D. and Levy, O. (2009). Theories of choice under risk: insights from financial markets. *Journal of Economic Behavior and Organization* 71(2): 330–46.
- Knetsch, J. L. (1989). The endowment effect and evidence of nonreversible indifference curves. *American Economic Review* 79(5): 1277–84.
- Knetsch, J. L. and Sinden, J. A. (1984). Willingness to pay and compensation demanded: experimental evidence of an unexpected disparity in measures of value. *Quarterly Journal of Economics* 99(3): 507–21.
- Knetsch, J. L. and Wong, W.-K. (2009). The endowment effect and the reference state: evidence

- and manipulations. *Journal of Economic Behavior and Organization* 71(2): 407–13.
- Knez, P., Smith, V. L., and Williams, A. W. (1985). Individual rationality, market rationality, and value estimation. *American Economic Review* 75(2): 397–402.
- Knight, F. J. (1921). *Risk, Uncertainty, and Profit*. Boston: Houghton Mifflin Company.
- Köbberling, V. and Wakker, P. P. (2003). Preference foundations for non-expected utility: a generalized and simplified technique. *Mathematics of Operations Research* 28(3): 395–423.
- Köbberling, V. and Wakker, P. P. (2005). An index of loss aversion. *Journal of Economic Theory* 122(1): 119–31.
- Kolm, S.-C. (1973). A note on optimum tax evasion. *Journal of Public Economics* 2(3): 265–70.
- Kolstad, C. D. and Guzman, R. M. (1999). Information and the divergence between willingness to accept and willingness to pay. *Journal of Environmental Economics and Management* 38(1): 66–80.
- Köszegi, B. and Rabin, M. (2006). A model of reference-dependent preferences. *Quarterly Journal of Economics* 121(4): 1133–65.
- Köszegi, B. and Rabin, M. (2007). Reference-dependent risk attitudes. *American Economic Review* 97(4): 1047–73.
- Köszegi, B. and Rabin, M. (2009). Reference-dependent consumption plans. *American Economic Review* 99(3): 909–36.
- Kothiyal, A., Spinu, V., and Wakker, P. P. (2011). Prospect theory for continuous distributions: a preference foundation. *Journal of Risk and Uncertainty* 42(3): 195–210.
- Kothiyal, A., Spinu, V., and Wakker, P. P. (2014). An experimental test of prospect theory for predicting choice under ambiguity. *Journal of Risk and Uncertainty* 48(1): 1–17.
- Kreps, D. M. (1990). *A Course in Microeconomic Theory*. Princeton, NJ: Princeton University Press.
- Kreps, D. M. and Porteus, E. L. (1978). Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* 46(1): 185–200.
- Kunreuther, H., Ginsberg, R., Miller, L., et al. (1978). *Disaster Insurance Protection: Public Policy Lessons*. New York: Wiley.
- Kunreuther, H., Meszaros, J., Hogarth, R. M., and Spranca, M. (1995). Ambiguity and underwriter decision processes. *Journal of Economic Behavior and Organization* 26(3): 337–52.
- Kunreuther, H. and Pauly, M. (2004). Neglecting disaster: why don't people insure against large losses? *Journal of Risk and Uncertainty* 28(1): 5–21.
- Kunreuther, H. and Pauly, M. (2005). Insurance decision making and market behaviour. *Foundations and Trends in Microeconomics* 1(2): 63–127.
- Langer, T. and Weber, M. (2008). Does commitment or feedback influence myopic loss aversion? An experimental analysis. *Journal of Economic Behavior and Organization* 67(3–4): 810–19.
- Lakshminarayanan, V., Chen, M. K., and Santos, L. R. (2008). Endowment effect in capuchin monkeys. *Philosophical Transactions of the Royal Society B* 363: 3837–44.
- Lakshminarayanan, V., Chen, M. K., and Santos, L. R. (2011). The evolution of decision-making under risk: Framing effects in monkey risk preferences. *Journal of Experimental Social Psychology* 47(3): 689–93.
- Lehrer, E. (2009). A new integral for capacities. *Economic Theory* 39(1): 157–76.
- Lerner, J. S., Small, D. A., and Loewenstein, G. (2004). Heart strings and purse strings: carry-over effects of emotions on economic transactions. *Psychological Science* 15(5): 337–41.
- Levin, I., Schreiber, J., Lauriola, M., and Gaeth, G. J. (2002). A tale of two pizzas: building up from a basic product versus scaling down from a fully loaded product. *Marketing Letters* 13(4): 335–44.
- Levitt, S. D. (2004). Understanding why crime fell in the 1990's: four factors that explain the decline and six that do not. *Journal of Economic Perspectives* 18(1): 163–90.
- Lichtenstein, S. and Slovic, P. (1971). Reversals of preference between bids and choices in gambling decisions. *Journal of Experimental Psychology* 89(1): 46–55.
- Lichtenstein, S. and Slovic, P. (eds.) (2006). *The Construction of Preferences*. Cambridge: Cambridge University Press.
- List, J. A. (2003). Does market experience eliminate market anomalies? *Quarterly Journal of Economics* 118(1): 41–71.
- List, J. A. (2004). Neoclassical theory versus prospect theory: evidence from the marketplace. *Econometrica* 72(2): 615–25.

- List, J. A. and Haigh, M. S. (2005). A simple test of expected utility theory using professional traders. *Proceedings of the National Academy of Sciences* 102(3): 945–8.
- Locke, E. A. and Latham, G. P. (1990). *A Theory of Goal Setting and Task Performance*. Englewood Cliffs, NJ: Prentice-Hall.
- Loewenstein, G. F., Weber, E. U., and Hsee, C. K. (2001). Risk as feelings. *Psychological Bulletin* 127(2): 267–86.
- Loomes, G. (2008). Modelling choice and valuation in decision experiments. University of East Anglia. Working Paper.
- Loomes, G., Starmer, S., and Sugden, R. (1989). Preference reversal: information-processing effect or rational non-transitive choice? *Economic Journal* 99(395): 140–51.
- Loomes, G., Starmer, C., and Sugden, R. (1991). Observing violations of transitivity by experimental methods. *Econometrica* 59(2): 425–39.
- Loomes, G., Starmer, C., and Sugden, R. (1992). Are preferences monotonic? Testing some predictions of regret theory. *Economica* 59(233): 17–33.
- Loomes, G., Starmer, C., and Sugden, R. (2003). Do anomalies disappear in repeated markets? *Economic Journal* 113(486): C153–C166.
- Loomes, G. and Sugden, R. (1982). Regret theory: an alternative theory of rational choice under uncertainty. *Economic Journal* 92(368): 805–24.
- Loomes, G. and Sugden, R. (1986). Disappointment and dynamic consistency in choice under uncertainty. *Review of Economic Studies* 53(2): 271–82.
- Loomes, G. and Sugden, R. (1987). Some implications of a more general form of regret theory. *Journal of Economic Theory* 41(2): 270–87.
- Lovallo, D. and Kahneman, D. (2000). Living with uncertainty: attractiveness and resolution timing. *Journal of Behavioral Decision Making* 13(2): 179–90.
- Luce, R. D. (1991). Rank-and sign-dependent linear utility models for binary gambles. *Journal of Economic Theory* 53(1): 75–100.
- Luce, R. D. (2000). *Utility of Gains and Losses: Measurement-Theoretical and Experimental Approaches*. Mahwah, NJ: Erlbaum.
- Luce R. D. (2001). Reduction invariance and Prelec's weighting functions. *Journal of Mathematical Psychology* 45(1): 167–79.
- Luce, R. D. and Fishburn, P. C. (1991). Rank-and sign-dependent linear utility models for finite first-order gambles. *Journal of Risk and Uncertainty* 4(1): 29–59.
- Luce, R. D. and Marley, A. A. J. (2005). Ranked additive utility representations of gambles: old and new axiomatizations. *Journal of Risk and Uncertainty* 30(1): 21–62.
- Luce, R. D. and Raiffa, H. (1957). *Games and Decisions*. New York: Wiley.
- Maccheroni, F., Marinacci, M., and Rustichini, A. (2006a). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* 74(6): 1447–98.
- Maccheroni, F., Massimo M., and Rustichini, A. (2006b). Dynamic variational preference. *Journal of Economic Theory* 128(1): 4–44.
- McClelland, G. H., Schulze, W. D., and Coursey, D. L. (1993). Insurance for low-probability hazards: a bimodal response to unlikely events. *Journal of Risk and Uncertainty* 7(1): 95–116.
- McCord, M. and de Neufville, R. (1986). Lottery equivalents: reduction of the certainty effect problem in utility assessment. *Management Science* 32(1): 56–60.
- MacCrimmon, K. and Larsson, S. (1979). Utility theory: axioms versus paradoxes. In: M. Allais and O. Hagen (eds.), *Expected Utility Hypotheses and the Allais Paradox*. Dordrecht: Reidel.
- MacDonald, D. N., Kagel, J. H., and Battalio, R. C. (1991). Animals' choices over uncertain outcomes: further experimental results. *Economic Journal* 101(408): 1067–84.
- Machina, M. J. (2009). Risk, ambiguity, and the rank-dependence axioms. *American Economic Review* 99(1): 385–92.
- Machina, M. J. and Schmeidler, D. (1992). A more robust definition of subjective probability. *Econometrica* 60(4): 745–80.
- Machina, M. J. and Siniscalchi, M. (2014). Ambiguity and ambiguity aversion. In M. J. Machina and W. K. Viscusi (eds.), *Handbook of the Economics of Risk and Uncertainty, Volume 1*. Oxford: North-Holland.
- Mankiw, N. G. and Zeldes, S. P. (1991). The consumption of stockholders and nonstockholders. *Journal of Financial Economics* 29(1): 97–112.
- Markowitz, H. (1952). The utility of wealth. *Journal of Political Economy* 60(2): 151–8.

- Marley, A. A. J. and Luce, R. D. (2001). Rank-weighted utilities and qualitative convolution. *Journal of Risk and Uncertainty* 23(2): 135–63.
- Marley, A. A. J. and Luce, R. D. (2005). Independence properties vis-à-vis several utility representations. *Theory and Decision* 58(1): 77–143.
- Marquis, M. S. and Holmer, M. R. (1996). Alternative models of choice under uncertainty and demand for health insurance. *Review of Economics and Statistics* 78(3): 421–7.
- Mas, A. (2006). Pay, reference points, and police performance. *Quarterly Journal of Economics* 121(3): 783–821.
- Mas-Collel, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. New York: Oxford University Press.
- Mehra, R. and Prescott, E. C. (1985). The equity premium: a puzzle. *Journal of Monetary Economics* 15(2): 145–61.
- Miao, B. and Zhong, S. (2012). Separating risk preference and time preference. National University of Singapore, Department of Economics. Working Paper.
- Mitton, T. and Vorkink, K. (2007). Equilibrium underdiversification and the preference for skewness. *Review of Financial Studies* 20(4): 1255–88.
- Morrison, G. C. (1997). Resolving differences in willingness to pay and willingness to accept: comment. *American Economic Review* 87(1): 236–40.
- Nakamura, N. (1990). Subjective expected utility with non-additive probabilities on finite state spaces. *Journal of Economic Theory* 51(2): 346–66.
- Nau, R. F. (2006). Uncertainty aversion and second-order utilities and probabilities. *Management Science* 52(1): 136–45.
- Nehring, K. (2001). Ambiguity in the context of probabilistic beliefs. Mimeo, University of California, Davis.
- Neilson, W. S. (1992). A mixed fan hypothesis and its implications for behavior toward risk. *Journal of Economic Behavior and Organization* 19(2): 197–211.
- Neilson, W. S. (2001). Calibration results for rank-dependent expected utility. *Economics Bulletin* 4(10): 1–5.
- Neilson, W. S. (2002). Comparative risk sensitivity with reference-dependent preferences. *Journal of Risk Uncertainty* 24(2): 131–42.
- Noussair, C. and Wu, P. (2006). Risk tolerance in the present and the future: an experimental study. *Managerial and Decision Economics* 27(6): 401–12.
- Novemsky, N. and Kahneman, D. (2005). The boundaries of loss aversion. *Journal of Marketing Research* 42(2): 119–28.
- Ober, J. (2012). Thucydides as prospect theorist. Princeton/Stanford. Working Papers in Classics.
- Odean, T. (1998). Are investors reluctant to realize their losses? *Journal of Finance* 53(5): 1775–98.
- Oettinger, G. S. (1999). An empirical analysis of the daily labor supply of stadium vendors. *Journal of Political Economy* 107(2): 360–92.
- Ok, E. A., Ortoleva, P., and Riella, G. (2015). Revealed (P)Reference Theory. *American Economic Review* 105(1): 299–321.
- Öncüler, A. and Onay, S. (2009). How do we evaluate future gambles? Experimental evidence on path dependency in risky intertemporal choice. *Journal of Behavioral Decision Making* 22(3): 280–300.
- Orwin, C. (1994). *The Humanity of Thucydides*. Princeton, NJ: Princeton University Press.
- Page, L., Savage, D. A., and Torgler, B. (2012). Variation in risk seeking behavior in a natural experiment on large losses induced by a natural disaster. CESifo. Working Paper 3878.
- Palacios-Huerta, I. and Serrano, R. (2006). Rejecting small gambles under expected utility. *Economics Letters* 91(2): 250–9.
- Pape, A. D. and Kurtz, K. J. (2013). Evaluating case-based decision theory: predicting empirical patterns of human classification learning. *Games and Economic Behavior* 82: 52–65.
- Pennings, J. M. E. and Smidts, A. (2003). The shape of utility functions and organizational behavior. *Management Science* 49(9): 1251–63.
- Peters, E., Slovic, P., and Gregory, R. (2003). The role of affect in the WTA/WTP disparity. *Journal of Behavioral Decision Making* 16(4): 309–30.
- Peters, H. J. M. (2012). A preference foundation for constant loss aversion. *Journal of Mathematical Economics* 48(1): 21–5.
- Pfanzagl, J. (1968). *Theory of Measurement*. New York: John Wiley and Sons.
- Plott, C. R. and Zeiler, K. (2005). The willingness to pay–willingness to accept gap, the “endowment effect”, subject misconceptions, and experimental procedures to eliciting valuations. *American Economic Review* 95(3): 530–45.

- Plott, C. R. and Zeiler, K. (2007). Exchange asymmetries incorrectly interpreted as evidence of endowment effect theory and prospect theory? *American Economic Review* 97(4): 1449–66.
- Polinsky, M. and Shavell, S. (2007). The theory of public enforcement of law. In M. Polinsky and S. Shavell (eds.), *Handbook of Law and Economics, Volume 1*. Amsterdam: Elsevier, ch. 6.
- Pope, D. and Schweitzer, M. (2011). Is Tiger Woods loss averse? Persistent bias in the face of experience, competition, and high stakes. *American Economic Review* 101(1): 129–57.
- Post, T., van den Assem, M. J., Baltussen, G., and Thaler, R. H. (2008). Deal or no deal? Decision making under risk in a large-payoff game show. *American Economic Review* 98(1): 38–71.
- Pöystia, L., Rajalina, S., and Summala, H. (2005). Factors influencing the use of cellular (mobile) phone during driving and hazards while using it. *Accident Analysis and Prevention* 37(1): 47–51.
- Prelec, D. (1990). A “Pseudo-endowment” effect, and its implications for some recent nonexpected utility models. *Journal of Risk and Uncertainty* 3(3): 247–59.
- Prelec, D. (1998). The probability weighting function. *Econometrica* 66(3): 497–527.
- Quiggin, J. (1981). Risk perception and risk aversion among Australian farmers. *Australian Journal of Agricultural Economics* 25(2): 160–9.
- Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior and Organization* 3(4): 323–43.
- Quiggin, J. (1993). *Generalized Expected Utility Theory*. Dordrecht: Kluwer.
- Rabin, M. (2000a). Risk aversion and expected-utility theory: a calibration theorem. *Econometrica* 68(5): 1281–92.
- Rabin, M. (2000b). Diminishing marginal utility of wealth cannot explain risk aversion. In D. Kahneman and A. Tversky (eds.), *Choices, Values and Frames*. Cambridge: Cambridge University Press.
- Radelet, M. L. and Ackers, R. L. (1996). Deterrence and the death penalty: the views of the experts. *Journal of Criminal Law and Criminology* 87(1): 1–16.
- Ramsey, F. P. (1931). Truth and probability. In *The Foundation of Mathematics and Other Logical Essays*. New York: Harcourt, Brace.
- Redelmeier, D. A., Koehler, D. J., Liberman, V., and Tversky, A. (1995). Probability judgment in medicine: discounting unspecified causes. *Medical Decision Making* 15(3): 227–30.
- Richard, S. F. (1975). Multivariate risk aversion, utility independence and separable utility functions. *Management Science* 22(1): 12–21.
- Rieger, M. O. and Wang, M. (2006). Cumulative prospect theory and the St. Petersburg paradox. *Economic Theory* 28(3): 665–79.
- Rieger, M. O., Wang, M., and Hens, T. (2014). Prospect theory around the World. Forthcoming in *Management Science*.
- Rizzo, J. A. and Zeckhauser, R. J. (2003). Reference incomes, loss aversion, and physician behaviour. *Review of Economics and Statistics* 85(4): 909–22.
- Roca, M., Hogarth, R. M., and Maule, A. J. (2006). Ambiguity seeking as a result of the status-quo bias. *Journal of Risk and Uncertainty* 32(3): 175–94.
- Rottenstreich, Y. and Hsee, C. K. (2001). Money, kisses, and electric shocks: on the affective psychology of risk. *Psychological Science* 12(3): 185–90.
- Rottenstreich, Y. and Tversky, A. (1997). Unpacking, repacking, and anchoring: advances in support theory. *Psychological Review* 104(2): 406–15.
- Routledge, B. R. and Zin, S. E. (2010). Generalized disappointment aversion and asset prices. *Journal of Finance* 65(4): 1303–32.
- The Royal Society for the Prevention of Accidents (2005). The Risk of Using a Mobile Phone while Driving. Available at http://www.rosipa.com/roadsafety/info/mobile_phone_report.pdf.
- Rubinstein, A. (1998). *Modeling Bounded Rationality*. Cambridge, MA: MIT Press.
- Rubinstein, A. (2006). Dilemmas of an economic theorist. *Econometrica* 74(4): 865–83.
- Rutten-van Mölken, M. P., Bakker, C. H., van Doorslaer, E. K., and van der Linden, S. (1995). Methodological issues of patient utility measurement. Experience from two clinical trials. *Medical Care* 33(9): 922–37.
- Safra, Z. and Segal, U. (2008). Calibration results for non-expected utility theories. *Econometrica* 76(5): 1143–66.
- Samuelson, L. (2005). Economic theory and experimental economics. *Journal of Economic Literature* 43(1): 65–107.
- Samuelson, W. F. and Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty* 1(1): 7–59.

- Santos, L. R. and Platt, M. L. (2014). Evolutionary anthropological insights into neuroeconomics: what non-human primates can tell us about human decision-making strategies. In P. W. Glimcher and E. Fehr (eds.), *Neuroeconomics* Elsevier Inc.
- Sarin, R. and Wakker, P. P. (1992). A simple axiomatization of nonadditive expected utility. *Econometrica* 60(6): 1255–72.
- Sarin, R. and Wakker, P. P. (2000). Cumulative dominance and probabilistic sophistication. *Mathematical Social Sciences*. 40(2): 191–6.
- Savage, L. J. (1954). *The Foundations of Statistics*. New York: Wiley and Sons.
- Sayman, S. and Öncüler, A. (2005). Effects of study design characteristics on the WTA-WTP disparity: a meta analytical framework. *Journal of Economic Psychology* 26(2): 289–312.
- Schmeidler, D. (1986). Integral representation without additivity. *Proceedings of the American Mathematical Society* 97(2): 255–61.
- Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica* 57(3): 571–87.
- Schmidt, U. (2003). Reference-dependence in cumulative prospect theory. *Journal of Mathematical Psychology* 47(2): 122–31.
- Schmidt, U., Starmer, C., and Sugden, R. (2008). Third-generation prospect theory. *Journal of Risk and Uncertainty* 36(3): 203–23.
- Schmidt, U. and Traub, S. (2002). An experimental test of loss aversion. *Journal of Risk and Uncertainty* 25(3): 233–49.
- Schmidt, U. and Zank, H. (2005). What is loss aversion? *Journal of Risk and Uncertainty* 30: 157–67.
- Schmidt, U. and Zank, H. (2007). Linear cumulative prospect theory with applications to portfolio selection and insurance demand. *Decisions in Economics and Finance* 30(1): 1–18.
- Schmidt, U. and Zank, H. (2008). Risk aversion in cumulative prospect theory. *Management Science* 54(1): 208–16.
- Schmidt, U. and Zank, H. (2012). A genuine foundation for prospect theory. *Journal of Risk and Uncertainty* 45(2): 97–113.
- Schunk, D. and Betsch, C. (2006). Explaining heterogeneity in utility functions by individual differences in decision modes. *Journal of Economic Psychology* 27(3): 386–401.
- Schunk, D. and Winter, J. (2009). The relationship between risk attitudes and heuristics in search tasks: a laboratory experiment. *Journal of Economic Behavior and Organization* 71(2): 347–60.
- Segal, U. (1987). The Ellsberg paradox and risk aversion: an anticipated utility approach. *International Economic Review* 28(1): 175–202.
- Segal, U. (1988). Does the preference reversal phenomenon necessarily contradict the independence axiom? *American Economic Review* 78(1): 233–6.
- Segal, U. (1989). Anticipated utility: a measure representation approach. *Annals of Operations Research* 19(1): 359–73.
- Segal, U. (1990). Two-stage lotteries without the reduction axiom. *Econometrica* 58(2): 349–77.
- Segal, U. (1993). The measure representation: a correction. *Journal of Risk and Uncertainty* 6(1): 99–107.
- Sen, S. and Johnson, E. J. (1997). Mere-possession effects without possession in consumer choice. *Journal of Consumer Research* 24(1): 105–17.
- Shefrin, H. and Statman, M. (1985). The disposition to sell winners too early and ride losers too long: theory and evidence. *Journal of Finance* 40(3): 777–90.
- Shogren, J. F., Shin, S. Y., Hayes, D. J., and Kliebenstein, J. B. (1994). Resolving differences in willingness to pay and willingness to accept. *American Economic Review* 84(1): 255–70.
- Siniscalchi, M. (2006). A behavioral characterization of plausible priors. *Journal of Economic Theory* 128(1): 91–135.
- Siniscalchi, M. (2009). Vector expected utility and attitudes toward variation. *Econometrica* 77(3): 801–55.
- Skinner, J. and Slemrod, J. (1985). An economic perspective on tax evasion. *National Tax Journal* 38(3): 345–53.
- Slemrod, J. and Yitzhaki, S. (2002). Tax avoidance, evasion and administration. In A. J. Auerbach and M. Feldstein (eds.), *Handbook of Public Economics, Volume 3*. Amsterdam: Elsevier Science.
- Smith, A. (2008). Lagged beliefs and reference-dependent utility. Department of Economics, University of Arizona, Working Paper 08-03.
- Smithers, S. (2015). Goals, motivation and gender. *Economics Letters* 131: 75–7.
- Snowberg, E. and Wolfers, J. (2010). Explaining the favorite-longshot bias: is it risk-love or misperceptions? *Journal of Political Economy* 118(4): 723–46.

- Spiegler, R. (2011). *Bounded Rationality and Industrial Organization*. Oxford: Oxford University Press.
- Spiegler, R. (2012). Monopoly pricing when consumers are antagonized by unexpected price increases: a “cover version” of the Heidhues-Koszegi-Rabin model. *Economic Theory* 51(3): 695–711.
- Stalmeier, P. F. and Bezembinder, T. G. (1999). The discrepancy between risky and riskless utilities: a matter of framing? *Medical Decision Making* 19(4): 435–47.
- Starmer, C. (1992). Testing new theories of choice under uncertainty using the common consequence effect. *Review of Economic Studies* 59(4): 813–30.
- Starmer, C. (2000). Developments in non-expected utility theory: the hunt for a descriptive theory of choice under risk. *Journal of Economic Literature* 38(2): 332–82.
- Starmer, C. and Sugden, R. (1989a). Violations of the independence axiom in common ratio problems: an experimental test of some competing hypotheses. *Annals of Operations Research* 19(1): 79–102.
- Starmer, C. and Sugden, R. (1989b). Probability and juxtaposition effects: an experimental investigation of the common ratio effect. *Journal of Risk and Uncertainty* 2(2): 159–78.
- Starmer, C. and Sugden, R. (1993). Testing for juxtaposition and eventsplitting effects. *Journal of Risk and Uncertainty* 6(3): 235–54.
- Starmer, C. and Sugden, R. (1998). Testing alternative explanations of cyclical choices. *Economica* 65(259): 347–61.
- Stott, H. P. (2006). Choosing from cumulative prospect theory’s functional menagerie. *Journal of Risk and Uncertainty* 32(2): 101–30.
- Strahilevitz, M. A. and Loewenstein, G. (1998). The effect of ownership history on the valuation of objects. *Journal of Consumer Research* 25(3): 276–89.
- Sugden, R. (1986). New developments in the theory of choice under uncertainty. *Bulletin of Economic Research* 38(1): 1–24.
- Sugden, R. (2003). Reference-dependent subjective expected utility. *Journal of Economic Theory* 111(2): 172–91.
- Sugden, R. (2004). Alternatives to expected utility: foundations. In S. Barbera, P. Hammond, and C. Seidl (eds.), *Handbook of Utility Theory*, Volume 2: Extensions. Dordrecht: Kluwer, pp. 685–755.
- Sutter, M. (2007). Are teams prone to myopic loss aversion? An experimental study on individual versus team investment behavior. *Economics Letters* 97(2): 128–32.
- Sydnor, J. (2010). (Over)insuring modest risks. *American Economic Journal: Applied Economics* 2(4): 177–99.
- Taylor, S. E. and Brown, J. D. (1988). Illusion and well-being: a social psychological perspective on mental health. *Psychological Bulletin* 103(2): 193–210.
- Teigen, K. H. (1974). Overestimation of subjective probabilities. *Scandinavian Journal of Psychology* 15(1): 56–62.
- Thaler, R. H. (1980). Toward a positive theory of consumer choice. *Journal of Economic Behavior and Organization* 1(1): 39–60.
- Thaler, R. H. and Johnson, E. (1990). Gambling with the house money and trying to break even: the effects of prior outcomes on risky choice. *Management Science* 36(6): 643–60.
- Thaler, R. H., Tversky, A., Kahneman, D., and Schwartz, A. (1997). The effect of myopia and loss aversion on risk taking: an experimental test. *Quarterly Journal of Economics* 112(2): 647–61.
- Trautmann, S. T. and van de Kuilen, G. (2016). Ambiguity attitudes. In G. Keren and G. Wu (eds.), *Blackwell Handbook of Judgment and Decision Making, Volume I*. Blackwell, pp. 89–116.
- Trautmann, S. T., Vieider, F. M., and Wakker, P. P. (2008). Causes of ambiguity aversion: known versus unknown preferences. *Journal of Risk and Uncertainty* 36(3): 225–43.
- Trautmann, S. T., Vieider, F. M., and Wakker, P. P. (2011). Preference reversals for ambiguity aversion. *Management Science* 57(7): 1320–33.
- Tversky, A. and Fox, C. R. (1995). Weighing risk and uncertainty. *Psychological Review* 102(2): 269–83.
- Tversky, A. and Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science* 211(4481): 453–58.
- Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4): 297–323.
- Tversky, A. and Koehler, D. J. (1994). Support theory: a nonextensional representation

- of subjective probability. *Psychological Review* 101(4): 547–67.
- Tversky, A. Slovic, P., and Kahneman, D. (1990). The causes of preference reversal. *American Economic Review* 80(1): 204–17.
- Tversky, A., and Thaler, R. H. (1990). Anomalies: preference reversals. *Journal of Economic Perspectives* 4(2): 201–11.
- Tversky, A. and Wakker, P. P. (1995). Risk attitudes and decision weights. *Econometrica* 63(6): 1255–80.
- van de Kuilen, G. and Wakker, P. P. (2011). The midweight method to measure attitudes towards risk and ambiguity. *Management Science* 57(3): 582–98.
- van Dijk, E. and van Knippenberg, D. (1996). Buying and selling exchange goods: loss aversion and the endowment effect. *Journal of Economic Psychology* 17(4): 517–24.
- Varian, H. (1992). *Microeconomic Analysis*. New York: W. W. Norton.
- Viscusi, K. W. (1998). *Rational Risk Policy: The 1996 Arne Ryde Memorial Lectures*. Oxford: Oxford University Press.
- von Gaudecker, H.-M., van Soest, A., and Wengstrom, E. (2011). Heterogeneity in risky choice behaviour in a broad population. *American Economic Review* 101(2): 664–94.
- von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Wakker, P. P. (1994). Separating marginal utility and probabilistic risk aversion. *Theory and Decision* 36(1): 1–44.
- Wakker, P. P. (2001). Testing and characterizing properties of nonadditive measures through violations of the sure-thing principle. *Econometrica* 69(4): 1039–59.
- Wakker, P. P. (2003). The data of Levy and Levy (2002) “Prospect theory: much ado about nothing?” actually support prospect theory. *Management Science* 49(7): 979–81.
- Wakker, P. P. (2005). Formalizing reference dependence and initial wealth in Rabin’s calibration theorem. Mimeo, Econometric Institute, Erasmus University Rotterdam, The Netherlands.
- Wakker, P. P. (2010). *Prospect Theory for Risk and Ambiguity*. Cambridge: Cambridge University Press.
- Wakker, P. P. and Deneffe, D. (1996). Eliciting von Neumann Morgenstern utilities when probabilities are distorted or unknown. *Management Science* 42(8): 1131–50.
- Wakker, P. P., Thaler, R. H., and Tversky, A. (1997). Probabilistic insurance. *Journal of Risk and Uncertainty* 15(1): 7–28.
- Wakker, P. P. and Tversky, A. (1993). An axiomatization of cumulative prospect theory. *Journal of Risk and Uncertainty* 7(2): 147–75.
- Wakker, P. P. and Zank, H. (2002). A simple preference foundation of cumulative prospect theory with power utility. *European Economic Review* 46(7): 1253–71.
- Walther, H. (2003). Normal-randomness expected utility, time preference and emotional distortions. *Journal of Economic Behavior and Organization* 52(2): 253–66.
- Walther, H. (2010). Anomalies in intertemporal choice, time-dependent uncertainty and expected utility—a common approach. *Journal of Economic Psychology* 31(1): 114–30.
- Watt, R. (2002). Defending expected utility theory. *Journal of Economic Perspectives* 16(2): 227–9.
- Weber, B. J. and Chapman, G. B. (2005). Playing for peanuts: why is risk seeking more common for low-stakes gambles? *Organizational Behavior and Human Decision Processes* 97(1): 31–46.
- Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology* 39(5): 806–20.
- Wicker, F., Hamman, D., Hagen, A. S., Reed, J. L., and Wieche, J. A. (1995). Studies of loss aversion and perceived necessity. *Journal of Psychology* 129(1): 75–89.
- Williams, A. F. and Lund, A. K. (1986). Seat belt use laws and occupant crash protection in the United States. *American Journal of Public Health* 76(12): 1438–42.
- Wu, G. (1994). An empirical test of ordinal independence. *Journal of Risk and Uncertainty* 9(1): 39–60.
- Wu, G. and Gonzalez, R. (1996). Curvature of the probability weighting function. *Management Science* 42(12): 1676–90.
- Wu, G. and Markle, A. B. (2008). An empirical test of gain-loss separability in prospect theory. *Management Science* 54(7): 1322–35.
- Yaari, M. E. (1987). The dual theory of choice under risk. *Econometrica* 55(1): 95–115.
- Yaniv, G. (1999). Tax compliance and advance tax payments: a prospect theory analysis. *National Tax Journal* 52(4): 753–64.

- Yitzhaki, S. (1974). A note on income tax evasion: a theoretical analysis. *Journal of Public Economics* 3(2): 201–2.
- Zank, H. (2001). Cumulative prospect theory for parametric and multiattribute utilities. *Mathematics of Operations Research* 26(1): 67–81.
- Zhang, W. and Semmler, W. (2009). Prospect theory for stock markets: empirical evidence with time-series data. *Journal of Economic Behavior and Organization* 72(3): 835–49.
- Zhao, J. and Kling, C. L. (2001). A new explanation for the WTP/WTB disparity. *Economics Letters* 73(3): 293–300.

PART 2

Other-Regarding Preferences

INTRODUCTION TO PART 2

Neoclassical economic theory typically relies on the assumption of *self-regarding* behavior, i.e., that individuals derive utility solely from their own pecuniary or non-pecuniary payoffs. We may term such preferences as *self-regarding preferences* and an individual with these preferences as a *self-regarding individual*.¹ This view of human behavior is often at odds with the evidence, particularly when interaction in small groups is considered but also, often in large-group interactions.

The evidence indicates that humans often exhibit direct or indirect concern about others for several reasons. As Gintis (2009, p. 49) puts it “We conclude from behavioral game theory that one must treat individuals’ objectives as a matter of *fact*, not *logic*.” Individuals may care for others on moral grounds, or they may directly care about the pecuniary and non-pecuniary payoffs of others, or they may derive a *warm glow* from helping those who are less unfortunate, or they may wish to punish others in the group whose treatment of them or of others in the group does not conform to expectations, or they may wish to signal desirable human qualities such as altruism or kindness to themselves and to others. In all these cases, humans exhibit *other-regarding preferences* (sometimes referred to as *social preferences*).

Part 2 of the book studies the behavior of individuals with other-regarding preferences.

The model with purely self-regarding preferences is unable to explain a range of phenomena from many diverse areas such as social dilemma situations, collective action, contract theory, the structure of incentives, political economy, and the results of several experimental games. By contrast, many of these phenomena can be explained by a model, which recognizes that individuals have other-regarding preferences. On the other hand, many phenomena such as the behavior of competitive experimental market subjects seem to be in conformity with the model of self-regarding preferences (Roth et al., 1991).² However, in richer competitive environments, one still recovers a role for other-regarding preferences (Fehr et al., 1998). One challenge for theoretical research is to describe why the same individual exhibits other-regarding preferences in small groups, or in bilateral interactions but self-regarding preferences in anonymous, large-group, market situations; see Fehr and Schmidt (1999) for a successful example of such a behavioral model.

Most experiments reveal a mixture of self-regarding and other-regarding preferences, but the relative fractions of players who exhibit the two types of preferences vary. For instance, Fischbacher et al. (2001) show that 30–40% may have self-regarding preferences while the remaining may have other-regarding preferences. An important and interesting issue for theoretical and empirical research is to examine the implications of such heterogeneity in preferences. We shall

¹ The literature sometimes uses the term *selfish preferences* for *self-regarding preferences*. For most of the models that we consider, individuals try to maximize their objective functions, whether they have self-regarding or *social preferences*. For this reason, we find the term self-regarding to be more appropriate.

² Interestingly, in the Roth et al. (1991) experiments, bilateral behavior of subjects is consistent with models of social preferences. This dichotomy of behavior in multilateral and bilateral settings is not simply a figment of experimental data. Bolton and Ockenfels (2014) confirm these experimental findings from a field study based on eBay transactions. This also adds to a growing number of field studies that increasingly establish the external validity of lab experiments, a theme that we consider in more detail in Section 5.5.

show that the introduction of even a small fraction of individuals with other-regarding preferences in a model with self-regarding preferences can substantially alter the results and vice-versa.

Evidence for other-regarding preferences comes from many sources. Consider the following examples; the reader is invited to ponder if the model with purely self-regarding preferences could account for these results.³

1. Wages are usually not lowered in a recession as firms fear that this action may be construed by the workers to be *unfair* and have a negative effect on worker's *morale* (Bewley, 1999, 2005). The experimental findings of Fehr et al. (1993) and Fehr and Falk (1999) demonstrate that downward wage rigidity is caused by the belief on the part of employers that workers will reciprocate lower wages with lower levels of effort. For instance, in a double auction experiment, Fehr and Falk (1999) showed that employers accepted high wage offers in excess of the competitive level when the effort was endogenous. However, when the workers' effort was exogenously fixed, firms accepted the lowest wage offers. Thus, firms did not accept high wage offers purely for fairness considerations but rather because they expected effort reciprocity on the part of workers. In actual practice, firms may prefer to hire workers who have other-regarding preferences because they may be expected to reciprocate positive actions by the firms and other co-workers in the team.
2. Consider a two-stage public goods game, where individuals make voluntary contributions to a public good, followed by the possibility of non-contributors being punished by the contributors. The model with purely self-regarding preferences predicts *free-riding*, i.e., zero contributions and no punishment of non-contributors by the contributors because bygones are bygones. In contrast, the experimental results show that individuals contribute more when others contribute more in previous rounds (Fischbacher et al., 2001). In experiments in many Western societies, we also observe the punishment of non-cooperators by cooperators and close to the first best level of contributions (Fehr and Gächter, 2000).
3. In one-shot interactions, in a range of experimental games considered below, such as the ultimatum game, the dictator game, and the trust game, people pass on far higher monetary amounts to others than are predicted by models of self-regarding preferences.
4. Kahneman et al. (1986) asked people if they would view a hardware store's action to increase the price of shovels following a snowstorm as fair? Most people thought that such a price increase was unfair. However, they did not perceive an increase in price following an increase in the cost of inputs as unfair. In actual practice, many firms may go to great lengths to establish their credentials for fair behavior.
5. Bargaining often stalls (*bargaining impasse*) in a range of situations, such as civil litigation, labor unrest, domestic strife, and clashes among religious and ethnic groups. An important psychological cause of bargaining impasse is that typically, individuals confuse between what serves their self interest and what is fair (Babcock and Loewenstein, 1997).⁴

³ It is possible that self-regarding preferences in conjunction with other auxiliary assumptions may explain these examples. However, if these assumptions are empirically rejected, then such explanations are ad hoc.

⁴ We do not specially focus on bargaining issues in Part 2; these are considered in more detail in Part 4. Some of the reasons for bargaining impasse are as follows. First, confusion between self-interest and what is fair, reduces the set of feasible agreements. Second, if each party believes that its (misplaced) notion of fairness is shared by the other party, then even potentially kind intentions by the other party could be perceived as being unkind. Third, psychological evidence indicates that parties in a bargaining situation are unwilling to settle below what they (mistakenly) perceive to be a fair entitlement. The reader might also wish to read a nice discussion of bargaining issues related to the issue of norms and social preferences in Crawford (1997, Section 5.3).

6. Individuals make donations to charity, donate blood at Red Cross blood drives, and often volunteer time and money for prosocial actions. Such prosocial actions are often carried out for the benefit of non-kin members at a personal cost. The provision of monetary incentives for prosocial behavior may even be counterproductive (Titmuss, 1971; Bénabou and Tirole, 2003, 2006).
7. In one-shot interactions, such as in the trust game, pre-game *promises* and *threats* are predicted to be ineffective in altering behavior. Yet, such promises and threats are effective in generating more prosocial behavior (Ellingsen and Johannesson, 2004). Indeed, humans seem to value virtuous behavior such as honesty, fairness, and justice for its own sake; in neoclassical economics such behavior is instrumental *but not intrinsic*.
8. In one-shot interactions, when firms have a choice between a bonus contract and an incentive contract, the model with self-regarding preferences predicts the choice of an incentive contract. Bonus contracts should be ineffective in eliciting any effort from workers. Yet, we observe the predominant choice of bonus contracts in such a situation and much higher effort levels as well as profitability relative to incentive contracts (Fehr et al. 2007).
9. Military training programs typically focus on creating a separate identity—the identity of a soldier, in contrast to a civilian identity. A soldier's identity is associated with norms of behavior in particular situations that can be different from civilian norms of behavior in similar situations. These include the norms of *service before self*, courage, sacrifice, and fortitude under adversity. Similarly, firms try very hard to alter the identity of workers into *insiders* from *outsiders*. Insiders typically buy into the culture of the firm and subscribe to its objectives.

There has been a traditional aversion in economics to altering preferences/tastes in order to explain economic phenomena. This aversion may have been historically well-placed because the alternative could have been a proliferation of ad-hoc models with preferences tailored to explain every possible situation. However, research in behavioral economics has provided a solid foundation, based on extensive experimental evidence, to propose models of other-regarding preferences. This is not ad hoc, and has enormously increased the ability of economic models to explain economic phenomena. Conversely, one may argue that in the face of overwhelming evidence that many humans exhibit some form of other-regarding preferences, it is ad hoc to focus exclusive attention on self-regarding preferences. As Gintis (2009, p. 49) puts it: “Self-regarding agents are in common parlance called sociopaths.”

The argument that neoclassical economics does not rule out the role of other-regarding preferences is correct. Indeed, an older generation of economic models considered ideas such as *keeping up with the Joneses*, and *snob consumption* (Veblen, 1899; Duesenberry, 1949). However, while plausible, these models were not empirically well-founded and would not be able to explain the results from modern experimental games on fairness and reciprocity. In contrast, the current approach in behavioral economics offers empirically well-founded and parsimonious models of other-regarding preferences.

Let us consider some terminology that is used in models of other-regarding preferences. This terminology is by no means standard, partly because its origin straddles several disciplines such as economics, biology, anthropology, and sociology. Hence, we simply clarify the sense in which various terms are used in this book, without claiming that such usage is unique or the best possible.

Preferences are *altruistic* (respectively, *spiteful*) if one's utility is increasing (respectively, decreasing) in the consumption or well-being of others. Altruism is *unconditional kindness*, while spite is *unconditional unkindness*, irrespective of the behavior of other players. In practice, however,

pure altruism is unable to explain a wide range of behaviors. For instance, it does not explain why many individuals are *conditional cooperators*, i.e., they cooperate only when others cooperate (as in successive rounds of a public goods game).⁵ An *inequity averse* individual not only derives utility from own material consumption but also derives disutility from differences in own material consumption relative to the consumption of others. Such an individual dislikes relatively greater consumption (altruism) or relatively lower consumption (envy).

Reciprocal behavior can also be shown to arise in traditional economic analysis with self-regarding preferences. In classical game theory, in infinitely repeated games (or finitely repeated games with some uncertainty), cooperative behavior can be sustained in equilibrium. Such behavior is said to be *instrumentally reciprocal* because it is motivated purely by the prospect of a trade-off between short-run and long-run material gains.⁶ By contrast, in a finitely repeated game, say, in a principal–agent relation, under self-regarding preferences, the outcome is that the worker shirks and output is low. However, under other-regarding preferences, theory and empirics support superior cooperative outcomes (Brown et al., 2004, 2012). Thus, the evidence violates instrumental reciprocity here, and in many other classes of games.

Reciprocity that takes the form of responding to kindness and unkindness of others with like behavior, even in the absence of long-term gains and under complete certainty, may be termed as *intrinsic reciprocity*. People who exhibit such behavior are often described as *conditional cooperators*. Many important behavioral explanations on human sociality rely on *intrinsic reciprocity*.

Altruism among kin members in humans is well documented (Daly and Wilson, 1988). However, among humans, one observes cooperation among non-kin group members and impressive division of labour at a level that is unprecedented in the animal kingdom. Cooperation is much lower in primate societies with whom we share a common ancestor. Cooperation is well documented among ants, bees, and wild dogs, but this seems limited to genetic relations. Furthermore, humans are possibly unique among all animals to consider the fairness implications of their actions for kin and non-kin members.

An important component of human sociality is that group members often engage in costly punishment (respectively, costly rewards) of actions by third parties that are inimical (respectively, favorable) to some social goal or group norm. The punishers may not be directly affected by the actions of the third parties, yet they chose to punish/reward. Such costly punishment also takes place in finitely repeated games with no uncertainty, and no kin relation among group members. This form of reciprocity may be termed as *indirect reciprocity* and plays an important role in upholding group norms and social goals.

Strong reciprocity is a more encompassing concept that includes conditional cooperation and indirect reciprocity. Bowles and Gintis (2003) provide a working definition of strong reciprocity:

A strong reciprocator comes to a new social situation with a predisposition to cooperate, is predisposed to respond to cooperative behavior on the part of others by maintaining or increasing his level

⁵ Axelrod and Hamilton (1981) conducted the classic experiments in reciprocity in the context of a sequence of repeated prisoner's dilemma games where there is some probability of the continuation of the game to the next period; see, also Axelrod (1984). They found evidence for the superiority of the tit-for-tat strategy when subjects are randomly matched in each round of the game. In this strategy, subjects initially start by cooperating and then reciprocate in subsequent rounds, i.e., based on their experience in previous rounds, play defect in response to defect and play cooperate in response to cooperate. Such a strategy did best among the ones that were considered.

⁶ A related concept from biology is that of *reciprocal altruism*. This involves a tendency to reward or punish others in response to their kind or unkind behavior, when such net rewards are in the long-term interests of the animal/individual. See, for instance, Axelrod and Hamilton (1981) and Trivers (1971).

of cooperation, and responds to free-riding behavior on the part of others by retaliating against the offenders, even at a cost to himself, and even when he cannot reasonably expect future personal gains from such retaliation. The strong reciprocator is thus both a conditionally altruistic cooperator and a conditionally altruistic punisher whose actions benefit other group members at a personal cost.

In what follows, when there is no confusion, we shall mostly use the generic term reciprocity to denote various forms of reciprocity. How can we distinguish among various reciprocity concepts such as instrumental reciprocity, intrinsic reciprocity, and strong reciprocity? The solution is provided by carefully constructed experiments that are able to conduct sharp tests of the competing hypotheses and disentangle the various confounding factors. We shall review a range of such experimental studies below.

There are four chapters in Part 2. We give a very brief introduction here; a more detailed introduction follows in each chapter.

Chapter 5 considers the empirical evidence on other-regarding preferences using a range of experimental games. These games include the *ultimatum* and *dictator games*, the *trust* and *gift exchange games*, and the *public goods game* with and without punishments. In this chapter, we also consider the external validity of lab experiments and the issue of the student subject pool (see also the introduction to the book). We show that lab evidence is representative of real-world human behavior, and if anything, the student subject pool delivers results in favor of the self-regarding model.

Chapter 6 considers models of *social preferences*, *type-based reciprocity*, *intentions-based reciprocity*, and other *hybrid models*.⁷ We not only consider the leading theoretical models but also use empirical evidence to evaluate the relative success of the models. We show how these models can be used to explain the empirical evidence in Chapter 5. We give applications to behavioral political economy, and develop appropriate stochastic dominance concepts in the presence of other-regarding preferences. We also draw the implications of other-regarding preferences for a Walrasian competitive equilibrium as well as for the fundamental welfare theorems. A wide range of further applications is given in each of the successive chapters.

Chapter 7 considers two new topics, *human morality* and *social identity*, both of which are not traditionally considered in courses in neoclassical economics, which is amoral. We first give the evidence on human morality, such as the aversion to lying, aversion to breaking promises, and wishing to uphold one's integrity. This is followed by an application to the hold-up problem in industrial organization. We then consider the evidence on social identity and explore applications on incentives in organizations.

Chapter 8 is the longest chapter in Part 2. It considers issues of moral hazard and asymmetric information in the presence of other-regarding preferences, and the literature on complete and incomplete short- and long-term contracts. Unlike many treatments of contract theory, we also closely consider the evidence on alternative contractual forms as well as the theoretical models and the relevant empirical evidence on *extrinsic* and *intrinsic motivation*.

⁷ The formal development of intentions-based reciprocity requires the use of *psychological game theory*, so it is postponed to Part 4 of the book.

CHAPTER 5

The Evidence on Human Sociality

5.1 Introduction

In this chapter, we review the evidence from several kinds of experimental games. These include the *ultimatum game*, the *dictator game*, the *gift exchange game*, and *public goods games* with and without punishments. In small group interaction, particularly in bilateral interaction, these experiments reject the notion that humans have purely self-regarding preferences. The behavior of a majority of people is consistent with other-regarding preferences and conditional reciprocity, although a significant minority exhibits purely self-regarding preferences. An important finding, from these experiments, is that the introduction of a small fraction of individuals with other-regarding preferences in a population of self-regarding individuals can result in a large change in the outcomes (Fischbacher et al., 2001; Fehr and Fischbacher, 2005; Fischbacher and Gächter, 2010).

Section 5.2.1 considers the ultimatum game, one of the most widely replicated experimental games in economics. In this two-player game, originally introduced by Güth et al. (1982), a *proposer* proposes to share a part of his endowment with a *responder*. The only action of the responder is to accept or reject. If accepted, the split proposed by the proposer is implemented, but if rejected, both get nothing. The prediction under self-regarding preferences is that the responder will be offered the smallest indivisible part of the endowment by the proposer. By contrast, the results are consistent with other-regarding preferences. Mean offers are 30–40% of the endowment and the median offers are 40–50% of the endowment.

These results continue to hold when the stakes are raised, say, up to an equivalent of about 62 hours of work (Slonim and Roth, 1998). But, for extremely large stakes, equivalent to 1600 hours of work, the responders' rejection rates fall but the median offer by the proposer is still about 20% of the endowment (Andersen et al., 2011). However, these stakes are too large relative to those that are commonly observed in real life among individuals. The evidence also indicates that there is some interaction between experience and the choices made in the ultimatum game but the self-regarding outcome is still not achieved. It would seem, therefore, that for most people other-regarding preferences are very hard to switch-off and most of us may be intrinsically disposed to exhibit at least some altruism and reciprocity, even when the stakes are reasonably high. We postpone a formal definition of *reciprocity* to Chapter 13, where we introduce the machinery of psychological game theory. For the moment, humans exhibit reciprocity if they respond to the kind actions of others with kindness, and unkind actions with unkindness.

It has been alleged by some critics that subjects play prosocially in ultimatum games because they mistakenly apply repeated game heuristics to an essentially static but multi-stage game (Binmore, 1998). However, the evidence from ultimatum games does not support this view (Fehr and Fischbacher, 2002). There is wide variation in the offers made by proposers and in the rejection rates of responders in different societies (Henrich et al., 2001). This cautions us against generalizing results too much from one culture to another. Interestingly, contrary to what students of economics may have been led to believe, offers and rejection rates are higher in societies where market exchange is more prevalent, although the direction of causality is difficult to establish with the available evidence. Indeed, market exchange requires a high degree of interpersonal trust especially because contracts may often be verbal and informal.

In Section 5.2.2, we consider the dictator game, which is a variant of the ultimatum game in which the responder has no actions, so the proposer's proposed split is always implemented. Nevertheless, proposers still make a mean offer of about 15–20% of the endowment in this case (Forsythe et al., 1994). This establishes two results. (i) Since the mean offer is lower relative to the ultimatum game, high offers in the ultimatum game arise partly for strategic reasons, i.e., to ensure acceptance of the offer by conditionally cooperative responders. (ii) Since the mean offers are strictly positive in the dictator game, proposers exhibit altruism.

Recent scholarship gives a more nuanced view of the results from dictator games. A distinction has been drawn between *genuine altruists* and *reluctant altruists* (Dana et al., 2006). A reluctant altruist is someone who, for instance, may give money to a beggar should he unexpectedly run into one, but prefers to be on the other side of the road if he spots a beggar from the distance. There are many reasons why people seem to behave more altruistically in dictator games when they perceive they are being watched (Haley and Fessler, 2005). For instance, people may have beliefs about the expectations of others about their actions, and they may dislike falling short of these expectations (*guilt*). Emotions such as guilt may directly enter one's utility function. Formal modeling of these issues requires the machinery of *psychological game theory* that is considered in Part 4 (Geanakoplos et al., 1989; Battigalli and Dufwenberg, 2009). Another potential explanation is that *identity* and *self-esteem* may be important determinants of the desire to appear fair to oneself and to others (Bénabou and Tirole, 2011).

In Section 5.2.3, we consider dictator game experiments that allow us to draw inferences about the rationality of subjects (Andreoni and Miller, 2002). Subjects in these experiments reveal that they have other-regarding preferences because there are strictly positive offers in the dictator game, yet the same subjects satisfy the *generalized axiom of revealed preference* (GARP). Hence, there is nothing irrational about other-regarding preferences.

Section 5.3 considers two related games, the *gift exchange game* and the *trust game*. We first consider the gift exchange game (Fehr et al., 1993; Fehr et al., 1997) in Section 5.3.1. A firm (player 1) makes a contract offer (w_b, e_n) to a worker (player 2) that specifies a binding wage level, w_b , and a non-binding effort level, e_n . The contract may offer *positive rents* to the worker, i.e., a wage (net of effort costs) that is greater than the outside option. If the contract is accepted, the worker chooses an actual effort level, e_a . Since the wage is binding, if workers have self-regarding preferences, they should shirk, i.e., $e_a = 0$. Anticipating this, firms should offer $w_b = 0$, so the rents of workers should be zero. By contrast, the experimental findings are that (i) workers are offered positive rents, and (ii) the rents and actual effort levels of workers are highly positively correlated. These results provide important confirmatory evidence for other-regarding preferences. In experiments that allow firms to punish shirking workers when $e_a < e_n$, results show that shirking is significantly reduced. If the firm were self-regarding, punishments should be ineffective because bygones are bygones in classical game theory.

The results of experiments in the gift exchange game have been extended in several directions. The effects of a wage decrease may have a more pronounced effect on effort as compared to wage increases (Kube et al., 2006). Although the sample sizes are small, this result is consistent with the findings of Bewley (1999) that firms are reluctant to cut wages in a recession. As in the ultimatum game, there are cultural differences in the degree of sociality in the gift exchange game (Charness et al., 2004). Fehr et al. (1998) embed a gift exchange game within competitive markets where quality is endogenous. They find substantial departures from the competitive equilibrium outcome. Thus, current competitive equilibrium experiments (see Part 4 of the book) that justify the equilibrium outcome may not be rich enough to deal with all the relevant aspects. The idea that sociality may survive and affect outcomes in competitive markets is one of the most important results in Part 2 of the book.

When the gift exchange game is played with a firm and two or more workers, then additional insights emerge, depending on the relative pay and effort levels of the workers. It turns out that the degree of sociality in gift exchange games is influenced by the degree of *social comparisons* in the workplace (Gächter et al., 2012; Abeler et al., 2010; Cohn et al., 2011). The degree of sociality also depends on whether the workers' rent is monetary or non-monetary (Kube et al., 2012). Recent work explores the *duration effects* of the gift exchange relation. Gneezy and List (2006) conclude that the effects are short-lived. However, Bellemare and Shearer (2007) find that the short-lived effects only apply to workers who are hired on the spot market. By contrast, workers with longer tenure exhibit stronger reciprocity and long-lived effects of gift exchange. The predictions of the gift exchange game are confirmed by survey evidence (Dohmen et al., 2009).

In the gift exchange game, the total surplus is controlled by player 2, the worker, who exerts mutually beneficial effort. In Section 5.3.2, we consider the *trust game* in which player 1 determines the surplus in the relation (Berg et al., 1995). Both players are given an identical initial endowment, $e > 0$, by the experimenter. Player 1, the *trustor*, chooses an investment, $i : 0 \leq i \leq e$. The experimenter then multiplies this investment by a number $m \geq 1$ and passes it on to player 2, the *trustee*. The trustee voluntarily returns an amount of money $r : 0 \leq r \leq mi + e$ to the trustor. The numbers i (investment) and r (return) are, respectively, taken as measures of *trust* and *trustworthiness*. If players are self-regarding, then the subgame perfect outcome of the trust game is $i = r = 0$. By contrast, the meta-study of Johnson and Mislin (2011) shows that, pooling across all regions, trustors invest 50.2% of their endowment and trustees return 37.2% of the feasible amount. But there is significant regional heterogeneity (e.g., North Americans send back significantly more than Africans), older people are more trusting, students are less trusting, stakes do not appreciably affect the results, and the results depend on the value m . Overall, these results confirm the presence of other-regarding preferences and intrinsic reciprocity.

When subjects have an opportunity to examine the history of the other player in trust games, most players choose to do so. Interestingly, subjects may use the past levels of "trust" exhibited by other players to also infer their "trustworthiness," suggesting that these two human virtues may be correlated (Charness et al., 2011). In an interesting experiment, Eckel and Petrie (2011) offer subjects a chance to view a photograph of their opponent at a price. They find that there is a stable demand curve for these photographs. Conditional on having observed the opponent's picture, players behave more prosocially. White trustors discriminate between white and black trustees, investing more if the trustee is white. Black trustors do not discriminate. There are hardly any gender differences in prosocial behavior. On the whole, there appears to be an information value in a face.

Section 5.4 considers *public goods games*. Section 5.4.1 considers the basic model *without punishment* and Section 5.4.2 *with punishments*. In these games, a group of individuals simultaneously

choose their level of contributions to a public good that is non-rival in consumption and non-excludable. For each extra unit of contributions to the public good, an individual receives some private per unit benefit from the public good that is less than the private marginal cost of the contribution. If all contributors have self-regarding preferences, then the unique prediction of the model is that all players will *free-ride*, i.e., each player will contribute zero. In the model with punishments, there is a second stage in which, once the contributions are revealed, players can engage in costly punishments of other players. However, under self-regarding preferences, in the second stage, bygones are bygones, so players should not punish others. Without the threat of punishments, the predicted results with and without punishments are identical, i.e., everyone free-rides.

The evidence from public goods games, in Section 5.4.3, rejects the predictions of the self-regarding model, irrespective of whether punishments are available or not (Fehr and Gächter, 2000). In the absence of punishments, contributions typically start off from a high level and then gradually decay over successive rounds of the experiment. Self-regarding preferences cannot explain the pattern of decay, nor the reason why contributions start off from a high level, once the game is restarted. When punishments are allowed, then the outcome is close to the first best. Contributors punish non-contributors when the game is played in most Western countries where the *rule of law* and the *norm of civic cooperation* are high (Herrmann et al., 2008).

The evidence on various aspects of the public goods game is rich and extensive. The negative net effects of punishments that are sometimes found when the experiment is repeated over a few rounds, typically 10, are reversed when the experiment is repeated a large number of times, say, 50 rounds (Gächter et al., 2008). This informs our view of why punishments may have evolved; they are efficiency enhancing in enforcing social norms, provided the game is played over sufficiently many rounds. A combination of rewards and punishments elicits even greater cooperation as compared to the use of these instruments on their own (Sefton et al., 2007). Non-monetary punishments such as ostracization may also be effective in securing cooperation (Masclét et al., 2003; Cinyabuguma et al., 2005). Endogenous group formation in which players or the experimenter choose group members based on their past contributions influences contributions (Chaudhuri and Paichayontvijit, 2006; Gächter and Thöni, 2005; Page et al., 2005). Face to face communication enhances cooperation, but it is predicted to be ineffective under self-regarding preferences (Bochet et al., 2006). Publicly given, intergenerational, advice achieves nearly 90% of the socially optimal contributions (Chaudhuri et al., 2006).

In Section 5.4.4, we consider an explanation of the pattern of contributions in public goods games based on heterogeneous other-regarding preferences (Fischbacher et al., 2001; Fehr and Fischbacher, 2005; Fischbacher and Gächter, 2010). This model highlights two fundamental features of many social dilemma situations. (i) There is a mix of people with self-regarding and other-regarding preferences. (ii) There are profound implications of having even a small fraction of subjects with one type of preferences in the presence of the other type. Suppose that in a public goods game, a majority of the subjects have other-regarding preferences and we begin with initially high levels of cooperation. Then the presence of a minority of subjects with self-regarding preferences may drive contributions to low levels. The reason is that over successive rounds of the game, as contributors discover that some of the other players are free-riders, negative reciprocity kicks in, and they respond with successively lower contributions. Thus, we cannot conclude, by observing a lack of cooperation, that most players have self-regarding preferences.

In Section 5.4.5, we consider games in which two rival groups expend costly resources to win a contest and the prize is shared equally among the members of the winning group (Abbink et al., 2010). Thus, the prize has the nature of a public good. Examples include countries going to war,

trade union strikes, and firms lobbying for a government contract. The winning probability, for any group, depends on the relative resources expended. If both groups increase their expenditures in the same proportion, their payoffs fall, but winning probabilities are unaltered. Thus, if the groups can coordinate, they can simultaneously, and proportionally, reduce their expenditures without altering winning probabilities, but increasing their payoffs; this is the sense in which the expenditures are unproductive. Each group, by incurring such expenditures, reduces the winning probability of the other group, a negative externality that it does not take into account. Empirical evidence shows that the unproductive expenditures are much higher than the Nash predictions.

Section 5.4.7 considers cultural differences in the nature of punishments in public goods games (Herrmann et al., 2008). In societies where the *rule of law* and the *norm of civic cooperation* is strong, one observes *prosocial punishments*, i.e., cooperators punish non-cooperators and this enhances prosociality. However, in other societies where one observes a weaker rule of law and a lower norm of civic cooperation, one observes *antisocial punishments*, i.e., non-cooperators may punish the cooperators in revenge for having been punished in the past. This has adverse effects on the level of contributions in public goods games. One reason for the net negative effect on outcomes in the presence of antisocial punishments is that experiments may overstate the extent of punishment opportunities that we observe in society. For instance, societies typically engage in collective punishments through a legal code. It is quite likely that private punishments may be harsher relative to those prescribed by the legal code, however, experiments typically do not distinguish between the two. When one allows for credible punishments over a reasonably high number of rounds, then the net effect of punishments is positive (Gächter et al., 2008).

Section 5.4.7 gives some observations on the evolutionary sources of reciprocity. Part 5 of the book considers this material in far greater detail.

Section 5.5 considers the issue of the *external validity* of lab experiments as well as other possible *biases* in such experiments. The introduction to the book also gives a detailed consideration of these issues, so the treatment here minimizes the degree of overlap between the two. Several concerns have been raised about the validity of lab experiments (Levitt and List, 2007). Section 5.5.1 considers the criticism that the typical experiment is conducted with students, and more prosocial students may have been self-selected into the experiments. Being watched by the experimenter may bias the behavior of subjects in experiments (*experimental scrutiny effects*). Furthermore, students may deliberately appear more prosocial, anticipating that this is what the experimenter desires (*experimenter demand effects*). An increasingly robust finding in the emerging experimental literature is that the student population is less prosocial as compared to the general population (Falk et al., 2013; Carpenter and Seki, 2011; Fréchette, 2015). If anything, this biases the results in favor of the textbook self-regarding model. More prosocial students do not self-select themselves into experiments (Cleave et al., 2012). Experimental scrutiny does not necessarily lead to more prosocial behavior (Barmettler et al., 2012).

In any case, even if experimental scrutiny were important, as is suggested in several experimental papers, then it too makes a powerful case for the importance of other-regarding preferences. If the presence of a stranger (the experimenter) in the room switches on other-regarding preferences, then it is quite plausible that most human activity, which is carried out in the direct presence or indirect presence¹ of others is similarly affected. It would appear that the negative implications of experimental scrutiny may have been exaggerated, even if such scrutiny were important.

¹ For instance, the party that is unfairly treated can disseminate information about an unfair person to others, including legal institutions if the behavior is illegal.

A rapidly expanding body of empirical literature, that we consider in Section 5.5.2, supports the external validity of lab experiments (Camerer, 2015). Japanese fishermen who operate in a collective manner in the field (sharing of costs and the fish catch) are also more prosocial when they play the public goods game in the lab (Carpenter and Seki, 2011). The conditional reciprocity of fishermen estimated in the lab is a significant predictor of their field productivity. Subjects who reveal more self-regarding preferences in the lab are also less likely to increase their effort in the real world, following a wage increase (Cohn et al., 2012).

Shrimpers and fishermen in northeastern Brazil who behave more selfishly in public goods games also behave more opportunistically in the field (Fehr and Leibbrandt, 2011). More prosocial behavior in the lab is correlated with better market performance (Leibbrandt, 2012). In Ethiopia, subjects who cooperate more in the public goods game also behave in a less opportunistic manner in their real-world activity that involves forest management (Rustagi et al., 2010). A greater degree of trustworthiness in trust games predicts a greater likelihood of repaying loans in a microcredit program in Peru (Karlan, 2005). Chicago MBA students who exhibit more trustworthiness in the lab are also more likely to make charitable donations (Baran et al., 2010). Similar results are true for Swiss students (Benz and Meier, 2008).

In a game played between two players, if a third party who has no material stakes in the game, is willing to engage in costly punishments and rewards, then it exhibits *indirect reciprocity* (Gintis, 2009). In the final part of the chapter, in Section 5.6, we consider the extensive evidence in favor of indirect reciprocity (Seinen and Schram, 2001; Engelmann and Fischbacher, 2009; Fehr and Fischbacher, 2004).

5.2 Ultimatum and dictator games

5.2.1 *The ultimatum game*

The *ultimatum game*, introduced by Güth et al. (1982), is played between two players, a *proposer* and a *responder*. The *proposer*, is given an integer amount $\$x$ by the experimenter. Suppose that the smallest indivisible amount of money is one cent. The proposer offers a share $s \in \{0, 0.01, 0.02, \dots, x.00\}$ to the *responder*. If the responder accepts the offer, then the proposer's initial offer is implemented and the payoffs of the proposer and the responder are, respectively, $y_P = x - s$, and $y_R = s$. If the responder rejects the offer, the respective dollar payoffs are $y_P = 0$, and $y_R = 0$.

If both players have self-regarding preferences, then in the subgame perfect outcome the dollar amount of payoffs are $y_P^* = x - 0.01$, $y_R^* = 0.01$. If the amount $\$x$ was infinitely divisible, then the subgame perfect outcome is $y_P^* = x$, $y_R^* = 0$.² The experimental results do not support this prediction. The results have been summarized in several places.³ The gist of the experimental results is as follows. (1) The mean offer is $s = 0.3x$ to $0.4x$. (2) The median offer is $s = 0.4x$ to $0.5x$. (3) There are rarely any unfair offers ($s = 0$ or $0.1x$) or over-fair offers ($s > 0.5x$). (4) Low offers are often rejected.

² If the amount $\$x$ were infinitely divisible with α being the smallest equal part, then the subgame perfect outcome is $y_P^* = \lim_{\alpha \rightarrow 0} (x - \alpha) = x$, $y_R^* = \lim_{\alpha \rightarrow 0} \alpha = 0$.

³ See for instance, Camerer (2003), Camerer and Thaler (1995), Thaler (1988), Güth and Tietz (1990), Roth et al. (1991).

These results continue to hold in several studies conducted in high stakes games. For instance, Cameron (1999) conducted high stakes experiments in Indonesia in which x ranged from 5,000 to 200,000 rupiahs (about three months' wages). The mean offer was $s = 0.4x$ to $0.45x$. Respondents were more likely to accept offers when stakes were high while the behavior of the proposers hardly changed.⁴

Slonim and Roth (1998) consider varying experiments with varying stakes in the Slovak Republic in which x equals 2.5, 12.5, and 62.5 hours of wages respectively. The rejection rates in the low and the high stakes conditions are similar in the first round and fall slightly in the high stakes condition by round 10, the last round of the experiment. This indicates that there is some interaction of experience and stakes in the rejection rates of offers. However, the mean and median offers are very similar in the low and the high stakes conditions. Exit interviews show that the main motive for the rejection of low but positive offers is that responders find them to be *unfair*.

Carpenter et al. (2005) find that for moderately high stakes, a given proportional increase in stakes leads to a much less proportionate fall in the offers made by the proposer.⁵ For instance, a ten-fold increase in stakes from \$10 to \$100 leads to a halving of the median offer from 40% to 20%, and to an even lower proportionate fall in the average offer from 33% to 25%. Andersen et al. (2011) raise the stakes beyond the levels of most previous studies. Their study, carried out in a northeastern Indian village, allowed stakes to vary between 20 rupees to 20,000 rupees (1.6 hours of work to 1600 hours of work).⁶ They find that although the median offer by the proposer is to give 20% of the share to the respondent, the rejection rate falls with the increase in the stake.⁷ As the authors correctly point out, these results do not invalidate the practical applicability of earlier results because one does not typically encounter real-life situations with such high stakes and the bulk of market transactions are low stake transactions.

The ultimatum game has been mostly played with university students. In an interesting study, Henrich et al. (2001) explore cross-cultural variations in the ultimatum game. In this study, evidence was gathered from 15 small-scale societies across 4 continents and 12 countries. The predictions of the ultimatum game with self-regarding preferences are typically not borne out in any of these societies. The mean offers were quite variable (26–58%) relative to experiments done previously with university students where the mean offers are 43–48%. In some societies, even low offers were rarely rejected, while in others, even extremely generous offers were rejected because in such societies, the rejection of a gift is considered tantamount to the rejection of subordination.

The two main factors that enhanced sociality in the results of Henrich et al. were the following. (1) *Market integration* in the community, i.e., the predominance of buying/selling and working for a wage. (2) *Cooperation in production*, i.e., whether production is carried out on an individual basis or it is cooperative, e.g., in a team environment. These two factors enhanced sociality in the sense of higher offers by proposers and accounted for 66% of the variation in the outcomes in the ultimatum game. On average, people from the least market-integrated communities offered about one-quarter of the stake. By contrast, those from the most integrated communities, such as rural Missouri, offered roughly half. Thus, *markets may foster sociality*, an idea that might seem surprising to some students of mainstream economics because this is hardly ever highlighted.

⁴ Similar findings are reported in Fehr et al. (2014) and Hoffman et al. (1996).

⁵ See also, List and Cherry (2008).

⁶ By way of comparison, in the study by Slonim and Roth (1998), the highest stake level is 62.5 hours of wages.

⁷ In one condition, they find that at the highest level of the stake the rejection rate falls to 0%.

The results of the ultimatum game are sensitive to how much proposers feel that they are entitled to their endowments—a sort of *entitlement norm*. A feeling of entitlement may be created by simply announcing an entitlement (Hoffman and Spitzer, 1982; Cherry et al., 2002), or by having a contest to choose proposers to play a high stakes ultimatum game (List and Cherry, 2000), or by having the proposer earn the endowment (Parrett, 2006). The greater is the entitlement offered to the proposers, the lower is their offer to the responder. Harrison and McCabe (1996) find that the outcomes in an ultimatum game depend on the history of offers and rejections in previous rounds of the game, which might serve to establish a *fairness norm*.

Binmore et al. (1995) and Roth and Erev (1995) have used learning models to explain the results of the ultimatum game in terms of purely self-regarding preferences. The basic argument is that there is an asymmetry between the learning propensities of proposers and responders when the proposers make low offers. The payoff loss to a responder from rejecting a low offer is very low but for a proposer it is very high. Hence, proposers quickly learn not to make very low offers and responders learn very slowly not to reject very low offers. Over time, which can involve thousands of iterations, the behavior with self-regarding preferences can be shown to converge close to the subgame perfect outcome ($y_P^* = x$, $y_R^* = 0$). There are two problems with this approach. First, there is no guarantee that convergence takes place in any sensible amount of time. Second, the structure of the ultimatum game is incredibly simple and it seems simply implausible that subjects are making systematic mistakes of this sort.

Binmore (1998) explains the results of the ultimatum game in terms of social norms. The idea is that in the real world, subjects in experiments may be used to playing repeated games. Hence, in the laboratory, they may mistakenly apply repeated game heuristics to the static ultimatum game and cooperate even when the subgame perfect outcome requires them not to cooperate. It is likely that individuals, in their role as subjects in experiments, bring with them real-world norms into the lab (see, for instance, the variation in offers in different cultures in the study by Henrich et al., 2001). However, whether subjects mistake a static context (the ultimatum game) for a repeated game is an empirical question. And the empirical evidence rejects this view. The same subjects in experiments behave very differently in the lab when they play static games and when they play repeated games where reputation building is important.⁸ Furthermore, neuroeconomic evidence does not support the view that subjects make mistakes in these simple games.⁹

Fehr and Fischbacher (2003) conducted experiments with the ultimatum game in which 10 proposers and 10 responders met each other over successive rounds. No players are paired more than once. In the *reputation condition*, the past behavior of the responders is made public, while in the *no-reputation condition* it is private information to the responder. In the reputation condition, responders have an opportunity to develop a reputation for being tough by rejecting low offers, which is costly in the short run. The gain from such an action is the possibility that future proposers may make high offers to them. The average acceptance thresholds were significantly higher in the reputation condition relative to the no-reputation condition.

Some related issues in ultimatum games are also considered in the section on bargaining games in Chapter 12.

⁸ See, for instance, Andreoni and Miller (1993), Engelmann and Fischbacher (2009), Fehr and Fischbacher (2003), and Gächter and Falk (2002).

⁹ A region of the brain, the dorsal striatum, associated with the reward circuits of the brain is activated when subjects punish behavior that is perceived to be unfair; see, for instance, de Quervain et al. (2004). Knoch et al. (2006) show that an impairment of the prefrontal cortex that is associated with rational thinking reduces the ability of subjects to reject unfair offers.

5.2.2 The dictator game

It is useful to compare the predictions of the ultimatum game with the *dictator game* for additional insights. The dictator game is similar to the ultimatum game except that the responder (sometimes referred to as the *receiver*) is passive and has no actions. Thus, the dictator's offer is always implemented. If all individuals had purely self-regarding preferences, then the prediction is that the dictator will keep all the money and offer none to the responder. On the other hand, a strictly positive, but relatively lower, offer in the dictator game would indicate the following. (1) Proposers are inequity averse, or altruistic.¹⁰ (2) Higher offers in the ultimatum game are made partly for strategic concerns such as fear of rejection because responders may have inequity averse preferences.

Forsythe et al. (1994) find that proposers' offers are higher in ultimatum games relative to dictator games. They find that the average offers made by the proposers in dictator games are between 10–25%. While these offers are lower, on average, relative to those found in ultimatum games, the offers are still significantly higher than zero. The results of this study have been replicated many times. The interested reader can consult Camerer (2003) for a useful table of results (see p. 57). The mean transfer, across studies, made by the proposer in dictator games is roughly 15–20% of the endowment.

List (2007) broadens the action set of the proposer in dictator games to allow for both positive and negative transfers. This is achieved by first giving an endowment to the proposer and the respondent. Negative transfers mean that the proposer is allowed to take money from the respondent's endowment. The main finding is that this reduces the average and median offers made by the proposer, yet the equilibrium outcome under self-regarding preferences (i.e., the proposer makes a offer of zero) does not arise. Levitt and List (2007) explain these findings in terms of the attempts by subjects in experiments to find *contextual clues* in the game in order to apply appropriate norms. For instance, the expansion of the action set of the proposer in dictator games to negative transfers, reduces the scope of what might be construed as a self-regarding offer.

Haley and Fessler (2005) find that when dictators make their proposal where they perceive that they are being watched, say, merely by having pictures of eyes in the room, then they tended to make more generous offers. Similar results are found in a related context, where a picture of eyes elicited greater voluntary contributions for a drink, relative to a picture of flowers (Bateson et al., 2006). However, in games with minimal strategic elements, such as in the trust game, these result do not survive. No effect of the eyes cue was found for either prosocial individuals or self-regarding individuals in the experiments of Fehr and Schneider (2010).

In general, there is wide variation in the offers made in dictator games that are quite sensitive to simple framing effects. These results point to the high degree of sensitivity of the results under dictator games to framing effects and experimental protocol relative to those of, say, the ultimatum games. Despite its popularity, the dictator game might not be a particularly good game to test alternative theories that require even a modicum of strategic interaction (Fehr and Schmidt, 2006).

Dana et al. (2006) consider a variant of the dictator game in which dictators were given an initial amount of \$10. In the first experiment, once the dictators had made their choices, but before these choices were communicated to receivers, dictators were given the option of exiting the game

¹⁰ Selten and Ockenfels (1998) consider a modified dictator game, called the *solidarity game*. In this game, players had a ex-ante chance to be proposers and responders and outcomes were risky. They find it difficult to explain the behavior of a set of their subjects (36% of the total) using altruism alone.

with \$9. If dictators wished to exit the game in this manner, responders were not informed of the dictator's offer, or the game they were playing. Despite the fact that dictators could get a higher payoff by playing the dictator game than exiting in this manner, a significant percentage, 28%, choose the exit option.

In a second experiment, in the private condition, receivers did not know if the money offered to them came from the dictator or the experimenter, thus, the dictator's intentions are not revealed. In this case, 1 out of 24 dictators exited. However, when the first experiment was replicated, 9 out of 24 dictators exited. These results show that at least for a significant minority of the dictator subjects, the offers of proposers may reflect their desire not to appear self-regarding to the receivers.

Dana et al. (2007) introduce a degree of opacity in a variant of the usual dictator game. In the *baseline treatment* with the usual and fully transparent dictator game, the dictator can choose one of two actions, A or B with respective payoffs (6,1) and (5,5) (the first payoff is the dictator's and the second is the responder's). Thus, by accepting a small reduction in his own payoff by a unit, the dictator can enhance the payoff of the responder by four units. In the *hidden payoff treatment*, the payoff of the responder is uncertain, so the respective payoffs from the actions A and B are shown to be (6,?) and (5,?). All subjects were told that the payoffs from A and B are equally likely to be either (i) (6,1) and (5,5), or (ii) (6,5) and (5,1). In each case, the proposer can, costlessly and voluntarily, choose to reveal the payoffs by clicking a button on a computer screen but this information could not be seen by the responder. Clearly, the dictator is more comfortable with his altruism by choosing B in case (i) and A in case (ii).¹¹

If inequity aversion is the dominant motive, then the proportion of dictators who "give" in the baseline treatment, i.e., choose B, should be the same as the proportion of dictators who click to reveal the payoffs in the hidden payoff treatment. A high proportion, 74% (14 out of 19), of the dictators in the baseline treatment chose the more generous option, B. This could be either because they are inequity averse or that they did not wish to appear non-altruistic. However, in the hidden payoff treatment, 56% of the proposers did not choose to click the button to reveal the payoffs. The difference in the proportion of dictators who choose to ensure a generous offer in the two treatments is statistically different. Thus, the authors conclude that the desire to appear fair to others and possibly to oneself is an important determinant in dictator games, which cannot be accommodated by current models of inequity aversion. These results resonate with the observation of why one might cross to the other side of the street to avoid a beggar. In conjunction these findings suggest that in dictator games, many dictators may be *reluctant altruists*. One way in which these findings and those on antisocial punishments in public goods games (see below) can be explained is by a more explicit theoretical model of the microfoundations of human identity and morality. For progress in this direction, see Bénabou and Tirole (2011).

However, it has proved hard to replicate results on moral wiggle room in games where some strategic interaction is involved. For instance, van der Weele et al. (2014) show that moral wiggle room in the trust and the moonlighting games does not reduce reciprocity. This is another instance of why results in the dictator game are quite special and non-robust with respect to other domains.

Ockenfels and Werner (2012) conduct Internet experiments with a large, diverse, subject pool, using the dictator game. The size of the cake that the dictator shares can be small or large. In the Noinfo treatment, the responder never discovers the size of the cake; he only observes the

¹¹ Some may argue that this substantially alters the dictator game, the essence of which is that the respondent's payoff equals the proposer's endowment minus the amount the proposer wishes to keep for himself.

dictator's offer. In the Info treatment, after the dictator has made his offer, the responder is informed about the size of the cake. In the Noinfo condition, when the cake is *large*, the dictator typically offers a share to the responder that is less than one half of the *small cake*. Hence, the dictator hides behind the "small cake effect." By contrast, when the cake is small there is very little difference between the results in the two treatments. These results suggest that giving in dictator games is influenced by second order beliefs, i.e., what the dictator believes the responder believes he will get. It may be the case that the dictator is overcome by guilt if the dictator believes that the responder believes that he will get more than the dictator actually gives (Battigalli and Dufwenberg, 2007). In other contexts, several papers have demonstrated an *audience effect* on giving.¹² These are all cases in which second order beliefs influence one's action. The appropriate framework to rigorously analyze these concerns is psychological game theory (see Part 4).

Andreoni and Bernheim (2009) consider the norm of sharing 50–50 in dictator games. About 20–30% of the dictators offer a 50–50 split to the responder. This can be explained by the dictator's desire to appear fair to the responder, particularly if there are shared social norms about equitable splits. It is important to consider the influence of social norms of sharing in many of the experimental games on other-regarding preferences, an area that is likely to gather more steam in the future.

Lotz et al. (2013) ask if there is a fraction of *genuine altruists* as distinct from *reluctant altruists*, and if we can predict the nature of altruism? They invoke an underlying measure of altruism known as justice sensitivity (JS) that has been used by several psychologists to explain individual-level differences in behavior (Schmitt et al., 1995; Schmitt et al., 2010). Subjects are asked 40 questions and asked to rate each on a scale of 0 (strongly disagree) to 5 (strongly agree). Examples of questions include the following: I cannot easily bear it when people profit unilaterally from me; It bothers me when someone gets something they don't deserve; I have a bad conscience when I receive a reward that someone else has earned. These responses are used to derive measures of sensitivity to becoming a victim of injustice (JS_{victim}), sensitivity of a neutral observer ($JS_{observer}$), sensitivity of an actual perpetrator ($JS_{perpetrator}$), and sensitivity of a passive beneficiary ($JS_{beneficiary}$). The last three measures are combined into a single measure- JS_{others} . The main findings are as follows. At the aggregate level, the authors replicate the result that dictators wish to avoid appearing unfair. However, at the individual level, those with a high JS_{others} score are genuine altruists with stable other-regarding preferences and those with a low JS_{others} score are reluctant altruists. This appears to be an important advance in the literature; it not only suggests a more nuanced view of the dictator game but could also help reconcile conflicting results in this game.

In Part 6 of the book, in the chapter on emotions, we make the distinction between *deliberative choices* and *automatic choices*. Deliberative choices involve the cognitive part of the brain, the prefrontal cortex (system II), that takes account of the long-run consequences, the costs and benefits, and the morality of one's action. By contrast, automatic choices are associated with the more primitive part of the brain, the limbic system (system I); these are quick, instinctive, reactive, pre-programmed choices that may respond to a range of emotions such as fear, anger, and anxiety. Are the choices of players in ultimatum and dictator games deliberative or automatic? One branch of the literature argues that preferences for fairness are deliberative choices. In support, this literature cites evidence that kindergarten children make more self-regarding choices (an automatic response), while adults are more likely to make other-regarding choices. Hence,

¹² See, for instance, Soeteven (2005) in the context of charitable giving in churches and Rege and Telle (2004) in the context of public goods games.

fairness preferences are acquired. A second literature, that is neurobiological and behavioral, suggests that fairness preferences are an automatic response in adulthood.

Neither of the two systems, system I or system II, has primacy over the other. However, system II can use a scarce, but easily depleted, resource, *willpower*, in order to control system I. Hence, willpower-depleted choices are more likely to be in greater conformity with system I. In Halali et al. (2013), subjects first participate in 20 incongruent trials of a *Stroop task*. For instance, they may be shown a red colored ink that is named blue. Subjects then have to suppress their automatic response of saying “red” by saying “blue” instead. This is presumed to deplete their willpower (particularly if the questions are posed rapidly). The evidence suggests that their automatic responses can then be elicited.

In the ultimatum game, following depletion of willpower, proposers make relatively fairer offers. Furthermore, the difference in the choices made under systems I and II (corresponding to relative depletion of willpower) is statistically significant. There are two possibilities. (P1) The automatic response is fear of rejection by the responder, so the offers are higher. (P2) The automatic response is more fair as compared to a deliberative response that is more self-regarding. But if P2 is the correct explanation, then we should also observe more fair offers in the dictator game. By contrast, the automatic response in dictator games is to make more self-regarding offers. This implies that P1 is the more likely explanation, i.e., the automatic response is strategic—fear of rejection. The authors also confirm similar findings for responders in other work.

5.2.3 *Rationality of other-regarding preferences in dictator games*

The rationality of other-regarding preferences hinges on whether they are complete and transitive. An obvious test is to observe the relevant restrictions on actual demand curves, such as the Slutsky restrictions on the cross partial derivatives of Hicksian demands. If this test is satisfied, then the demand curves are the outcome of a utility maximization exercise. But in actual practice, it is hard to observe demand curves.

An alternative approach, using *revealed preference*, is to directly observe the market prices $(p_1, p_2, \dots, p_n)$ and the associated levels of demands $(x_1^*, x_2^*, \dots, x_n^*)$.¹³ Thus, our observations are a sequence of price–quantity pairs $(p_i, x_i^*)_{i=1}^n$. A utility function, u , *rationalizes* these observations if for any $i = 1, \dots, n$ we have $u(x_i^*) \geq u(x_i)$ for all x_i such that $p_i x_i^* \geq p_i x_i$. Thus, x_i^* is revealed to be preferred to all other consumption bundles, x_i , that the consumer is able to afford at the price p_i . In this case, we can infer that the data was generated by a utility maximization process. Hence, the underlying preferences are rational, i.e., complete and transitive.

A sufficient condition to guarantee rationality and, hence, the existence of a utility function that rationalizes the data is the *generalized axiom of revealed preference* (GARP). Suppose that choices are generated by utility maximization. Suppose also that we observe that at price p_i , consumption bundle x_i^* was chosen when another bundle $\bar{x}$ was available and affordable, i.e., $p_i x_i^* \geq p_i \bar{x}$. In this case, x_i^* is *directly revealed preferred* to $\bar{x}$; we write this as $x_i^* D \bar{x}$. If choices were generated by utility maximization, then we must have $u(x_i^*) \geq u(\bar{x})$.

We may also, in our data, observe a series of binary choices of the following form: $x_i^* D x_j$, $x_j D x_k$, $\dots$, $x_n D \bar{x}$. In this case, we say that x_i^* is *revealed preferred* to $\bar{x}$; we write this as $x_i^* R \bar{x}$. If choices were indeed generated by utility maximization, then we may again conclude that $u(x_i^*) \geq u(\bar{x})$.

¹³ Each x_i^* may be a vector, in which case the corresponding price, p_i , is also a vector.

Suppose that we impose the restriction of non-satiation on the utility function, i.e., $u(x) < u(x + \varepsilon)$, $\varepsilon > 0$. In this case, if data are generated by utility maximization, the observation $p_i x_i^* > p_i \bar{x}$ would imply that $u(x_i^*) > u(\bar{x})$. In particular, it cannot be that $u(x_i^*) = u(\bar{x})$. For if it were, then, given $p_i x_i^* > p_i \bar{x}$, we can find some x in the neighborhood of $\bar{x}$ that is affordable, $p_i x_i^* \geq p_i x$, yet, it gives higher utility, $u(x_i^*) = u(\bar{x}) < u(x)$, on account of non-satiation. But this contradicts that x_i^* was chosen at a price p_i .

Now suppose that we observe that at price p_i , consumption bundle x_i^* was chosen when another bundle $\bar{x}$ was strictly cheaper, i.e., $p_i x_i^* > p_i \bar{x}$. In this case, we say that x_i^* is *strictly directly revealed preferred* to $\bar{x}$; we write this as $x_i^* SD \bar{x}$.

Definition 5.1 (*Generalized axiom of revealed preference, GARP*): Suppose that x_i^* is chosen at a price p_i when $\bar{x}$ is available and affordable. If $x_i^* R \bar{x}$, then it cannot be the case that $\bar{x} SD x_i^*$.

Andreoni and Miller (2002) test for GARP in the dictator game, using the following method. Suppose that a dictator is given an income y and he wishes to propose the split (y_P, y_R) . For every unit of income that the dictator wishes to give to the responder, he needs to give py_R units to the experimenter, where $p > 0$. The experimenter then credibly gives y_R units to the responder. The variable p may be thought of as the price of giving. If the dictator has other-regarding preferences, then he solves a standard utility maximization problem in which he maximizes $u(y_P, y_R)$, subject to the budget constraint $y_P + py_R = y$. The levels of income used were $y = 40, 60, 75, 80, 100$, and the prices used were $p = 0.25, 0.33, 0.5, 1, 2, 3, 4$.

Andreoni and Miller (2002) use Afriat's *critical cost-efficiency index* (CCEI) to measure whether a violation of GARP is severe or not. The CCEI gives an indication of how much of the budget must be shrunk in order to rationalize the data. The closer the index is to 1, the less severe are the violations. Out of 176 subjects, there were 18 violations of GARP, i.e., the choices of 18 subjects cannot be rationalized by utility maximization. Based on the CCEI index, there were only 4 violations. Random choice, on the other hand, would have led more than three-quarters of the subjects to violate GARP. With the exception of 3 subjects, the observed data can be rationalized by a quasiconcave utility function. This evidence suggests that the subjects are rational and one can find a utility function that rationalizes the choices that they have made.

At a price of $p = 1$, the dictator gave away, on average, about 23% of the endowment, which is similar to the results in Forsythe et al. (1994). There is heterogeneity among the subjects and only a small minority are unconditional altruists. A little less than a quarter of the subjects behaved as if they had self-regarding preferences, 14.2% of the subjects tried to equalize their payoffs with the other player (Leontief preferences) and 6.2% of the subjects had preferences of the perfect substitutes form.

Results along similar lines, i.e., consistency with GARP and heterogeneity at the individual-level, when CES utility functions are fitted to individual-level data in dictator games, are also reported by Fisman et al. (2007). In a more recent study, Iriberri and Rey-Biel (2013) also implement a similar framework to that in Andreoni and Miller (2002) and Fisman et al. (2007) in dictator games. They find that self-regarding preferences are the most common (47%), followed by inequity averse preferences (21%), social welfare maximizers (19%), and competitive preferences (13%).

Volk et al. (2012) elicit contribution rates of subjects in public goods games (see Section 5.4) at three different dates: in the base period, and 2.5 months and 5 months from the base period. They classify subjects into three categories: conditional cooperators, free-riders, and others. At the aggregate level, the distribution of the level of cooperation is unchanged at the three dates.

This is important because most inferences in public goods games are drawn at the aggregate level (see Section 5.4 below). There is variation in intertemporal stability of preferences at the individual level. While two-thirds of the subjects stay in their category at all three dates, a third switch categories.¹⁴ Personality traits, in particular, the trait of being *agreeable* is a significant predictor of temporal stability of preferences at the level of the individual.

In conclusion, there is nothing irrational about other-regarding preferences.

5.3 Gift exchange and trust games

5.3.1 *The gift exchange game*

The *gift exchange game* was introduced by Fehr et al. (1993). Its objective was to test for the presence of *intrinsic reciprocity*, using a labour market setting. It drew on the insights of Akerlof (1982) and Akerlof and Yellen (1990), who had already explored departures from competitive market outcomes in macroeconomic models by using notions of fairness and reciprocity in labour markets. In exploring the effects of wage rates greater than the market clearing wage rates, as in models of efficiency wages, Akerlof (1982, p. 544) writes: “On the worker’s side, the ‘gift’ given is work in excess of the minimum work standard; and on the firm’s side the ‘gift’ given is wages in excess of what these workers could receive if they left their current jobs.” These critical insights pervade the literature that followed.

Consider the gift exchange game used in the experiments conducted by Fehr et al. (1997). Subjects in experiments were assigned to one of two roles: a principal or an agent. The identities of trading partners were not revealed to any experimental subject, ruling out repeated game effects and any attempts to build a reputation for cooperative behavior. Principals offer a contract (w_b, e_n) to agents that specifies a *binding* wage level, w_b , and a *non-binding* effort level, e_n . Since a binding effort level cannot be specified, there is contractual incompleteness in the model. The wage level can be chosen from the set $\{0, 1, 2, \dots, 100\}$, while the effort level can be chosen from the set $\{0.1, 0.2, \dots, 1\}$. Let e_a be the *actual* effort level chosen by the agent. The payoff of the principal is $100e_a - w_b$. The payoff of the agent is $w_b - c(e_a)$, where $c(e_a)$ is an increasing and convex function of effort with the endpoints given by $c(0.1) = 0$ and $c(1) = 18$. The outside option of the agent is zero.

If all agents possess self-regarding preferences, then the outcome is very simple. Because the effort level is non-contractible, each self-regarding agent chooses the lowest effort level, $e_a = 0.1$. The rational, forward-looking, principal correctly anticipates this, so he offers a wage $w_b = 0$, which ensures that the individual rationality constraint of the agent binds, i.e., $w_b = c(0.1) = 0$.

Suppose that the agent has *other-regarding preferences* or *intrinsic reciprocity*, and that this is known to the principal. In this case, the principal anticipates that the intrinsic reciprocity of the agent would induce him to reciprocate higher wage offers with higher actual effort levels. If the cost of eliciting higher effort is not too high, then it may be optimal to offer $w_b > 0$ and, on account of gift exchange, one would expect to observe $e_a > 0.1$.

The experimental results are depicted in Figure 5.1, where the vertical axis measures the actual effort, e_a , and the non-binding contractual effort level, e_n . The horizontal axis measures the level of rents, $w_b - c(e_n)$, embodied in a contract. The following conclusions arise.

¹⁴ These estimates on stability are higher than some of the other estimates; see, for instance, Muller et al. (2008).




Figure 5.1 The relation between effort and rent in the gift exchange game.

Source: Reprinted from *European Economic Review*, 46: 687–724. Fehr, E. and Falk, A. "Psychological foundations of incentives." Copyright © 2002, with permission from Elsevier.

1. There is a direct relation between e_a and $w_b - c(e_n)$. Hence, workers respond to higher rents by actually working harder. These results are consistent with at least some fraction of the workers having a preference for intrinsic reciprocity but are consistent with principals having either self-regarding or other-regarding preferences. (a) Even principals with self-regarding preferences who are aware that the agents have intrinsic reciprocity will wish to tap into such reciprocity by offering higher wages that elicit higher effort, provided that eliciting effort is not too costly. (b) So long as a high enough fraction of workers has intrinsic reciprocity, the principal will have an incentive to offer rents, i.e., wages greater than the market clearing level even if other workers have purely self-regarding preferences.
2. Conditional on any given level of rents, there is some shirking, i.e., $e_a < e_n$.
3. There is a significant fraction of self-regarding agents who put in the minimum effort level, $e_a = 0.1$. Fehr and Falk (2002) summarize the weight of the evidence from several studies and suggest an indicative figure for the fraction of self-regarding subjects that lies between 40% and 60% percent. The remaining fraction is accounted for by subjects who have other-regarding preferences.¹⁵

Suppose now that, ex-post, the principals can punish workers who work less than the contracted level, $e_a < e_n$, and reward those who work more than the contracted level, $e_a \geq e_n$. If all principals had purely self-regarding preferences, then introducing such an option should have no effect. This is because at the end of the game, bygones are bygones and punishments do not directly increase the material utility of individuals. However, if principals are believed to have other-regarding preferences or intrinsic reciprocity, then agents rationally believe that they could

¹⁵ For an estimation of the relative proportions of self-regarding and other-regarding economic agents, see also Berg et al. (1995), Fehr and Falk (1999), McCabe et al. (1998), Abbink et al. (2000).

be punished for underperforming.¹⁶ Hence, we may expect an increase in actual effort, e_a . The main results in Fehr et al. (1997) are as follows.

1. If $e_a < e_n$, then principals punish in 68% of the cases.
2. If $e_a \geq e_n$, then principals reward in 41% of the cases.
3. Workers seem to anticipate the principal's reciprocity. When $e_a < e_n$, workers anticipate being punished in 54% of the cases. When $e_a > e_n$, workers expect to be rewarded in 98% of the cases.
4. The presence of rewards and punishments decreases shirking from 83% of the cases to only 26% and increases the average level of effort from 0.37 to 0.65.

The gift exchange game has been replicated a large number of times. The general finding of a positive relation between mean effort level and the rent offered by the principal is quite robust.¹⁷ Many studies, some reviewed below, indicate that the positive relation between effort and wages is less strong than originally thought but that there is a much stronger reduction in effort following a wage decrease. Thus, wage decreases appear to be much more costly relative to the gains to firms from equivalent increases in wage. For instance, the elasticity of effort with respect to a wage increase is 0.15 in Gneezy and List (2006) and 0.30 in Kube et al. (2006). However, the elasticity is 0.82 for a wage decrease in Kube et al. (2006).

Fehr et al. (2009) conjecture two reasons for the finding of a weak positive effect of rents on effort in the empirical literature. First, the sample sizes are quite small. Second, the base wage over which rents are offered is already too high. Evidence indicates that workers are more likely to make a greater increase in effort if the base wage is perceived as unfair.

Falk (2007) ran the following field experiment with a non-student population. A Swiss charity made two different kinds of gifts in solicitation letters sent out to 10,000 potential donors. In the baseline case of the *no gift condition*, no gift was included. In the *small gift condition*, the request for donations was accompanied by a single card from a child in Bangladesh. In the *large gift condition*, four such cards were included. The relative frequency of donations in the no gift condition, small gift condition, and the large gift condition was respectively 12%, 14%, and 21%. These differences were statistically significant, which confirms the insights of the gift exchange game. Further, there was insignificant intertemporal substitution in the sense that larger donations in the current period did not crowd out future donations.

Indirect evidence of gift exchange is also found in situations where strikes or other labour relation disruptions in companies lead to a perception of unfairness on the workers' part. If workers perceive that they are unfairly treated, they then tend to reduce effort. Such indirect evidence has been found by Mas (2006) at the New Jersey police department, Mas (2008) at Caterpillar, Krueger and Mas (2004) at Bridgestone/Firestone tires, and Lee and Rupp (2007) for airline wage settlements.

Do the results of the gift exchange experiments survive with an increase in stakes? Fehr et al. (2002) conducted experiments in Russia where subjects earned, on average, the monetary equivalent of 1 week's worth of wages in one condition, and 10 week's worth of wages in another condition. No significant differences were found in the results between the two conditions. These stakes are high enough to cover a very large fraction of real-world economic transactions.

¹⁶ Principals may well have such preferences. Fehr and List (2004) use the trust game in Fehr and Rockenbach (2003), which was performed on students, to examine the behaviour of CEOs in Costa Rica. They find that CEOs are actually more trusting and trustworthy relative to students.

¹⁷ See, for instance, Charness (2000), Fehr et al. (1993, 1998), Fehr et al. (2007), Fehr and Falk (1999), Gächter and Falk (2002), and Hannan et al. (2002).

The gift exchange experiments of Charness et al. (2004) show that there are possibly cultural differences between European and US subjects in experiments. The US subjects engage in lower levels of gift exchange. They also find negative effects on the magnitude of gift exchange when payoff tables are provided. However, it is not clear why the provision of these tables should matter, e.g., does it clarify the structure of the problem or induce particular framing effects?

GIFT EXCHANGE AND COMPETITIVE EQUILIBRIUM

The results of Fehr et al. (1998) speak to the class of results that find support for the predictions of competitive market equilibria in experimental markets (Smith, 1962; Davis and Holt, 1993).¹⁸ Fehr et al. (1998) embed a gift exchange game within competitive markets by allowing sellers to choose a variable quality for a good (see the introductory chapter for more details). If buyers pay a higher price, then sellers can choose to supply them with a higher quality. Under self-regarding preferences, the sellers should supply the lowest quality, and anticipating this, buyers should offer low prices. By contrast, the experimental evidence shows that buyers offer a higher price in anticipation of reciprocal behavior on the part of sellers, who reciprocate by supplying higher quality. The outcome of the experiments departs significantly from the competitive market equilibrium predictions, due to fairness and reciprocity considerations. Furthermore, such reciprocal behavior has a positive effect on the profits of buyers. In conjunction with the results from ultimatum games that reveal more prosocial behavior in the presence of greater market integration, these results challenge the conventional view that the market necessarily produces self-regarding outcomes.

GIFT EXCHANGE AND SOCIAL COMPARISONS

The standard gift exchange experiments assume a bilateral employment relation, i.e., there is one principal and one agent. This feature is relaxed in several experiments. Maximiano et al. (2007) extend the standard framework to allow for four workers working for the same principal and find that the effort levels of the workers in the bilateral and multilateral conditions are very similar.

Social comparisons seem to be an important determinant of reciprocity. Gächter et al. (2012) modify the gift exchange game by considering one employer and two employees. The employer announces a publicly observed wage. The employees then choose their effort levels sequentially. Employee 1 chooses his effort level first, which is observed by employee 2, who then chooses his effort level. When the wage level is high enough, then a higher effort level by employee 1 is followed by a higher effort level by employee 2. This effect is over and above the normal effort levels in the gift exchange game. Cohn et al. (2011) consider randomized field experiments in which teams of two agents are paired with a principal. They consider two separate treatments. In the first, the wages of both agents are cut by 25%, while in the second, the wage of only one of the two agents is cut by 25%. The effort level of the agent whose wage is cut in the second treatment decreases twice as much (by 34%) relative to that in the first treatment (by 15%). This demonstrates that social comparisons play an important role in gift exchange experiments.

In an interesting set of experiments, Abeler et al. (2010) consider repeated interaction between a principal and two agents. Neither effort nor wages are contractible. In one treatment, the equal wage treatment (EWT), both workers are paid equal wages, so there is no horizontal inequity. In

¹⁸ The results of experimental competitive markets are discussed in more detail in Part 4 of the book.




Figure 5.2 The effect of the equity norm on effort levels.

Source: "Gift exchange and workers' fairness concerns: When equality is unfair," Abeler, J., Altmann, S., Kube, S., and Wibrall, M., *Journal of the European Economic Association*, 8(6): 1229–1324. Copyright © 2010 European Economic Association.

a second treatment, the individual wage treatment (IWT), unequal wages are allowed. EWT may be viewed as fair if both workers put in identical levels of effort. However, in the case of unequal efforts, an EWT may be perceived as unfair and, hence, may prove detrimental to effort. The experimental results support these findings; see Figure 5.2 for the average effort level across 12 rounds. In the EWT treatment, the worker who initially puts in a higher level of effort responds to equal wages by putting in a lower effort level in the next round. Effort, in this treatment, moves downwards over successive rounds. By contrast, in the IWT, the worker who puts in a lower effort level in one round, receives a lower wage from the principal and responds by putting in a higher effort level in the next round. In this case, effort increases over successive rounds. Average wage is positively correlated with the effort levels as in gift exchange games.

Kube et al. (2012) find that gift exchange is more effective when the firm offers a non-monetary reward to the worker as compared to a monetary reward. Students are hired for a data entry task and earnings of 36 euros for a three-hour task are announced. In the *monetary gift treatment*, subjects discover that their wage has been unexpectedly raised to 43 euros. In the *material gift treatment*, workers unexpectedly receive a material gift worth 7 euros. Workers respond by entering 6% more data in the first treatment but 30% more data in the second treatment. Interestingly, in a binary choice comparison, workers preferred to receive 7 euros in preference to the gift. This suggests that the gift is not perceived to be more valuable, however, the gift seems to have a signaling value, which elicits greater effort.

GIFT EXCHANGE AND DURATION EFFECTS

Gneezy and List (2006) make two extensions to the standard gift exchange game. First, they look at field experiments rather than lab experiments. Second, they focus on studying the effort provision of workers over an extended period of time. In one of their tasks, the library task, workers were recruited with an offer of an hourly wage of \$12 for six hours of work. They were not told that they were participating in an experiment. In the no-gift treatment, workers were paid a flat wage




Figure 5.3 Behavior of workers' effort levels in the library task.

Source: Gneezy and List (2006) with permission of the Econometric Society.

of \$12/hour as promised. In the gift treatment, they were unexpectedly paid an increased, but flat, wage of \$20/hour. Effort was initially higher in the second treatment, possibly on account of the gift exchange motive. However, gradually, over six hours, the effort levels in the two treatments converged, as shown in Figure 5.3.

These results beg two questions. First, what are the reasons for the downward movement in effort in the gift treatment? The authors' preferred explanation is that workers in the second treatment have not had a chance to figure out the game with a cool head (the so-called *hot state*, *cold state* comparison). However, an alternative explanation, based on the adjustment of the reference points, also appears plausible. The reference point of workers whose wage has gone up, may adjust upwards, so that the unexpected increase in wage no longer seems like a gift after the adjustment has been made. Second, if real-world firms are rational, then firms should eventually learn that the effects of gift exchange do not last long. So instead of offering a flat unexpected increase in wage to \$20/hour, it might well be optimal to offer a gradual increase in the wage rate in small increments to elicit a positive incremental effort level over time. In this sense, the Gneezy and List (2006) results enrich the insights from the gift exchange game by offering insights about duration effects.

Bellemare and Shearer (2007) point out that the Gneezy and List (2006) results rely on observations taken from spot markets but do not speak to the issue of the long-term relations between workers and firms. They examine the effect of a one-off surprise bonus of \$80 in a tree planting firm in British Columbia where the average daily earnings were \$200. There are workers of varying tenure length within the firm. The results are shown in Figure 5.4 for two typical profiles of workers who begin their employment at the age of 20 and 28, respectively. The effort levels of workers, measured in terms of the number of trees planted daily, are observed over six weeks. Significant long-term effects on the effort levels are found in workers with longer tenure, after controlling for planter-fixed effect, weather, and day-of-week effect (to control for fatigue). Such effects are insignificant for workers with very short tenures (as in the spot markets of Gneezy and List, 2006). In general, there is a positive relation between reciprocity and tenure. This complements nicely the Gneezy–List results suggesting that the gift exchange relationship survives over the long term, provided there are long-term relations between the firm and the workers. Indeed this might even provide a new rationale for long-term relations in firms.




Figure 5.4 The relation between reciprocity and tenure.

Source: Reprinted from *Games and Economic Behavior*, 67(1): 233–44. Bellamere, C. and Shearer, B., "Gift giving and worker productivity: evidence from a firm level experiment." Copyright © 2009, with permission from Elsevier.

REFERENCE DEPENDENCE AND GIFT EXCHANGE

The evidence shows that other-regarding preferences can be reference dependent (Fehr et al., 2009). This may interact with gift exchange issues that future research can examine. Thus, labour market policy interventions such as minimum wage laws could alter the perception of workers about a fair wage and they may respond in the effort decisions, accordingly. Hence, after the introduction of minimum wage laws, firms may find that they must pay higher wages or be perceived unfair. The experiments of Brandts and Charness (2004) show that market competition and minimum wage laws do not have a significant effect on the results of gift exchange experiments, although effort is somewhat lower in the presence of minimum wages. In contrast to these results, Owens and Kagel (2010) find that in gift exchange experiments, the effect of a minimum wage is to increase effort levels. Falk et al. (2006) explore the relation between minimum wages and fairness and find that the minimum wage raises the reservation wage of workers. It is interesting that if minimum wage laws are repealed, then workers do not reduce their notion of a fair wage by as much as the drop in the minimum wage, giving rise to a sort of *ratchet effect*. Hence, firms may find it difficult to reduce wages sufficiently, giving rise to downward rigidity in wages.

Kube et al. (2013) use a field study to examine the contention by Bewley (1999) that firms are reluctant to cut wages for fear of reducing morale and effort. The authors allow for an exogenous variation in the wage in a one-shot interaction between a firm and its workers. Student-workers are hired to undertake a data entry task and informed that the projected wage is EUR 15/hour for six hours. This was intended to serve as a reference point for the wage that workers could expect. In the baseline treatment, students were actually paid EUR 15/hour. In the PayCut treatment, they were paid EUR 10/hour, and in the PayRise treatment they were paid EUR 20/hour. These variations in the wage levels were not related to the past level of wages, a problem that often arises in actual firm-level data.




Figure 5.5 Productivity of workers in different wage conditions.

Source: "Do wage cuts damage work Morale? Evidence from a natural field experiment," Kube, S., Maréchal, M. A., and Puppe, C., *Journal of the European Economic Association*, 11(4): 853–70. Copyright © 2013 by the European Economic Association.

The results are shown in Figure 5.5. In the left panel, each quarter on the horizontal axis represents a 90 minute period. The average number of books marked was significantly lower in the PayCut treatment as compared to the Baseline treatment, and this difference is stable over time. On average, the number of books marked was 21% lower in the PayCut treatment, relative to the Baseline treatment. By contrast, the difference between the PayRaise and the Baseline treatments was not statistically different. Thus, the evidence is not consistent with positively reciprocal behavior but it is consistent with negatively reciprocal behavior and with Bewley's findings. In the right panel of Figure 5.5, we plot the cumulative distribution functions of the number of book entries. The PayCut distribution is first order stochastically dominated by the distributions in the other two treatments.

These results can be explained by models of inequity aversion if one assumes that all players consider only the incomes earned during the field experiment and ignore outside incomes (narrow bracketing). The results can also be explained by models of reciprocity. The authors show that the model by Cox et al. (2007) can explain their findings.

SURVEY EVIDENCE FOR GIFT EXCHANGE

Most of the results on the gift exchange game have been based on experimental findings. In a novel study, Dohmen et al. (2009) use panel survey data from the German Socio-Economic Panel Study (SOEP) from the 2005 wave. Many economists have aversion to survey evidence on a-priori grounds,¹⁹ but a-priori grounds can also be used to defend survey evidence, particularly if people respond to survey questions, using their intrinsic motivation, and have some aversion to lying

¹⁹ The standard arguments are that survey evidence is typically not as tightly controlled as lab experiments, there is no guarantee that the responses are incentive compatible, and face to face interaction while collecting the evidence may bias choices.




Figure 5.6 Average responses for the positive reciprocity questions (left panel) and negative reciprocity questions (right panel).

Source: "Homo reciprocans: Survey evidence on behavioral outcomes," Dohmen, T., Falk, A., Huffman, D., and Sunde, U., *The Economic Journal*, 119(536): 592–612. Copyright © The Authors. Journal compilation © Royal Economic Society 2009.

(see Chapters 7 and 8). Indeed, survey evidence is routinely used by government agencies and firms as the basis of important decisions. Six interesting questions on reciprocity were asked in the 2005 wave. These were as follows.

- (1) If someone does me a favor, I am prepared to return it.
- (2) If I suffer a serious wrong, I will take revenge as soon as possible, no matter what the cost.
- (3) If somebody puts me in a difficult position, I will do the same to him/her.
- (4) I go out of my way to help somebody who has been kind to me before.
- (5) If somebody offends me, I will offend him/her back.
- (6) I am ready to undergo personal costs to help somebody who helped me before.

Questions 1, 4, and 6 elicit positive reciprocity and 2, 3, and 5, negative reciprocity of the responders. Responders had to rank their response on a scale of 1 (does not apply) to 7 (applies perfectly).

The results are shown in Figure 5.6, averaged across the positive reciprocity questions (1,4,6) and the negative reciprocity questions (2,3,5); the vertical scales in the left and the right panels are different. For the three positive reciprocity questions, less than 5% of the responders answer 1 (does not apply) and the modal response is 6 or 7. For the three negative reciprocity questions, the modal response is a 1 or 2 and 10% choose a 6 or 7. There is heterogeneity in responses in both cases but there is significantly greater heterogeneity in the negative reciprocity questions. Positive reciprocity weakens as the costs of reciprocal behavior increases. The response 7 (fully agree) is chosen by 60% when there is no cost (Q1), by 35% when responders need to go out of their way (Q4) and by 22% when there is a personal cost (Q6). While modal responses are low for negative reciprocity, these responses do not appear to be sensitive to the costs of reciprocity. An important finding in this study is that the correlation between positive and negative reciprocity is quite low, 0.021. Hence, the extent of negative reciprocity in a responder does not allow one to predict very well the extent of positive reciprocity.

Suppose that we interpret overtime work as a measure of reciprocity and assign to it a binary variable (1 if overtime, 0 if no overtime).²⁰ A probit regression shows that individuals who score high on the positive reciprocity questions are also more likely to do overtime work and

²⁰ Obviously, overtime work could be entirely supply side driven, and reflect responder circumstances unrelated to reciprocity.

individuals who score highly on the negative reciprocity questions less likely to do so. Not only does the effect of reciprocity, measured in 2005, show up in overtime work in the current year 2005 but also in future years, 2006, 2007. Positive reciprocity has an even stronger effect on the probability of overtime work, when responders perceive that they are paid a fair wage. A similar effect of reciprocity on absenteeism is found.

Mincer-type equations are traditionally used in economics to explain earnings as a function of several explanatory variables such as individual characteristics and educational attainment. The authors get all the classical effects²¹ but, in addition, positive reciprocity enhances earnings in current and future years. However, there is no significant effect of negative reciprocity. Furthermore, individuals who are positively reciprocal are less likely to be unemployed, more likely to have a larger number of close friends, and more likely to report greater satisfaction with life.

5.3.2 The trust game, trust, and trustworthiness

The trust game was introduced by Berg et al. (1995), which in turn is a simplified version of a game in Camerer and Weigelt (1988). Consider a static game, played between an *investor* and a *responder*, sometimes known, respectively, as the *trustor* and the *trustee*. Each player is endowed with some fixed level of endowment, say, 10 MUs (an MU is a unit of money in the experiments). The investor must decide to send, or invest, an amount $i \in \{0, 1, 2, \dots, 10\}$ to the responder, via the experimenter. The experimenter multiplies the amount by a number $m > 1$, known as the *multiplier*, which captures the market return and passes it on to the responder; $m = 3$ in many experiments. The responder then returns an amount $r \in \{0, 1, 2, \dots, mi\}$ back to the investor. The respective payoffs of the investor and the responder at the end of the game are $10 - i + r$ and $10 + mi - r$.

Definition 5.2 (*Trust and trustworthiness*): The numbers i and r are, respectively, measures of trust and trustworthiness.

RELATION BETWEEN THE GIFT EXCHANGE GAME AND THE TRUST GAME

The reader may have noted the similarity in the structure of the trust game with the gift exchange game. The surplus in the trust game, found by adding the payoffs of the two players is $20 + (m-1)i$, which is increasing in the investment made by the investor. By contrast, the surplus in the gift exchange game is $\theta e_a - c(e_a)$, where $\theta > 0$ is the marginal product of effort; for instance, $\theta = 100$ in the Fehr et al. (1997) experiments described above. If θ exceeds the marginal cost of effort (which is the case in all experiments), then the surplus is increasing in the effort level. From these observations, the main difference between the two games is that in the trust game, the first mover (investor) can increase the surplus through his actions (investment). By contrast, in the gift exchange game, it is the second mover (the worker) who can increase the surplus through his action (effort level).

If subjects in experiments have self-regarding preferences and are rational, then the responder will never send any money back to the investor ($r = 0$). The rational investor correctly anticipates the responder's behavior, and sends an amount, $i = 0$. Thus, the unique subgame perfect outcome is $i = r = 0$.

²¹ For instance, men are paid more than women, experience has an increasing and concave effect on earnings, and education has a positive effect on earnings.




Figure 5.7 Measures of trust and trustworthiness in a meta study.

Source: Reprinted from *Journal of Economic Psychology*, 32(5): 865–89. Johnson, M. D. and Mislin, A., "Trust games: a meta-analysis." Copyright © 2011, with permission of Elsevier.

By contrast, if subjects in experiments have other-regarding preferences (or intrinsic reciprocity), then responders are likely to reciprocate higher levels of i with higher levels of r . One would then expect $i > 0$ and $r > 0$. This prediction is borne out by the evidence. Responders typically return back the investment that they have received from the investor. The investor finds that his investment is, at least, mildly, profitable, so he sends a strictly positive amount. However, small upward jumps in i are typically ineffective in eliciting greater r , relative to larger jumps in i .

A large number of replications of the trust game have been conducted. The interested reader may consult the meta-study by Johnson and Mislin (2011) that draws on 161 replications of the trust game with a total of 23,000 respondents. They measure trust by a variable pSent, and trustworthiness by a variable pReturn, defined respectively as $i/10$ and $r/(10 + mi)$, i.e., the fraction of total amounts available that are sent by each of the players. These measures are shown in Figure 5.7.

Across all studies, the mean amount sent by the investor, pSent, is .502 (so investors send about half their wealth), while the mean amount sent by the responder, pReceive, is .372 (so responders send about 37.2% of their wealth). For instance, if initial endowments are 10 each, $m = 3$, and $i = 5$, then $r = 9.3$. The meta-study reveals many other interesting features. The amounts sent are larger when interacting with a human, as compared to a non-human, subject such as a computer; older people send larger amounts; significantly lesser amounts are sent in Africa as compared to North America; students send a significantly lower amount as compared to the non-student population; reducing the multiplier from 3 to 2 reduces the amount returned by the respondent. An important finding from the meta-study is that stakes do not have any affect on the results.

The experimental results in Charness et al. (2011) show that providing a history of the behavior of respondents alters the amounts transferred by the investors. When given a choice of exploring past data on the history of the responders, 70% of the investors choose to check the relevant history. Past behavior that is more trusting and trustworthy gives rise to greater trust. However, subjects also use past data on trust to infer greater trustworthiness, suggesting that these two measures may be correlated. Does it pay to trust based on past data? The answer is yes, in both treatments, although in only one treatment is it statistically significant.

Eckel and Petrie (2011) consider a trust game in which investors and responders can, at a price, look up a photograph of the other player.²² They explore if the amount of trust and trustworthiness changes as a result of looking at a photograph of the other player. The motivation for this paper is earlier work that discovered a link between economic decisions and human characteristics such as beauty, ethnicity, and gender.²³ In their trust game experiments, both parties are endowed with 10 tokens each (each token is equivalent to \$1.5). In various sessions, the subjects face different prices for buying photographs of the other player. Not all subjects decide to purchase the photographs. At the lowest price of \$0.20, 50% of the subjects buy the photographs, but they respond to an increase in price by reducing their demand for photographs. Hence, there appears to be a downward sloping demand curve for photographs.

The difference between the demands for photographs of investors and responders is not statistically significant. However, investors are willing to pay more for the photographs, relative to responders, because their decision in a trust game is relatively more strategic. There is no gender difference in the demand for photographs, but there is an ethnic difference. White investors are 40% more likely to wish to buy photographs as compared to subjects of any other ethnicity. Conditional on buying a photograph, white investors discriminate in their investment, i , based on the ethnicity of the responder. On average, they send 3.6 tokens to a black responder and 5.5 to a white responder; the difference is statistically significant. Black investors send significantly fewer tokens than whites, but they do not discriminate in their investment choices based on ethnicity.²⁴ There is no difference in the number of tokens sent based on the relative genders of the investor and the responder.

For responders, the response, r , to the amount sent by the investor, i , is upward sloping. The highest slope of the response function is for white responders. In general, conditional on observing a photograph, responders send more tokens to the investor, relative to the case where they do not see a photograph. In their roles as responders, both blacks and whites send more tokens back to an investor of the same ethnicity. The authors conclude that the “informational value of a face is non-zero.” However, the informational value differs for different ethnic groups.

We have used the term *trust* in an informal manner, without a formal definition. This is clearly unsatisfactory. In applied economic work, various proxies for trust are used. This literature indicates that measures of trust are correlated with GDP, growth, inflation, volume of trade, and the propensity to invest in the stock market. What sort of trust does the trust game capture? There have been attempts to answer this question in terms of altruism and *betrayal aversion*. For a fuller discussion of these very interesting issues and the formal definitions of the relevant terms, see the excellent survey by Fehr (2009).

The trust game is not the only context in which *trust* and *trustworthiness* have been measured. Bolton et al. (2004) consider these concepts in the context of feedback systems in online electronic purchases such as those on Amazon and Ebay.²⁵ In these experiments, a buyer first chooses to buy or not. If the buyer chooses to buy, he needs to pay the money upfront. The seller then chooses to

²² In related research, dictators give only a third as much to an anonymous responder relative to the case where they are told that the responder is the Red Cross; see, for instance, Eckel and Grossman (1996).

²³ See Croson and Gneezy (2009) for a survey of the link between decisions and gender. For the link between beauty and decisions, see, for instance, Hamermesh and Biddle (1994), Mobius and Rosenblat (2006), and Andreoni and Petrie (2008).

²⁴ The authors conjecture that whites are in a majority, so they are more likely to be “teachers in school or bosses in workplace” hence, they are less likely to be discriminated against.

²⁵ See also the follow-up paper by Bolton et al. (2013) that speaks to the limitations of field studies in online auctions in engineering changes that might promote more trust through various feedback systems.

dispatch the good or not. The seller's action is subject to moral hazard because he might choose to run with the money and not dispatch the good. This setting characterizes many forms of online purchases. Three markets are considered. In the *strangers market*, buyers and sellers meet only once and do not have access to each other's transactions history. In the *feedback market*, an online system tracks the previous shipping decisions of the seller and reveals them to the buyer at the time of the buying decision. However, buyers and sellers are randomly paired in each round. Finally, in the *partners market*, feedback is as in the feedback market, but the buyer and seller pairings are unaltered over the experiment.

In this online experimental market, we may define *efficiency* by the percentage of potential transactions completed. *Trust* is measured by the percentage of buyer orders and *trustworthiness* by the percentage of seller dispatches, once the buyer has decided to buy. On all these indicators, the partners market does the best, followed by the feedback market, while the strangers market performs the worst. The reason for the difference is that each of these markets has different incentives for building reputations, and there are differences in the quality and interpretation of feedback. For instance, we may use information on the last item shipped, or not shipped, by the seller as a predictor of trust in each market. Such information has no effect on trust in the strangers market (because such information is not available). However, conditional on the last item being shipped, in the feedback market, the probability of trust doubles from 33% to 65%. The effect on trust in the partners market is even stronger. The quality of feedback in the partners market is superior because current feedback improves one's own future buying decisions. By contrast, in the feedback market, the benefit accrues to others. Thus, one market form is superior, relative to the others, in the internalization of information externalities arising on account of the feedback.

The experimental results of the ultimatum game show that market integration and cooperation in production are associated with greater sociality. Huck et al. (2012) consider whether increased competition can enhance sociality and efficiency in trust games.²⁶ They consider the game in Figure 5.8, originally due to Bohnet et al. (2005). The trustor (player 1) has two strategies: Trust (T) or not trust (N). The trustee (player 2) also has two strategies: Honor trust (h) or exploit trust (e). The game has a unique subgame perfect Nash equilibrium (N,e). The *trust rate* is defined as the percentage of times the trustor plays T and the *honor rate* is the percentage of times the




Figure 5.8 The trust game in Huck et al. (2012).

Source: Reprinted from *Games and Economic Behavior*, 76(1): 195–209. Huck, Steffen, Lünses, Gabriele, and Tyran, Jean-Robert, "Competition fosters trust." Copyright © 2012, with permission from Elsevier.

²⁶ Trust can also be influenced negatively by the degree of inequality in outcomes in a game; see for instance, Hargreaves Heap et al. (2013).

trustee plays h . The efficiency rate is defined as the percentage of times one observes the outcome (T, h) .

In the *control treatment*, the stage game is played repeatedly over 30 rounds with randomly matched players. In the original experiments by Bohnet et al. (2005), in the control treatment, the efficient outcome (T, h) is played only 5% of the time and, as expected, the trust and honor rates are very low. Huck et al. (2012) introduce additional treatments to isolate the effect of competition. Over the 30 rounds, subjects are assigned either to the role of trustors or trustees over the entire duration of the experiment. A “market” has four players in each role. In the private information treatment (pi), each subject is informed of only their own past interactions (their actions and the actions of the opponent) in the game. In the full information treatment (fi), all trustors have access to the histories of all the trustees. There is a second dimension to these treatments. In the random matching treatment each trustor is randomly assigned to a trustee in each round. In the endogenous matching treatment, the trustors choose which trustees they would like to be matched with. Since trustees compete for trustors, this is known as the *treatment with competition*.²⁷

The results are as follows. In the control treatment, the results are similar to those in Bohnet et al. (2005). The efficiency rate is only 8% and the trust rate and the honor rates are identical at 28%. In the case of random matching, private information (pi) and full information (fi) give similar results but with vastly improved rates of sociality relative to the control treatment—the trust rate almost doubles to 50%, the honor rate doubles to 69%, and the efficiency rate quadruples to 36%.²⁸ However, the most dramatic effect arises in the treatment with competition, irrespective of whether information was private or full. For instance, in the full information treatment with competition, the trust rate, honor rate, and the efficiency rates are respectively 0.90, 0.94, and 0.85, roughly triple those in the control treatment.²⁹

Burnham et al. (2000) postulate that humans have a friend or foe (FOF) mental mechanism. Depending on whether humans perceive their opponent to be a friend or a foe, they make different choices. The authors test for such a mechanism by priming subjects in two-player trust games. For instance, if the other player is termed as a “partner” he might be perceived as a friend and if he is termed as an “opponent” he might be perceived as a foe. They find that partners trust more than opponents. They also find twice the level of trustworthiness among partners relative to opponents.

5.4 Public goods games

In this section, we separately consider the public goods game with and without punishments, followed by an examination of the evidence on such games.

²⁷ If a trustee is matched with more than one trustor in any round, the trustee can choose only one action. His total payoff for that round is the sum of his payoff across all matchings.

²⁸ These findings are in line with those in Bolton et al. (2004) who also find that giving feedback to players about past play improves sociality.

²⁹ In other contexts, with similar endogenous matching of participants, Page et al. (2005) show that contributions in public goods games improve and Bohnet and Kübler (2005) show that cooperation improves in social dilemma games.

5.4.1 Public goods game without punishments

There are $i \in N = \{1, 2, \dots, n\}$, $n \geq 2$, risk-neutral individuals and each has an initial endowment $y > 0$, expressed in some monetary units. They simultaneously decide on their respective contributions, $g_i \geq 0$, to a public good. The total monetary contribution towards the public good is given by $G = \sum_{i=1}^n g_i$. The monetary payoff, m_i , of individual $i \in N$ is given by

$$m_i = y - g_i + rG; \quad \frac{1}{n} < r < 1, \quad (5.1)$$

where r is the monetary equivalent of the return, or benefit, to an individual, arising from one unit of the public good, G .

Definition 5.3 *The variable r in (5.1) is called the marginal per capita return (MPCR).*

The Nash equilibrium outcome of this game when individuals have self-regarding preferences involves *free-riding*; it is derived in the next proposition.

Proposition 5.1 *(i) In the unique Nash equilibrium, all self-regarding individuals make zero contributions, i.e., $g_i^* = 0$ for any $i \in N$. (ii) By contrast, the social optimum is characterized by $g_i^* = y$ for all i .*

Proof: (i) Recall that all individuals are risk-neutral. To check that in a Nash equilibrium $g_i^* = 0$ for any $i \in N$, let the vector of contributions of others $\mathbf{g}_{-i}^* = \mathbf{0}$. Using (5.1),

$$\frac{dm_i}{dg_i} = r - 1 < 0; \quad \because r < 1,$$

hence, $g_i^* = 0$, which proves the result.

(ii) The social objective function is $W = \sum_i m_i = ny - G + rnG$. To find the social optimum, note that

$$\frac{dW}{dg_i} = rn - 1 > 0; \quad \because \frac{1}{n} < r,$$

hence, in the social optimum, $g_i^* = y$ for all $i \in N$. ■

5.4.2 Public goods game with punishments

In the public goods game with punishments, there are two stages. In the first stage one plays the public goods game without punishments. In the second stage, once the individual contributions, $g_1, g_2, \dots, g_n$, are public information, each player simultaneously chooses a punishment vector $\mathbf{p}_i = (p_{i1}, \dots, p_{in})$ where p_{ij} is the punishment inflicted by player i on player j at a cost of c per unit and $0 < c < 1$. In this game, the utility function of a risk-neutral individual i with self-regarding preferences is given by

$$m_i = y - g_i + rG - c \sum_j p_{ij} - \sum_j p_{ji}. \quad (5.2)$$

The first three terms in (5.1), (5.2) are identical. The fourth and fifth terms, respectively, in (5.2) capture the cost of punishing others and the cost of being punished by others.

Definition 5.4 (*Efficiency of punishment*) : Equation (5.2) assumes a linear punishment technology. For any player i , it costs c units (in some appropriate units, such as experimental tokens) to punish player j , $j \neq i$, by 1 unit. The efficiency of punishment is defined by $1/c$. So, a lower cost of punishment/unit corresponds to a higher efficiency of punishment.

Proposition 5.2 When all players have self-regarding preferences, the Nash equilibrium of the public goods game with punishments and without punishments is identical, i.e., in equilibrium the contribution of any player is $g_i^* = 0$ for all $i \in N$.

Proof: Conditional on any vector of first stage contributions $(g_1, g_2, \dots, g_n)$ consider the subgame that begins with the decision by players to engage in punishment. Conditional on the punishment vectors of others, $\mathbf{p}_{-i}$, the choice of $\mathbf{p}_i = \mathbf{0}$, maximizes m_i . Hence, it is a dominant strategy, for each player, to not engage in any punishment in the subgame. Given that the second stage outcome is $\mathbf{p}_i = \mathbf{0}$ for all $i \in N$, the first stage payoffs in (5.1), (5.2) are identical. Hence, using Proposition 5.1, $g_i^* = 0$ for all i . ■

Proposition 5.2 shows that in this multi-stage game, punishments are never optimal for players with self-regarding preferences because bygones are bygones. Hence, any threats to engage in punishment are not credible, thus, punishments carry no bite.

5.4.3 Empirical evidence on public goods games

How do the results in Propositions 5.1, 5.2 compare with the empirical evidence? Experiments on public goods games without punishments typically proceed as follows. There are a fixed number of rounds, say, 10, and a certain number of subjects in experiments, say, $n = 6$. At the start of the experiment, subjects are given some endowment (tokens or points), say, $y = 20$, that can be redeemed at the end of the experiment for money. In each period, subjects are asked to contribute to a common pool or a public account, which captures the idea of a public good. After the contributions have been made, the experimenter places some percentage, say 30% (so $r = 0.3$), of the total amount in the public account into the private account of each individual. This captures some level of uniform benefits from the provision of the public good.

At the end of each period, all subjects are able to observe the total amount in the common pool as well as the individual contributions of others. In public goods games with punishments, a second stage is added to the first in which contributors can punish non-contributors at some personal cost. Typically, for each token surrendered by a punisher to the experimenter, the recipient of the punishment has to surrender more than one token depending on the efficiency of punishment (see Definition 5.4). The same sequence of play then takes place in the subsequent rounds.

The socially efficient outcome is for each individual to put all the endowment into the common pool (Proposition 5.1(ii)). In the socially efficient outcome, each subject parts with the 20 units of endowment but receives $0.3(20 \times 6) = 36$ from the experimenter, for a net gain of 16. However, the essence of the privately optimal solution is that conditional on others contributing, any individual can do better by not contributing, i.e., free-riding. For instance, if the first 5 players contribute, then by not contributing, the 6th player's payoff is $20 + 0.3(20 \times 5) = 50 > 36$. Thus, in a Nash equilibrium, nobody contributes (see Proposition 5.1(i)).

Experimental evidence in public goods games without punishments typically shows that subjects begin by contributing about half of their endowments to the common pool. The contributions then begin to drop in subsequent rounds and towards the last round, the contributions for most subjects are nearly zero, as predicted by the model with self-regarding preferences.³⁰ This evidence from the later rounds is sometimes taken as supportive of the predictions of the model with self-regarding preferences (Proposition 5.1). However, on the whole, this pattern of contributions is not consistent with self-regarding preferences because it does not explain (i) why contributions are initially high, and (ii) why if the game is restarted, the contributions typically jump up to a higher level (Fehr and Fischbacher, 2003).

The results of the public goods game with punishments under self-regarding preferences (see Proposition 5.2) are rejected by the evidence. For instance, Fehr and Schmidt (1999) report that, on average, across various studies considered by them, 73% of the subjects choose $g_i = 0$ in the public goods game without punishments. By contrast, on average, 80% cooperated fully in the public goods game with punishments.

The experimental results of Fehr and Gächter (2000) reject the predictions in Proposition 5.2. First, cooperators in public goods games are willing to engage in costly punishment of the free-riders, although such punishments do not give them any material benefits. Second, the extent of punishment inflicted on the free-riders depends on the magnitude of free-riding. Third, in the presence of punishment, there is less free-riding. Thus, potential free-riders seem to realize that they are likely to be punished if they free-ride.

More fully, Fehr and Gächter (2000) used ten rounds and four subjects in each group. They consider the following three treatments. In the *partner treatment*, the same four subjects remained in the group for the ten periods. In the *stranger treatment*, subjects were randomly reassigned at the end of each period, not ruling out the possibility that some subjects may be matched with each other in more than one period. Finally, the *perfect stranger treatment* is similar to the stranger treatment except that it was ensured that no two subjects get matched in the same group in more than one period. The public goods game with punishments was played in the first ten rounds followed by the public goods game without punishments in the next ten rounds. The results are shown in Figure 5.9.³¹

In the first ten rounds, when costly punishment is allowed, cooperation as measured by average contributions is mostly increasing in all treatments. In the last ten rounds, when the punishment option is removed, we get the standard result that cooperation deteriorates over time. In the partner treatment, when punishment is permitted, one achieves close to full cooperation (i.e., the socially efficient outcome) by round 10. Furthermore, the cooperation level in the partner treatment is greater than any of the other two treatments. Once a group member is punished, it is likely that all subjects in a group believe that inadequate contributions will almost certainly invite credible punishment. The cooperation levels in the other treatments, while lower than the partner treatment, are still impressive and much higher than the prediction of zero cooperation under self-regarding preferences (see Proposition 5.2).

The punishment levels remain at high levels in the final rounds when there is little or no future to worry about. Thus, these punishments are not carried out for instrumental reasons, such as preventing future levels of shirking. Rather, these punishments appear to reflect the innate desire to discipline violators of the norm of cooperation, and/or to derive hedonic utility from punishing

³⁰ See, for instance, Andreoni (1988), Isaac et al. (1985). The results from a wide range of studies on this issue are surveyed in Dawes and Thaler (1988) and Ledyard (1995).

³¹ The gaps in the curve representing the perfect stranger treatment arose because of the finiteness of the number of participants in the experiment.




Figure 5.9 Evolution of average contributions in three different treatments in the presence and in the absence of punishments.

Source: Fehr and Gächter (1999).

non-cooperators. In Part 9 we show that the neuroeconomic evidence supports this view (de Quervain et al., 2004).

We now give a brief and selective overview of some further results on public goods games. By varying the MPCR (see Definition 5.3), Brandts and Schram (2001) find that conditional cooperation is not caused by errors on the part of subjects in experiments. Kurzban and Houser (2004) find that the fraction of subjects who exhibit a preference for conditional cooperation is stable across various public goods games; about 65% of subjects appear to be conditional cooperators. Thus, the results of the public goods game are not simply errors, they indicate stability of reciprocal preferences.

As in the other experimental games on human sociality, there are cultural variations in prosocial behavior for the public goods game. Kocher et al. (2008) find variation in the fraction of subjects who are conditional cooperators in a sample from the US, Austria, and Japan. The highest percentage of conditional cooperators are in the US, at 80.6%; the respective figures for Austria and Japan are 44.4% and 41.7%. Brandts et al. (2004), however, find little variation in the fraction of conditional cooperators in the US, Spain, Japan, and the Netherlands. Herrmann and Thöni (2009) find fairly stable fractions of conditional cooperators (about 56%) that do not depend on rural/urban regions, or on several socioeconomic variables.

Gächter et al. (2008) consider the effects on the degree of cooperation by varying the number of replications of the public goods game (10 and 50 rounds respectively) with two treatments, punishment, P, and no punishment, N. This gives four sets of experiments: 10P, 50P and 10N, 50N. They wish to test if the costly consequences of punishments on payoffs in 10P experiments, relative to the benefits of increased cooperation, are eliminated or reversed in 50P experiments. They find that the highest levels of cooperation and the maximum payoffs among the four cases arise in the 50P game; in particular, these exceed the observed levels for the 10P case. Towards the final rounds of the 50P game, the contribution levels stabilize and one observes very little punishment.

Thus, an observation of little punishment does not mean that punishments are ineffective. Rather very little punishment may be required if individuals make high contributions to the public good; see also Gintis and Fehr (2012).

Sefton et al. (2007) consider the possibility of both positive punishment (sanctions) and negative punishment (rewards) in public goods games. They find that a combination of rewards and punishments elicits greater cooperation relative to punishments or rewards alone. However, in their setup, the efficiency of punishment (see Definition 5.4) is only 1. Nikiforakis and Normann (2008) consider the effect of the degree of efficiency of punishment on the extent of cooperation. They consider four levels of efficiency, $1/c$, namely, 1, 2, 3, and 4. They find that the contributions to the public good are increasing in the efficiency of punishment. For very efficient punishments ($\frac{1}{c} = 3, 4$) they achieve contributions that are close to the socially optimal solution. Using data on Dutch subjects in experiments, Egas and Riedl (2008) also find that the level of contributions are at the socially optimal level for $\frac{1}{c} = 3$, and older subjects are more likely to punish non-contributors.

There are two opposing effects on cooperation as group size increases under the assumption that the cooperation of group members can only be elicited through costly monitoring. In a larger group, there could be a greater incentive to free-ride because the probability one may be monitored by any single player is smaller. However, a countervailing force in a larger group is that there are more monitors. The net effect on the probability of detection of non-cooperation, and, hence, the incentive to free-ride is not clear. Carpenter (2007) considers simulated automata and finds that larger groups are more efficacious in making higher contributions, and they exhibit a lower level of free-riding.

In the typical public goods experiments, punishments are monetary. In traditional societies, however, punishments often take a non-monetary form, such as ostracization and other social sanctions. Masclet et al. (2003) consider a public goods game in which, in addition to monetary punishments, subjects can assign between 1 and 10 “disapproval points” to non-contributors that do not affect their monetary payoffs. They find that both monetary and non-monetary punishments produce greater cooperation as compared to the case of no punishments. However, while in the initial rounds of the experiment, non-monetary punishment is initially quite effective, subsequently, monetary punishment is more effective.

Cinyabuguma et al. (2005) consider another form of non-monetary punishment, namely, expulsion of free-riders by a majority vote. Expelled members can play subsequent rounds of the public goods game but their endowments are halved. Individuals who are expelled, or if a vote is taken on their expulsion, respond by contributing more in subsequent periods. Not many expulsions take place to ensure high contributions, possibly because the threat of expulsion is a sufficient deterrent. The expelled members are typically those who contribute the least, so punishments are not antisocial. There is a fairly rich and varied literature on the differential responses of individuals to democratically and non-democratically chosen policies. For instance, Dal Bó et al. (2010) show that players cooperate more in prisoner’s dilemma games when a policy, say, punish if others do not cooperate, is chosen democratically rather than randomly by a computer.

In a public goods game, Chaudhuri and Paichayontvijit (2006) find that 62% of the subjects are conditional cooperators. If information is given to subjects about the conditional cooperation of other subjects, then conditionally cooperative subjects are found to increase their contributions. Page et al. (2005) allow for the endogenous formation of groups in the sense that subjects can express a preference for who they would like to be paired with after being provided with information about the contributions of others. This also facilitates high levels of contributions to public goods, even in the absence of punishment.

Several other papers study the implications of matching *like-minded players*, in the sense of similarity in contributions in a public goods game, into separate groups. It is found that the contributions in the initial rounds of the public goods game are a useful measure of cooperative disposition (Gunnthorsdottir et al., 2007). Gächter and Thöni (2005) conduct an initial round of a public goods experiment and based on the contributions, they group the top three contributors in the first group, the next three in the second group and so on. Each group then plays a public goods experiment over ten rounds with an MPCR (Definition 5.3) chosen to be 0.6. Even in the absence of any punishment, the average contributions in the first group are nearly socially optimal.³²

Monetary and non-monetary punishments are not the only mechanisms to elicit greater contributions. Bochet et al. (2006) allow for several kinds of communication devices such as *face to face* (direct communication for five minutes prior to the game), *chat-room* (communication via a computer monitor), and *numerical cheap talk* (each person types in a number indicating his/her intended contribution to the public good that is seen by others but no further communication is allowed). Face to face communication is the most effective form and achieves total contributions that are about 96% of the socially optimal solution.

Chaudhuri et al. (2006) explore the implications of several forms of intergenerational advice given by one generation of players playing the public goods game to the next generation. They allow for several treatments that are differentiated on the basis of whether the advice is given to one member of the next generation, or to all members, and if this advice is publicly announced or not. Publicly announced advice is most effective in eliciting the greatest level of contributions, that is 90% of the socially optimal level.

Iturbe-Ormaetxe et al. (2011) consider *threshold public goods games* in which the public good materializes only when total contributions exceed a certain threshold level. In one treatment, the public goods provision game (PGP), subjects simultaneously contribute towards a public good and if the total contributions exceed a threshold, the public good is provided to all participants. In a second treatment, the public goods deterioration (PGD) treatment, subjects are initially endowed with a public good. But they are told that the good will completely deteriorate in the future, unless a threshold amount of voluntary contributions are raised to prevent the deterioration. In standard theory, one should not expect any difference in the results in the PGP and PGD treatments. By contrast, the authors find that the contributions are significantly higher in the PGD treatment. They explain their results using prospect theory. The presence of an initial entitlement to the public good in the PGD treatment creates a higher reference point. Since losses bite more than equivalent gains in prospect theory (loss aversion), subjects are willing to pay relatively more to avoid degradation of the public good. This result is analogous to the explanation of the endowment effect (see Part 1 of the book).

5.4.4 *Heterogenous preferences and the pattern of contributions*

In public goods games, the *expectations of the contributions of others* is a critical factor in determining one's own contributions. Subjects in experiments initially contribute about 40–60% of

³² The level of contributions in the other groups are quite high as well. While the contributions in the second group are also high, the average contributions among the lowest group from round 4 onwards are about half the social optimum. The authors conjecture that the results for the lowest group arise on account of strategic considerations while those among the higher groups arise on account of conditional reciprocity.

their endowment.³³ In the absence of punishments, why does the level of cooperation drop off in successive rounds of the public goods game? One possible explanation is that subjects gain experience and, hence, learn how to maximize their self-regarding payoffs. But this does not explain why these subjects start with high cooperation rates if they are asked to play a repeated public goods contribution game once again (Andreoni, 1988).

The leading candidate for the explanation of these stylized facts is a theory that recognizes the role of reciprocity and the heterogeneity of players. The evidence is consistent with the view that a fraction of the subjects are conditional cooperators, while the remaining have self-regarding preferences. In repeated public goods games, the evidence in Fischbacher et al. (2001) indicates that 10% of the subjects always match the expected average contribution of others, 14% fully match the expected contributions of others until they spend half their endowments, 40% match the expected contributions of others to a lesser extent, 30% are free-riders, and the remaining 6% have no well-defined behavioral pattern.³⁴

For the moment, let us ignore the last group which has no well-defined pattern, and find the rational expectations equilibrium of the model. Denote the expected average contribution of others by $\bar{y}^e$. Suppose that subjects have spent less than half their endowments. Then the behavior of the first two groups, constituting 24% of the population, is given by $y_1 = a_1 + \bar{y}^e$ and that of the third group, constituting 40% of the population, is given by $y_2 = a_2 + b_2\bar{y}^e$, where $0 < b_2 < 1$ and $a_1, a_2 \geq 0$ are constants that are independent of $\bar{y}^e$. For the 30% who are free-riders, the contributions are $y_3 = 0$. Hence, conditional on $\bar{y}^e$, total contributions $y = 0.24y_1 + 0.40y_2 + 0.30y_3$ can be written as

$$y = f(\bar{y}^e) = a + b\bar{y}^e, \quad (5.3)$$

where $a = 0.24a_1 + 0.4a_2 \geq 0$ and $b = 0.24 + 0.4b_2 < 1$.

A rational expectations equilibrium requires us to find the fixed point of $y = f(\bar{y}^e)$. For one parameterization of the model, $y = f(\bar{y}^e)$ is depicted in Figure 5.10 by the line labeled “average conditional contribution.” Crucially, this line is flatter than the 45° line, which is also shown. Suppose that we begin at an initial level of cooperation of 80% and subjects best respond to the observed contributions in the previous period. There is then a gradual temporal decay in cooperation until an equilibrium is achieved with an aggregate contribution of about 10%. The following important lessons emerge from this analysis.

1. Despite the fact that the self-regarding individuals are in a minority (30%), they drive down the level of cooperation of the conditionally cooperative players, who in this case are in a majority (64%), from a high level (80%) to a very low level (10%).
2. One cannot infer the incidence of reciprocity/altruism in a population, just by observing the final outcome. In this case, the aggregate equilibrium contribution of 10% hides the fact that more than half the players are conditional cooperators.
3. Upholding the belief that others will cooperate is very important for a cooperation norm to develop. If the possibility of punishment of non-cooperators is introduced, and the

³³ See, for instance, Dawes (1980) and Fischbacher et al. (2001). For an extensive survey of the results till about 1995 see Ledyard (1995).

³⁴ For an early use of the term “conditional cooperation” in the context of public good experiments, see Sonnemans et al. (1999). Keser and van Winden (2000) find that as many as 80% of the experimental subjects behave in a conditional manner. Fischbacher and Gächter (2010) find that 55% of their subjects are conditional cooperators and 23% are free-riders.




Figure 5.10 Explaining the pattern of contributions when there is heterogeneity in social preferences.

Source: Fehr and Fischbacher (2005).

punishment technology is effective enough, then cooperation can be sustained.³⁵ The credibility of punishment, which arises from intrinsic reciprocity, prevents the non-cooperators from free-riding, hence, strengthening the belief that others will cooperate.

Gürerk et al. (2006) conduct public goods experiments in which individuals self-select themselves into two possible “institutions” over 30 rounds. In any round, there are three stages. In stage 1, subjects choose among two institutions, a *sanctioning institution* (SI) and a *sanction-free institution* (SFI). Only in the SI can subjects engage in ex-post rewards and punishments. In stage 2, subjects are grouped by their chosen institution, given an endowment of 20 MUs, and play a simultaneous move public goods game. In stage 3, subjects may positively or negatively sanction other players in their group, if their institutional choice allows it. At the end of each round, subjects learn the payoffs of other players in each institution before playing the next round, where they get an opportunity to once again choose one of the two institutions.

The results are shown in Figure 5.11. In the first round, the results are as follows. The majority institutional choice is to shun sanctions; 36.9% choose SI to SFI. Of those who chose SI, almost half contribute at least 15 MUs (high contributors) and nearly three-quarters engage in punishment of low contributors. The neoclassical prediction with self-regarding preferences is that players should not engage in any ex-post punishments in this finitely repeated, full information game. High contributors in the SI comprise 13.1% of the total population of 84 subjects, and may be termed as strong reciprocators. This is in contrast to the outcome in the SFI, where 43.4% are free-riders and only 11.3% are high contributors. Only 2.4% of the subjects in SFI may be classified as strong reciprocators. Across both institutions, 15.5% are strong reciprocators in the first round.

³⁵ See Fehr and Schmidt (1999), Fehr and Gächter (2002), Ostrom et al. (1992), and Yamagishi (1986).




Figure 5.11 Choice of institutions and contributions.

Source: From Güerk, Ö., Irlenbusch, B., and Rockenbach, B. "The competitive advantage of sanctioning institutions," *Science* (2006), 312(5770). Reprinted with permission from AAAS. Readers may view, browse, and/or download material for temporary copying purposes only, provided these uses are for non commercial personal purposes. Except as provided by law, this material may not be further reproduced, distributed, transmitted, modified, adapted, performed, displayed, published, or sold in whole or in part, without prior written permission from the publisher.

Over time, subjects endogenously choose the SI, so that it becomes the predominant institution. Furthermore, average contributions in the SI reach almost first best levels of 20 MUs, while those in the SFI converge to near 0 MUs. In round 30, for instance, the average contribution in the SI is 19.4 MUs, while that in the SFI is 0 MUs.

What causes the switch in institutions to a more cooperative form? Suppose that individuals mimic the high payoff players. After round 5, the payoff of high contributors in the SI is higher than the payoff of high/low contributors in the SFI and also higher than the payoff of low contributors in the SI. This provides a powerful incentive for free-riders in the SI to switch to high contributions in subsequent rounds. Furthermore, 80.3% of subjects who switch institutions from the SFI to the SI, immediately increase their contributions as they gain membership of the SI. In a significant finding, 27.1% of subjects who make this institutional switch, convert from being free-riders (0 MUs) to full cooperators (20 MUs). Thus, subjects quickly internalize the norm of contributions and punishments when they switch to a cooperating group; this is reminiscent of conformist behavior in models of cultural selection. Conversely, those who switch in the opposite direction, SI to SFI, immediately reduce contributions in a similar manner.

5.4.5 Unproductive expenditures on public goods

In the standard public goods experiments, punishments enhance efficiency and improve individual payoffs. However, in another class of public goods games, punishments may deter efficiency. Consider wasteful expenditures that might be incurred by groups of players in rent-seeking

games to obtain a prize. The prize has the nature of a public good; it is non-rival and non-excludable for members of the winning group. For instance, two countries fighting a war over rights to a territory, a trade union striking for better wages, ethnic conflicts in countries, R&D competition among firms, several firms lobbying for a government contract, and two political parties fighting an election. These disputes could be between groups of people, or between individuals, or between individuals and groups. Abbink et al. (2010) consider this class of problems using Gordon Tullock's rent-seeking model.

Suppose that there are two competing groups, $j = A, B$. The number of group members are $n_j = 1, 2, 3$, or 4 players depending on the treatment; when $n_j = 1$, the group comprises a single individual. The composition of the groups stays the same over successive rounds of the experiment. Each player is given an initial endowment of $E = 1000$ points, which can either be used to buy contest tokens (one point can be exchanged for one token), or be added to one's points balance, which can be redeemed for money at the end of the experiment. The objective is to win a prize of 1000 tokens, which is given to each player in the winning group. The total prize received by a winning group of size $N \in \{n_A, n_B\}$ is $1000N$. Thus, the prize has the characteristic of a public good (non-rival in consumption and non-excludable).

Each team member, in each group, simultaneously chooses to buy a contest token. Denote the number of tokens purchased by player i in group $j = A, B$ by t_{ij} . Suppose that the total number of contest tokens purchased by groups A and B is $T_A = \sum_{i=1}^{n_A} t_{iA}$ and $T_B = \sum_{i=1}^{n_B} t_{iB}$. As in Tullock's model, the probability that any of the two groups, $j = A, B$, wins the prize equals $\frac{T_j}{T_A + T_B}$. The payoff of any risk-neutral player i in any of the groups $j = A, B$ in any round of the experiment is given by

$$\pi(t_{ij}, T_A, T_B) = E + \frac{T_j}{T_A + T_B} 1000 - t_{ij}. \quad (5.4)$$

There are 20 rounds of the experiment. In each round, each player from each group, selects the number of contest tokens to buy. Using (5.4), the first order conditions with respect to t_{ij} , $j = A, B$, for an interior solution, give rise to two equations: $(T_A + T_B)^2 = 100T_B$ and $(T_A + T_B)^2 = 100T_A$. It follows that, in a Nash equilibrium, $T_A = T_B$. Setting $T_A = T_B$ in any of the first order conditions, one gets the Nash prediction $T_A^* = T_B^* = 250$.³⁶ The Nash equilibrium does not impose any restrictions on the contest tokens purchased by individuals, so long as $\sum_{i=1}^{n_j} t_{ij} = 250$, $j = A, B$.

In a Nash equilibrium, both groups spend the same amounts on buying tokens, thus, the winning probability for each group is $\frac{1}{2}$. Since $\frac{T_j}{T_A + T_B}$ is not defined for $T_A = T_B = 0$, if both groups coordinated by mutually reducing their respective total expenditures to any smallest possible amount $\epsilon > 0$, then their winning probability is also $\frac{1}{2}$. However, since expenditure on tokens reduces utility (see (5.4)), the total expenditure on tokens across both groups $\sum_{j=1}^{n_j} \sum_{i=1}^{n_i} t_{ij}$ may be viewed as the extent of wasteful or unproductive expenditures.

We denote a generic treatment by $x : y$, where $x = n_A$ and $y = n_B$. The first three treatments are the following: 1:1, 1:4, and 4:4. In the fourth treatment, denoted by 4:4P, a punishment option is allowed and each group has four members. In treatment 4:4P, at the end of each round, after observing the contributions of other members of one's group, a group member can engage in costly punishment. Sacrificing a token reduces the number of tokens of the punished player by

³⁶ If one incorporates risk aversion in the payoff functions of individuals, then it can be shown that the equilibrium purchase of contest tokens is smaller.

three tokens. As in public goods games with self-regarding preferences and punishments, bygones are bygones, so we should not observe any punishments in equilibrium. Thus, the Nash equilibrium of the games 4:4 and 4:4P is predicted to be identical when individuals have self-regarding preferences.

We next define a measure of the total unproductive expenditures (tokens plus punishment expenditures) across both groups as a percentage of the total amount won, 1000N.

Definition 5.5 (*Rent dissipation*): The “rent dissipation not accounting for punishment costs” is given by $\sum_j \sum_i t_{ij}/1000N$, where N is the number of members in the group that wins the prize. Suppose that player i in group j inflicts a total punishment of p_{ij} tokens on other group members. Then “rent dissipation accounting for punishment costs” is given by $\sum_j \sum_i (t_{ij} + p_{ij})/1000N$.

In the fifth, and last, treatment denoted by 4:4/40, the treatment in which $n_A = n_B = 4$ is conducted over 40 rounds, instead of 20 rounds.

The results are as follows. Figure 5.12 considers treatments 1:1, 1:4, and 4:4P in which there is no punishment. It shows the average contest expenditures over the 20 rounds as well as the benchmark Nash equilibrium prediction of a contest expenditure of 250 tokens. When a single individual plays against a team of players (as in the treatment 1:4), then in the curve labeled (Teams), the sum of the expenditures of all team members is shown, while in the curve labeled (Individuals) only the individual expenditure of the singleton member is shown.³⁷




Figure 5.12 Contest expenditures in the treatments without punishment.

Source: Abbink et al. (2010) with permission from the American Economic Association.

³⁷ Thus, the curve labeled 1:4 (Individuals) shows the expenditure by the single team member of group A, while the curve labeled 1:4 (Teams) shows the expenditure by the combined team members of group B. The curve labeled 4:4 (Teams) shows the average across each of the two teams.

In each of the treatments, the difference between the Nash prediction and the actual expenditures is statistically significant.³⁸ On average, individuals spend about twice the Nash prediction and teams spend about four times the Nash prediction. The average contributions reduce in successive rounds in the direction of the Nash prediction. Yet in the last five rounds, when contributions are lowest, in treatments 1:4 and 4:4, they are still significantly higher relative to the Nash prediction in a statistical sense. Teams contribute significantly more towards conflict expenditures relative to individuals.

The downward trend in contributions over successive rounds makes one wonder if with sufficiently many rounds, the contributions would converge to the Nash prediction. However, when the 4:4/40 treatment is run, the contributions stabilize by round 30, so that there is no significant difference in contributions between rounds 31–35 and 36–40. The average contributions still remain significantly above the Nash prediction.

The effect of punishments is seen by comparing the two treatments, 4:4P, and 4:4. The punishment option significantly increases the level of conflict expenditures. Expenditures in treatment 4:4P are everywhere above the levels of the 4:4 treatment and higher in a statistically significant sense, even in the last five rounds. Indeed, if one adds the conflict expenditures across both teams of four players, they are nearly equal to the prize that the winner gets (4×1000). This dissipates almost the entire gains from the costly contest.

The punishment option is frequently used by the subjects in experiments in experimental games when it is available, as in public goods games. In classic public goods games experiments, due to free-riding, the actual contributions are lower than the socially optimum level, which is typically full contribution of the endowments (Fehr and Gächter, 2000). In this case, punishments nudge the contributions towards the social optimum and increase the payoffs of players. However, in the case of conflict expenditures, even though the prize has the features of a public good, the players contribute higher than the target (Nash equilibrium) and punishments drive the contributions further away from equilibrium and reduce the payoffs of players.

Figure 5.13 shows the extent of rent dissipation in the presence and the absence of punishment for each of the four treatments, for each quarter of the 20 rounds. For the 4:4P treatment we show, separately, rent dissipation, while accounting for and not accounting for punishment (Definition 5.5). Rent dissipation is very high for the 1:1 treatment.³⁹ Interestingly, rent dissipation in the 4:4P treatment is higher than the 4:4 treatment and even higher than the 1:1 treatment. Thus, unlike the public goods game, punishments lead to huge, non-productive, expenditures in this case.

5.4.6 *Antisocial punishments*

The efficiency enhancing role of punishments in public goods experiments has mainly been documented in experiments conducted in the US and in Western Europe. In these societies, the indices for the *norms of civic cooperation* and *rule of law* are high.⁴⁰ Norms of civic cooperation encompass attitudes to tax evasion and abuse of the welfare state, while rule of law reflects peoples' trust in law enforcement institutions. Herrmann et al. (2008) consider the role of punishment in

³⁸ This is not a new result; see Abbink et al. (2010) for the relevant references.

³⁹ Recall that in the definition of rent dissipation, the denominator is $1000N$, where N is the number of members in the group that wins the prize. Ceteris paribus, rent dissipation is highest when $N = 1$.

⁴⁰ These indices can be found, for instance, in the World Values Survey (see <http://www.worldvaluessurvey.org/>).




Figure 5.13 Actual and equilibrium rent dissipation.

Source: Abbink et al. (2010) with permission from the American Economic Association.

public goods games in 16 diverse societies (see Figure 5.14 for a list) with varying norms of civil cooperation and the rule of law.

The subject pool is undergraduate students who were similar in age and socio-demographic characteristics. The main finding is the presence of *antisocial punishments*, i.e., punishment of cooperators by non-cooperators, in public goods games that can reduce the level of cooperation. Groups of four players played the public goods game over ten rounds in the no punishment (N) and the punishment (P) treatments. Each subject had an initial endowment of 20 tokens, and for each token contributed to the public good, each player received a return of 0.4 tokens (in the terminology of Section 5.4.1, $r = 0.4$).

Figure 5.14 (left panel) shows the mean level of *prosocial punishments*, i.e., punished players contribute lower amounts, and *antisocial punishments*, i.e., punished players contribute higher amounts. Five categories of punishment are shown. The category labeled $[x, y]$, $x < y < 0$ shows the punishment of someone who contributed between x to y tokens less than the punisher (prosocial punishment). When $0 < x < y$, the category $[x, y]$ shows the punishment of players who contributed between x to y tokens more than the punisher (antisocial punishment). Prosocial punishment is widespread and the degree of variability is low. However, there is high variability in the level of antisocial punishment and differences in such punishments across groups are statistically significant.⁴¹ The right panel of Figure 5.14 shows a negative relation between the mean levels

⁴¹ In a follow-up study, Gächter and Herrmann (2009) find large cross-cultural differences in the antisocial punishment of prosocial cooperators in four cities in Russia and Switzerland.




Figure 5.14 Prosocial and antisocial punishment in the public goods game (left panel) and the relation between antisocial punishment and contributions (right panel).

Source: From Herrmann, B., Thöni, C., and Gächter, S., "Antisocial punishment across societies," *Science* (2008), 319(5868). Reprinted with permission from AAAS. Readers may view, browse, and/or download material for temporary copying purposes only, provided these uses are for noncommercial personal purposes. Except as provided by law, this material may not be further reproduced, distributed, transmitted, modified, adapted, performed, displayed, published, or sold in whole or in part, without prior written permission from the publisher.

of antisocial punishment and contributions to the public good (Spearman's correlation coefficient is -0.87). Thus antisocial punishments are associated with detrimental social outcomes.

One possible determinant of antisocial punishment is the human desire for revenge. Indeed, controlling for the level of contributions, the amount of punishment in period $t - 1$ is a significant determinant of the level of antisocial punishment in period t . This effect is highly significant in 7 groups, weakly significant in 2 groups and insignificant in the remaining groups. In all the groups, the level of cooperation stabilized over successive rounds in the P condition (LHS panel of Figure 5.15). However, unlike the earlier experimental results, punishments did not necessarily lead to high levels of contributions. By contrast, in the N condition (RHS panel of Figure 5.15), the level of contributions converged towards low levels in most groups, although there is still significant variability by round 10. In both treatments, history plays a powerful role. Groups that start off at higher levels of contributions continue at a higher level throughout the 10 rounds of the experiment. On the whole, the cooperation level was relatively higher in the P treatment, in a weakly significant sense, in 11 groups, while there was no significant difference in the other 5 groups.

Thus, a central message here is that when one allows for cultural variations and departures from Western subject pools, the efficacious effect of social sanctions cannot be taken for granted. A regression analysis shows that antisocial punishments are more likely to be carried out in groups where the norm of civic cooperation and the rule of law is weak. Both variables are statistically significant when antisocial punishments take place. However, with prosocial punishments, the rule of law is no longer significant in explaining the differences in contribution rates. In interpreting these results, Gintis (2009, p. 81) writes: "The success of democratic market societies may depend critically upon moral virtues as well as material interests, so the depiction of economic actors as 'homo economicus' is as incorrect in real life as it is in the laboratory."




Figure 5.15 Mean contributions in the P treatment (left panel) and the N treatment (right panel).

Source: From Herrmann, B., Thöni, C., and Gächter, S., "Antisocial punishment across societies," *Science* (2008), 319(5868). Reprinted with permission from AAAS. Readers may view, browse, and/or download material for temporary copying purposes only, provided these uses are for noncommercial personal purposes. Except as provided by law, this material may not be further reproduced, distributed, transmitted, modified, adapted, performed, displayed, published, or sold in whole or in part, without prior written permission from the publisher.

Falk et al. (2005) explore the motivation for private sanctions in a three-player prisoner's dilemma game. After the strategy choices have been made, by giving up a unit of their payoffs, players who sanction can reduce the payoff of another player by an amount $\gamma \geq 1$. In the high sanctions treatment (high γ), they find that 68.5% of the cooperators punish defectors. Indeed, cooperators almost exclusively punish defectors. About 40% of defectors imposed sanctions on other defectors. However, 34% of defectors punished other cooperators—a form of antisocial punishment. At the individual level, cooperators punish more than defectors. In the low sanctions condition ($\gamma = 1$), sanctions cannot reduce payoff differences. Only 2.2% of the defectors punished others but nearly 60% of the cooperators punished other defectors; intentions based-theories (see Chapter 6) may provide one possible explanation.

5.4.7 Evolution and reciprocity

Why does reciprocity arise? One strand of the literature arising from theories of evolution argues in terms of a *selfish gene* where altruism reduces the fitness of a player (Dawkins, 1989). Reciprocity in this framework is a *maladaptation* (Ghiselin, 1974; Hamilton, 1964; Trivers, 1971; Williams, 1966). However, the propensity of individuals to cooperate with others and act in a prosocial manner, reviewed above, especially those who are not related by kin, suggests that self-interest is not the only motive. In other words, there is support for conditional reciprocity, which includes intrinsic reciprocity. Gintis (2000) shows that starting from a population of purely self-regarding individuals, mutants who have regard for others can successfully invade the population, i.e., purely self-regarding preferences is not an evolutionary stable strategy. In equilibrium, there is a polymorphic population that consists of both self-regarding and reciprocal types. The intuition is that during times of hardship, e.g., wars and famine, and in cooperative activities, such as hunting, the ability to cooperate is essential for survival. Purely self-regarding phenotypes will find it difficult to survive such periods. Hence, a certain critical number of prosocial types may become essential to tide over such times. Furthermore, in group conflict, altruism may result in an increase in the fitness of the group, which is a controversial idea in biology. These issues are

sufficiently vexed, and important at the same time, that they require a fuller treatment and the development of the relevant analytical machinery (Bowles and Gintis, 2013; Gintis et al., 2001; Gintis et al., 2003). We devote an entire chapter in Part 5 to these issues.

5.5 How representative is the lab evidence?

Can the experimental evidence provide a useful guide to the actions of individuals in the field? This is the issue of *external validity* of the experiments that is already considered in detail in the introductory chapter; the interested reader must read both treatments in conjunction because they often consider non-overlapping issues. Consider some possible objections to the external validity of the experimental evidence in the context of other-regarding preferences (Levitt and List, 2007). First, there may be adverse selection among students who volunteer for experiments. For instance, those with preferences for reciprocity may self-select themselves to participate in these experiments leading to an overestimation of reciprocity among humans. Second, subjects in experiments may simply be acting in a reciprocal manner because they believe that this is what the experimenter wishes to demonstrate (*experimenter demand effect*).⁴² Third, subjects in experiments interact anonymously with others in the experiment whereas most real-world data is generated by personal interaction among people. These are legitimate concerns. The aim of this section is to examine the empirical evidence for such concerns. Recent scholarship allays fears over many of these concerns.

5.5.1 The student subject pool and experimental scrutiny

Falk et al. (2013) examine the criticism that the experimental results derived from the student population may have biased the findings. From 2004 to 2009, 89% of the studies considered by them used the student population as subjects in experiments. This is a particularly valuable study because it elicits fairness concerns of the student and the non-student populations, using the same experimental protocol. A variant of the trust games is used to elicit fairness preferences.

The main findings are as follows. First, students with reciprocal preferences are not more likely to be self-selected into experiments. Second, the non-student subject pool exhibits more reciprocity in the sense that responders return significantly more money to the proposer in the trust game. So if anything, the student population may lead to downwardly biased inferences about the degree of reciprocity. A range of other evidence supports the assertion that the non-student population is relatively more prosocial; see, Fehr and List (2004), List (2004), Bellemare et al. (2008), Carpenter et al. (2008), Burks et al. (2009), and Baran et al. (2010). In a gift exchange game, Hannan et al. (2002) also find lower reciprocity among undergraduate students relative to MBAs under identical experimental conditions. Their explanation is that MBAs have greater experience in job environments where gift exchange is important.

Cleave et al. (2012) conduct experiments with 1,173 students that are recruited to play a trust game and a lottery choice task. The resulting data are assumed to be the preferences of the general student population. These 1,173 students are then invited to participate in a lab experiment. Those who turn up for the lab experiment exhibit preferences similar to those of the general student

⁴² For a thought provoking discussion of the issues involved in such *experimental demand effects*, and how they might be minimized, see Zizzo (2010).

population. Thus, the authors conclude, it is not necessarily the case that more prosocial students select themselves into lab experiments.

Could students simply be making more prosocial choices in the lab because they perceive that this is what the experimenter desires (experimenter demand effects and experimental scrutiny)? Or, perhaps, they may feel that self-regarding or greedy choices are likely to be frowned upon by other observers (*audience effect*), hence, they may, falsely, exhibit other-regarding preferences in the lab? These concerns are articulated in Levitt and List (2007). Bandiera et al. (2005) find that the monitoring of workers, relative to non-monitoring, induces more prosocial behavior. Recall also the discussion on Haley and Fessler (2005), above, in which simply showing subjects a pair of eyes in dictator games may induce them to give more, yet these results do not carry over to games with strategic concerns (Fehr and Schneider, 2010).

Barmettler et al. (2012) directly address the question of experimental scrutiny in a carefully controlled experiment using three experimental games—a dictator game, an ultimatum game—and a trust game. In the dictator and the ultimatum games, the first mover, respectively, the dictator and the proposer, was given CHF 20 to allocate between himself and the second mover (the responder). In the trust game, the trustor and the trustee were given an endowment of CHF 10 each. Each subject played each game only once in a between-subjects design.

The typical method of conducting these experimental games is through a computer interface. However, this requires the presence of an experimenter who may scrutinize the choices that students make. For this reason, the authors conduct their experiments with paper and pen in a classroom, and the students were seated sufficiently far apart to ensure anonymity of choices. The students were provided packages with colored envelopes and pens and a personal ID number that only they knew. They made their choices on paper and stuffed the choices back into the envelopes. The envelopes were collected by the experimenters, and the relevant payments based on the choices were put back into the envelopes.

In the single-anonymous treatment (S), at the end of the experiment, the subjects reported their id numbers to the experimenter who matched it to the envelopes and handed the envelopes to the subjects. In this case, the experimenter could have observed the choices of the students. In the double-anonymous treatment (D), the subjects collected their payments directly, and privately, from a table on which the envelopes with ID numbers were placed. In this case, the experimenter could not identify the choices made by the participants.

The results are shown in Figure 5.16. In none of the three games and for any player in any role (proposer or responder, trustor or trustee), is there any statistically significant difference in the choices of the players across treatments S and D. In the dictator game, in treatment S, the median amount transferred by dictators is CHF 3.9, while in treatment D it is 3.6; the difference is statistically insignificant. These results are consistent with the results in Hoffman et al. (1994) for dictator games who varied the degree of experimental anonymity but the results were not statistically different.

For the ultimatum game, the mean offer of the proposer is CHF 8.6 in treatment S and CHF 8.3 in treatment D. The difference is statistically insignificant and the median offers are identical. The behavior of responders in treatments S and D is not statistically different.⁴³ In the trust game, the

⁴³ Bolton and Zwick (1995) find significant effects of anonymity on the results for ultimatum games in some of the trials. But anonymity only explains a small percentage of the departures from equilibrium play in their experiments. Barmettler et al. (2012) also note procedural differences across the single- and double-anonymous treatments in Bolton and Zwick (1995).




Figure 5.16 Results for the single- and double-anonymous treatments.

Source: Reprinted from *Games and Economic Behavior*, 75(1): 17–34. Barmettler, F., Fehr, E., and Zehnder, C., “Big experimenter is watching you! Anonymity and prosocial behavior in the laboratory.” Copyright © 2012, with permission from Elsevier.

median transfer by the trustor in treatment D is CHF 8, while in treatment S it is CHF 7. In this case as well, the differences are statistically insignificant and the average transfers are identical.

Critics of the effect of experimental scrutiny on other-regarding preferences have perhaps been too harsh. If experimental scrutiny turns out to be important, it would serve to illustrate the powerful role of other-regarding preferences. Some critics argue that simply by having an observer (an experimenter) in the room or having a picture of a pair of eyes of an anonymous individual in the room seems to switch-on other-regarding preferences. Many real-world interactions are either directly observed by third parties or allow for credible information transmission to third parties. Using the critics’ argument, institutions that wish to promote human sociality then appear to have an extremely low-cost method of switching-on other-regarding preferences in most cases. This is in line with anthropological findings of cooperation in small groups of individuals; see Gintis (2009).

5.5.2 The external validity of experiments on other-regarding preferences

Cohn et al. (2012) address issues of the external validity of lab experiments in the gift exchange game. Using a Swiss subject pool, where subjects distributed newspapers at train stations and other public places, working on a regular hourly pay of CHF 22, they implemented an increase in salary of CHF 5 to CHF 27. They found that only workers who felt that the original salary was




Figure 5.17 Workers' performance response to the wage increase (vertical axis), conditional on their fairness perceptions.

Source: Cohn et al. (2012). Reproduced with permission. Copyright, INFORMS, <http://www.informs.org>

unfair responded by putting in more effort. Those who felt that the original salary was fair, did not increase their effort in a statistically significant manner; Figure 5.17 shows the workers' average performance and the standard errors. In a lab experiment, conducted on the same subjects, they were able to elicit the other-regarding preferences of the subjects. Those subjects who behaved selfishly in the lab did not increase their effort in the field, following a wage increase, regardless of whether they considered the original salary (CHF 22) to be fair or unfair. Thus, the effect shown in Figure 5.17 is driven by workers who were prosocial in the lab. In this case, behavior in the lab appears to have a high degree of external validity.

Carpenter and Seki, (2011) study the behavior of two types of fishermen in Toyama Bay, Japan, who mainly fish for shiroebi (Japanese glass shrimp). The *poolers* who have been operating as such since the early 1960s, pool their incomes, expenditures, and the fish catch. This activity requires substantial coordination, cooperation, and concern among the individuals. The second group of fishermen, the *non-poolers*, who have been operating since 1992, work on an individual basis; they do not pool their resources, or the fish catch. The aim of the study is to see if there is a link between the other-regarding preferences of the two groups, as elicited by a *modified public goods game* (MPG game) and the choice of their real-life activity, poolers or non-poolers. Thus, the study directly addresses the issue of the external validity of lab experiments on other-regarding preferences.

In the MPG game, conducted in 2003, groups of four individuals are formed, and the experimenter gives each individual an initial endowment that is equivalent to about USD 4.30. The experiment lasted less than an hour and each participant earned USD 73.19, on average. The experiment is conducted separately with poolers, non-poolers, and students recruited from nearby universities.

In the first five rounds of the experiment, the participants play a standard public goods game without punishments. The experimenter doubles the total contributions of the group and then redistributes the resulting total among the four members. Since the return on a unit of contribution is 0.5 for each group member, it is optimal to free-ride if one has purely self-regarding preferences. The efficient outcome is for each individual to contribute their full endowment. In the remaining five rounds of the experiment, a *social disapproval protocol* is unexpectedly introduced at the end of each round. According to this protocol, the contributions of all group members are randomized and shown to the others. Conditional on having observed the contribution of others, group members can then buy, at a cost of 10 yen each, either a smiley face (to indicate approval) or a sad face (to indicate disapproval). These faces are then shown to the group at the beginning of the next round, so that group members can gauge the approval levels of others.

The MPG game allows the experimenter to identify several aspects of the other-regarding preferences of poolers, non-poolers, and students. For instance, when individual contributions in round t are regressed on total contributions in round $t - 1$, the intercept term recovers the degree of *unconditional cooperation*, while the slope term recovers the degree of *conditional cooperation*. Similarly, the propensity to send sad faces recovers *disapproval*, and the reaction in round t to the number of sad faces based on contributions in round $t - 1$, recovers *response to disapproval*.

One can now directly examine the productivity of fishing boats belonging to poolers and non-poolers and examine if productivity is affected by (1) conditional cooperation, (2) unconditional cooperation, (3) propensity to show disapproval, and (4) response to disapproval. A range of other control variables, e.g., the experience of the fishermen, is carefully introduced. Conditional cooperation, recovered from the lab evidence, is significantly positively correlated with actual productivity of the fishermen. This tendency is stronger for poolers as compared to the non-poolers. Conditional cooperation in the lab is the single most important variable in explaining the higher productivity of poolers relative to non-poolers.

The propensity to show disapproval is also a statistically significant factor in explaining higher productivity, although the response to disapproval is not positively related to productivity. Finally, poolers are found to exhibit greater unconditional cooperation in the lab, relative to non-poolers. It is eminently possible that many decades of being involved in team production has shaped the other-regarding preferences of poolers. As in Falk et al. (2013), the finding here is that the student population is less cooperative. The authors remark in their conclusion that: “This sort of result—students being less pro-social than field participants—has become a “stylized fact” of the field experimental literature.”

Fehr and Leibbrandt (2011) study the behavior of fishermen in rural fishing villages in north-eastern Brazil. There are two kinds of individuals. Those who catch shrimp (shrimpers) and those who catch fish (fishermen). The shrimpers and the fishermen operate on a lake that has common property rights with no legal limits on the amount of the catch and the costs of fishing are relatively small. The main focus is on the fishing instruments used. Shrimp traps with small holes or fishing nets with smaller holes increase the catch but create an immediate and future negative externality for others. Shrimpers and fisherman are aware of this channel of causality.

The size of the holes is taken as a measure of the cooperation of the shrimpers and the fishermen. The smaller the holes, the less cooperative is the observed behavior. Shrimpers and fishermen first play a public goods game (PG). Each participant is also asked to state, prior to making their contributions, the contributions that they expect of others. It is found that the expected contributions of others are the best predictor of one's contributions in the PG game, which suggests




Figure 5.18 Checking the external validity of lab data on social preferences.

Source: Reprinted from *Journal of Public Economics*, 95(9–10): 114–55. Fehr, E. and Leibbrandt, A., “A field study on cooperativeness and impatience in the Tragedy of the Commons.” Copyright © 2011, with permission from Elsevier.

the presence of conditional cooperation. In a separate experiment, the time preference parameters of the shrimpers and fishermen are also elicited because the negative externalities accrue now and in the future.

Figure 5.18 illustrates the main insights of the paper. The horizontal axis classifies the contributions in the PG game into three categories—low, medium, and high. The vertical axis measures the actual hole size used by shrimpers and fishermen. More cooperative behavior in the PG game is also associated with a larger hole size, which corresponds to greater cooperation in the field data. Fehr and Leibbrandt (2011, p. 1146) write: “We show that there is no insurmountable gap between the laboratory and the field, even though the context in our laboratory environment differs from the field context in important ways.” Impatience, as measured in the time preference experiments, is not related to contributions in the public goods game. Hence, while both cooperation and impatience are important in the degree of exploitation of common resource problems, they operate through separate channels.

The subject pool in Leibbrandt (2012) is shrimp fisherman in Brazil who sell their shrimp, individually, in open air markets. The fishermen participate in a public goods experiment (PGE) and a risk aversion experiment in which the identity of the opponents is kept anonymous. In the PGE, each player is given 10 MUs (monetary units). In the unique Nash equilibrium, each player should free-ride and in the socially optimal equilibrium, each player should contribute all 10 MUs. The number of MUs that are contributed is used to measure the *degree of prosociality*. This allows the segregation of all subjects into two pools.

Those that contribute less than 4 MUs (54.6%) are classified as *less prosocial* and those who contribute 4 MUs or more (45.4%) are classified as more prosocial. Following the collection of lab data, two kinds of real-life data are collected from the shrimp sellers. First, survey data on various aspects of their sociality. Second, actual data on prices and quantities sold. The results are as follows.

1. Less prosocial sellers are more likely to sell their shrimp catch at a relatively lower price. Except for a minimal range of prices, the distribution function of the price of shrimps sold

by more prosocial sellers first order stochastically dominates the distribution function of less prosocial sellers. Sellers who completely free-ride in the PGE, achieve the least shrimp price per unit. Controlling for several other factors (shrimp size, risk aversion, experience, gender, and color of shrimp) sociality, as observed in the lab, has a significantly positive effect on the prices of the individual shrimp sellers.

2. More prosocial sellers are less likely to lose important buyers and more likely to have trade relations of greater duration. For instance, 59.3% of the less prosocial sellers had trade relations less than or equal to 3 years, while the corresponding figure for more prosocial sellers is 37.7%. The average duration of the trade relationship for more prosocial sellers is 6 years, and the corresponding figure for less prosocial sellers is 4.5 years.
3. In the self-reported surveys, more prosocial sellers also report greater confidence in the extent to which they can signal their trustworthiness.

These results are intriguing. Future research can check for the robustness of these findings in other contexts and clarify the transmission mechanism between other-regarding preferences and the relevant economic variables. However, this research rules out two of the channels that were considered to be important. (1) Prosocial players self-select themselves into jobs with higher earnings and productivity. (2) Prosocial individuals earn more because they interact better with colleagues. Explanation 1 does not work because both types of sellers are found in the same profession and the second explanation does not work because this is a market in which the sellers are individual players. The authors conjecture that the most likely transmission mechanism is the one pointed out by Bowles et al. (2001). Since prosocial sellers are less likely to engage in opportunistic behavior with buyers, they “interact more smoothly,” which could have reputational benefits.

Rustagi et al. (2010) consider issues of forest conservation, a common property resource, in Ethiopia. The Bale Oromo people in the Bale region were given property rights over the forest in return for managing the forest cover through conservation efforts, sustainable tree felling, and monitoring misuse by others. In the first stage of the experiment, a public goods game was played with the subjects. Groups of two players were formed. Each of the players had to give a response to how much they would contribute, conditional on the seven possible levels of contributions of the other players; this is the *strategy method* of eliciting choices. One of the players was then picked at random and asked to make an offer, following which the other player's offer, revealed from the strategy method, is implemented.

A player is defined to be a conditional cooperator if his contribution is increasing in the perceived contribution of the other player. Based on the experimental results, players are categorized into four types. The most interesting types in the current context are the *free-riders*, or type 0 players (11.49%), *weak conditional cooperators*, or type 2 players (11.63%), and *strong conditional cooperators* or type 3 players (34.02%); the remaining players, type 1 players, have mixed preferences and it is difficult to classify them cleanly into any of these three categories. The main difference between weak and strong conditional cooperators is the difference in statistical significance of the correlation between their contributions and the contributions of others.

In terms of forest management, the main indicator of cooperation is the extent of costly personal monitoring of the forest, say, through forest patrols, to prevent misuse of the common resource. The main hypothesis is that cooperation in the public goods game and the degree of cooperation in managing the common forest resource are correlated. This is borne out by the evidence, as shown in Figure 5.19. There is a statistically significant relation between cooperation in the lab and in the field. The authors conclude on p. 964: “In sum, better forest management outcomes are not only a result of conditional cooperators being more likely to abide by the local rules of the group but also being more willing to enforce these rules at a personal cost.”




Figure 5.19 Average time spent on monitoring by different behavioral types. The bars represent standard errors of the means.

Source: From Rustagi, D., Engel, S., and Kosferd, M. "Conditional cooperation and costly monitoring explain success in forest commons management," *Science* (2010), 330(6006). Reprinted with permission from AAAS. Readers may view, browse, and/or download material for temporary copying purposes only, provided these uses are for noncommercial personal purposes. Except as provided by law, this material may not be further reproduced, distributed, transmitted, modified, adapted performed, displayed, published, or sold in whole or in part, without prior written permission from the publishers.

Karlan (2005) considers participants in a microcredit program in Peru. The participants are members of a non-profit group lending organization in Ayacucho, Peru that allows poor women to borrow and save. Loans are given for a period of four months and the organization actively encourages members to develop a savings habit. In the first step, these participants play a standard trust game. The amount given by the first mover (investor or trustor) is an indication of the degree of *trust*, while the amount sent back by the second mover (the responder or trustee) is often termed as the degree of *trustworthiness*.

The aim of this study is to see if actions taken in the trust game can predict the outcomes of the financial decisions a year later. In relation to the outcomes in the trust game, the two main hypotheses are the following.⁴⁴ (1) A greater degree of trustworthiness should be associated with a greater likelihood of repaying loans. (2) A higher degree of trust should be associated with a greater gambling instinct. Such players should save less. Both hypotheses are confirmed by the evidence.

Hargreaves Heap et al. (2012) test the external validity of lab experiments on prosociality on the Gisu people in eastern Uganda. The Gisu are known for individualistic behavior and relatively low levels of prosociality, even with insiders in the same village or people sharing a common lineage. The authors played the Trust game with the Gisu and tested two hypotheses. The first is that the Gisu are expected to exhibit relatively less prosocial behavior compared to the typical lab findings. The second is that even when one observes prosocial behavior, insiders will be favored

⁴⁴ The author also considers a public goods game and results of a General Social Survey, that we do not report here.

less relative to the typical experimental findings. The authors are unable to reject any of these hypotheses.

Baran et al. (2010) consider the entire cohort of 2008 Chicago MBA students. These students play the trust game in one of their courses. As a measure of trustworthiness, 83.8% of the students in their roles as the trustee, return at least 30% of the amount sent to them by the trustor. Eighteen months later, the same students are followed up to observe the amount of donations they make back to the school. Trustworthiness in the trust game is a significant predictor of the actual donations made by the students. Thus, lab experiments have high external validity and subjects do not just exhibit more trustworthiness in experiments to appear nice to the experimenter. The authors feel that their results are strong enough to conclude that: "In sum, our paper provides evidence that lab measures are a reasonably accurate way of predicting individual behavior in the field, bolstering support for using laboratory experiments to measure other-regarding preferences. It also supports the claim that reciprocity is an important—even characteristic—preference of many, but not all, individuals."

Benz and Meier (2008) consider University of Zurich students who are asked each semester, at the time of renewal of their registration, if they would like to anonymously and voluntarily contribute to one of two funds. Fund A helps students in financial difficulties and one can contribute CHF 7 (about USD 5) and Fund B supports foreign students at the University of Zurich and one can contribute CHF 5 (about USD 4). The students can also choose not to contribute to any fund. The decisions of the students are taken as the relevant field data.

These students subsequently participate in one of two experiments where they have to decide to allocate an endowment of CHF 12 given to them by the experimenter. In the first experiment, the social funds experiment, 99 students are asked to contribute to the same fund (or not contribute at all) that they did in the field experiment in the university. In the second experiment, 83 students are asked to choose to contribute between two randomly chosen charities outside the university. The students were not aware that they were participating in a scientific study that was comparing their field and experimental actions. The field and the experimental data are in close agreement. For instance, the average contribution of the students in the social funds experiment was CHF 9.46, while in the field study they had contributed CHF 9.07 in the previous semester and CHF 9.55 in the four semesters after the experiment took place. This indicates stable preferences that are picked up well by the field data and the experiments.

List (2006) reports on lab and field experiments conducted with sports card traders.⁴⁵ The lab experiments are carefully conducted in a location very near the field experiments and are closely matched in protocol. The study examines whether sellers reciprocate higher price offers of buyers by offering sports cards of higher quality, and whether this response differs in the lab relative to the field. Overall, quality for each price is lower in the field relative to that in the lab. For instance, the average quality offered by sellers for price offers of \$20 by buyers in the lab and the field are respectively 3.1 and 2.3; the corresponding figures for a price offer of \$65 are 4.1 and 3.1. However, the response of price to quality is comparable in the lab and in the field.

Levitt and List (2008) interpret the results in List (2006) as greater reciprocity in the lab relative to that in the field. However, Camerer (2015) re-examines this data in a within-subjects analysis in which the same individuals participated in the lab and the field experiments. He shows that the data reported in Levitt and List (2008) is for a smaller sample of non-local dealers only who account for just 30% of the sample. It turns out that the non-locals are indeed relatively more reciprocal in the lab. However, this result reverses for the local dealers (70% of the sample) who

⁴⁵ See also the introduction of this book for a more detailed account of this research.

are relatively more reciprocal in the field. Pooling across the two kinds of dealers, and using three different statistical tests of significance, no firm conclusions can be drawn for the difference between lab and field data. Camerer concludes (p. 271): “To be crystal clear, I am not sure whether dealers behave differently in the lab and the field, in the most closely matched comparison of lab-market and field (cards).” Camerer (2015) reviews a set of papers where the lab features are deliberately matched with field features to ensure comparability between the two. He finds only one paper, by Stoop et al. (2012), based on data collected from fishermen in Holland, where prosociality in the lab is weaker relative to the field experiments. Even in this case, he believes that (p. 275): “it is not strong evidence because there are many differences between field and lab.”

5.6 Indirect reciprocity

A variety of games studied above, e.g., the ultimatum game, gift exchange game, trust game, public goods game, illustrate the powerful role of reciprocity. *Indirect reciprocity*, on the other hand, refers to the reciprocation of kindness or unkindness between two parties, by a third party whose payoffs are not necessarily directly influenced. The term indirect reciprocity was used by Alexander (1987) who used it to provide an evolutionary basis for cooperation and punishment in human societies.

Seinen and Schram (2001) use the *helping game* to illustrate the importance of indirect reciprocity. Subjects in the experiment were paired randomly in repeated bilateral interactions. In each interaction, one of the subjects was given the role of a donor and the other, the role of a recipient. The only decision for the donor is to decide to *give* an amount of money, $m > 0$, at a cost $c > 0$ such that $c < m$, or to *pass* the opportunity to give, in which case, the game ends. The recipient is passive in each interaction. The act of giving (or helping) is what gives this game its name.

All individuals are given “scores” that reflect their social status. A score at any point in time, indicates the number of times that an individual chose to give (helpful choice) as a donor, in the past six decisions. Hence, the social status of an individual belongs to seven possible categories: 0, 1, . . . , 6. For instance, a social status of 4 means that the individual chose to give four times in the past 6 rounds, but declined the opportunity to give twice. Cost, c , can be either high (*HC* treatment) or low (*LC* treatment), and information on social status can either be public (*I* treatment), or private (*N* treatment). The interactions between these four categories are then considered. For instance, *LCI* is the treatment in which costs are low and information about social status is public. The results are as follows.

1. The highest fraction of helpful choices is in the *LCI* treatment and the lowest in the *HCN* treatment.
2. The difference in the fraction of helpful choices between the following pairs of treatments is statistically significant at the 5% level: *LCI-HCN* and *HCI-HCN*.
3. Subjects with high social status are almost always reciprocated.⁴⁶ This provides evidence of indirect reciprocity. Thus, it seems entirely rational for subjects to build a reputation for helping using the channel of indirect reciprocity.

⁴⁶ Hargreaves Heap and Varoufakis (2002) first consider the effect of endogenously generated social status in a hawk-dove game. Then, their subjects play a modified hawk-dove game that adds an extra pure strategy—cooperate. Subjects who were discriminated against in the first game were more likely to cooperate with each other in the second game.

4. Subjects are often found to be using cutoff rules, e.g., help the recipient if her social status is greater than some cutoff level.
5. There are interesting dynamics, which suggest that the initial level of social status in a critical minimum level of the population is crucial in determining whether, over time, the subjects will evolve towards higher or lower cooperation.

Engelmann and Fischbacher (2009) argue that the study by Seinen and Schram (2001) confounds two factors—indirect reciprocity and strategic considerations. So, donors might wish to give for two possible reasons. (1) For strategic reasons, in order to build a reputation in terms of social status, so that in their future roles as recipients, they receive greater help from the donors. (2) On account of indirect reciprocity, to reward recipients whose social status is high. In their experimental design, the donor can give up 6 units of the experimental money to enable the recipient to receive 15 units. The game was repeated over 80 rounds. For some subjects, the *social status* based on scores (number of times they helped in 5 previous times in the role of donors) was made public in the first 40 rounds, and for the rest, in the last 40 rounds.

Thus, in any round, one could observe the interaction among subjects with different social status. When information was made public, the donor was told about the social status of the recipient before making the decision to give. Evidence was found for both confounding factors—indirect reciprocity and strategic considerations. For instance, 25% of the subjects did not help when the social status was not made public but did help when it was made public. This is evidence of strategic considerations. There is a positive relation between the giving actions of donors and the social status of recipients, irrespective of the social status. However, the amount of help given to those with high social status is twice as much as that given to those with low social status, which provides evidence of indirect reciprocity.

In an interesting variant of the dictator game, Fehr and Fischbacher (2004) introduce the possibility of third party punishment. If the third party is willing to incur a private loss, but achieves no personal gain to enforce a “norm of giving,” then it furnishes evidence for indirect reciprocity. In this third party punishment game (TPG), there are three players. Player A (the proposer) has an endowment of 100 MU (unit of money in the experiments), player B (the respondent) has 0 MU, and player C (the third party) has 50 MU. A dictator game is played between players A and B. Player C observes the game, and in his role as the third party, decides whether to punish or reward player A for this actions. By spending 1 MU, C can reduce A’s payoff by 3 MU, however, there is no direct personal benefit to C from punishing player A. All players know the extensive form of the game. The experimenters conjecture that there is some sort of *fairness norm* which applies to the TPG, so that player C is predicted to punish low offers by player A.

Figure 5.20 shows that the magnitude of punishment meted out by player C is higher if the dictator makes a lower transfer. The very fact that punishments are undertaken by player C indicates that preferences are not self-regarding. Since 50–50 is a potentially salient fairness norm, in 55% of the cases, player C punishes A when offers to B fall below 50%. It is interesting that 70–80% of player Bs expect A to be punished when the offer falls below 50%. Many institutions in society, such as the police, law, third party intervention in bargaining, and regulation, ensure that powerful parties do not exploit their monopoly positions.

Charness et al. (2008) consider a trust game in which the material payoff of a third party is unaffected by the decisions of the other two parties who play a trust game. They consider three separate treatments. Treatment 1 is a baseline treatment, which is the standard trust game in the absence of a third party. Treatment 2 allows the third party to punish the trustee for not sending enough money back to the trustor. Treatment 3 allows the third party to reward the trustor for receiving an insufficient amount of money from the responder. In Treatments 2 and 3, the actions




Figure 5.20 Third party punishments in the dictator game.

Source: Reprinted from *Evolution and Human Behavior*, 25(2): 63–87. Fehr, E. and Fischbacher, U., "Third-party punishment and social norms," Copyright © 2004, with permission from Elsevier.

of the trustee and the trustor are found to be altered towards greater trust and trustworthiness, relative to the baseline treatment in which the third party is absent. For instance, the trustor sends 60% more to the trustee in Treatment 2. However, there is no appreciable difference in the amount sent when there is the possibility of third party reward. This is the main result in the paper; recall that in the trust game, the trustor controls the surplus in the relationship, so he is the more important player. The interpretation of the results for the trustee are more complicated because he is influenced not just by the amounts sent by the trustor, but also by third party punishment. Yet, on the whole, trustees too return a higher amount to the trustor in the presence of third party punishment.

CHAPTER 6

Models of Other-Regarding Preferences

6.1 Introduction

An impressive range of theoretical models have been developed to address the empirical findings on social preferences. To organize the discussion in this chapter, we have found it useful to employ the following distinction among the alternative models.

In *models of social preferences*, decision makers care not only about their own payoffs but also the payoffs of others in some reference group. In *models of interdependent preferences*, decision makers respond to the actions of others, conditional on the perceived type of others. For instance, altruistic individuals may respond in an altruistic manner if they perceive others to be altruistic, but not if others are perceived to be self-regarding. In *models of intention-based reciprocity*, the perceived *intentions* of others are the critical determinant of how one responds to the actions of others. The types of others (e.g., altruistic or self-regarding) are irrelevant in these models. Finally, there are *hybrid models* that combine two or more of the earlier models and/or use a richer array of social objectives that individuals might have, such as the *maximization of efficiency* or the *utility of the worst-off member* of society.¹

We begin with some simple examples. Consider a set of individuals $i \in N = \{1, 2, \dots, n\}$ whose respective monetary payoffs are denoted by $y_1, y_2, \dots, y_n$. Suppose that the preferences of any individual $i \in N$ can be represented by a utility function.

Example 6.1 (*Altruism*): Individual $i \in N$ has purely self-regarding preferences if, his utility function is $u_i = u_i(y_i)$. Such preferences are unable to explain the results of the experimental games considered in Chapter 5 such as the ultimatum game, the trust game, the gift exchange game, and the public goods game.

Let $\mathbf{y}_{-i}$ be the vector of payoffs of all individuals other than individual i . Preferences of individual $i \in N$, are said to be altruistic if the utility function can be written as $U_i = U_i(y_i, \mathbf{y}_{-i})$, such that (1) there exists at least one element y_j in $\mathbf{y}_{-i}$ such that $\frac{\Delta U_i}{\Delta y_j} > 0$, and (2) there does not exist any y_k in $\mathbf{y}_{-i}$ such that $\frac{\Delta U_i}{\Delta y_k} < 0$. From the evidence reviewed in Chapter 5 we know the following. Behavior in dictator games seems to accord with such preferences, and

¹ This distinction is by no means unique, or universally agreed upon (see footnote 1 in Dufwenberg and Kirchsteiger 2004, p. 269). For instance, our suggested terminology does not readily admit emotions such as *guilt* and *anger*. However, psychological game theory is treated separately in Part 4 of the book.

is, in some cases, enhanced when players feel that they are being watched. In a modified dictator game, Andreoni and Miller (2002) found that the behavior of subjects was consistent with the generalized axiom of revealed preference (GARP), so such preferences are not irrational. There was substantial heterogeneity in the actions of subjects. The behavior of 23% of the subjects was spiteful, i.e., they acted so as to reduce the payoffs of others. Purely altruistic preferences cannot also explain the behavior in trust games as well as in public goods games with punishments.

Example 6.2 (Envy): Suppose that we have only two players, $n = 2$. Consider the utility function due to Bolton (1991),

$$U_i = U_i \left(y_i, \frac{y_i}{y_j} \right); i, j = 1, 2 \text{ and } i \neq j,$$

where $\frac{\partial U_i}{\partial y_i} > 0$, and $\frac{\partial U_i}{\partial \frac{y_i}{y_j}} > 0$ if $y_i < y_j$, which captures envy, but $\frac{\partial U_i}{\partial \frac{y_i}{y_j}} = 0$ if $y_i \geq y_j$.

However, envy alone cannot explain why individuals give a positive amount in the dictator and the ultimatum games. It also does not explain behavior in the trust game or the gift exchange game.

Examples 6.1 and 6.2 suggest that neither altruism, nor envy, alone, can provide the foundations for a model of social preferences. A more complete description of human behavior requires an integration of selfishness, altruism, and envy. The model of Fehr and Schmidt (1999) (henceforth, FS) that we consider in Section 6.2, incorporates these diverse concerns. Furthermore, any model that aspires to explain social preferences must, at a minimum, explain the findings of the basic experimental games in Chapter 5.

In the FS model, inequity is self-centered. Individuals compare their own payoff with the payoffs of those who are better off (envy) and those who are worse off than them (altruism). The FS model does not invoke the beliefs of players about others (first order beliefs) or the beliefs of players about the beliefs of others (second order beliefs), so, it does not model the role of *intentions*; a rigorous analysis of these concerns requires the machinery of psychological game theory (see Part 4 for details).

The FS model can successfully explain many of the findings of the experimental games in Chapter 5; these include the ultimatum game (Section 6.2.1, below), the dictator game, the gift exchange game, and the public goods game with and without punishments (Section 6.2.2, below). The channel through which the FS model explains positive giving in ultimatum and dictator games is altruism. A similar channel facilitates *trust* and *trustworthiness* in the trust game. However, envy and altruism, jointly, need to be invoked to explain the results of the public goods games with punishments. One difficulty that the FS model faces is in the explanation of some of the results from the dictator games that we reviewed in Chapter 5. This includes, in particular, the empirical finding of reluctant altruism; yet, these results are far less important in games that include a modicum of strategic concerns such as the ultimatum game.

Another great merit of the FS model is that it can simultaneously explain other-regarding behavior in bilateral or small group interaction and self-regarding behavior in competitive market settings. We explore these issues in Section 6.2.1 through an example of the ultimatum game by successively allowing for *proposer competition* and *responder competition*. Once we allow for such competition, the outcome of the ultimatum game with social preferences, under certain restrictions on the magnitude of altruism and envy, is identical to the outcome with self-regarding preferences. For instance, in the presence of responder competition, each responder is reluctant to

reject an unfair offer by the proposer because another responder might accept it. This allows the proposer to get away with offering the minimum possible amount. Will the proposer wish to offer the minimum amount? The answer is affirmative, if the altruism of the proposer is low enough.

A second model of inequity aversion that we consider, in Section 6.3, is the *equity, reciprocity, and competition* (ERC) model of Bolton and Ockenfels (2000). In this model, players care about the absolute magnitude of their payoff, and their payoff relative to the total payoff of others, but do not care about the distribution of payoffs of others. In two-player games, such as ultimatum and dictator games, the ERC model and the FS model often yield similar predictions.² In games with three or more players, however, the predictions of the two models may differ substantially. The reason is that in the ERC model, a player does not care about the payoffs received by each of his opponents individually as long as the average payoff remains constant. For example, in the public goods game with punishment, ERC predicts that a contributor is indifferent whether to punish a free-rider or another contributor, because punishing either of them affects the average payoff in the same way, while FS predicts that the free-rider will be punished.

The predictive ability of many theoretical models of self-regarding preferences hinges critically on answering the following question: Given two income distributions P and Q , which does the individual prefer? In a major advance, Rothschild and Stiglitz (1970, 1971) showed that the relevant concepts are *first and second order stochastic dominance*. However, these concepts are not sound for individuals who have social preferences, such as those of the Fehr–Schmidt form. In Section 6.4, we outline newer statistical concepts that are appropriate for individuals with Fehr–Schmidt preferences (Dhami and al-Nowaihi, 2013). We also clarify the relation of these newer dominance concepts with Lorenz dominance, which should facilitate applications of social preferences in the income distribution literature.

The FS model fits well the evidence on the behavior of subjects in experiments on societal redistribution. This evidence suggests that explanations for societal redistribution could be underpinned by social preferences and the innate desire of humans to be fair. By contrast, the existing theoretical literature on political economy assumes self-regarding preferences. Section 6.5 explores the implication of other-regarding preferences for understanding societal redistribution in a simple general equilibrium model with endogenous labor supply (Dhami and al-Nowaihi, 2010a,b). Two main results are established in this section. First, the conditions for the existence of a *Condorcet winner* under social preferences are natural and intuitive. Second, when there is a mixture of self-regarding and other-regarding voters, then the outcomes can be counter intuitive. Even if self-regarding voters are in a majority, the preferred policy of the other-regarding voters may prevail, i.e., be a Condorcet winner, and vice versa. An important determinant of the eventual redistributive outcome is the skill distribution and the inequality aversion of other-regarding voters.

Section 6.6 considers the Walrasian competitive equilibrium and the fundamental welfare theorems when consumers have other-regarding preferences (Dufwenberg et al., 2011). Using a separability condition on preferences, we establish a link between the Walrasian equilibria of an economy with other-regarding preferences and a related economy with self-regarding preferences. In general, the fundamental welfare theorems need not hold. But with an appropriate *social monotonicity condition*, restrictions on preferences to *well-being externalities*, and lump-sum transfers, we establish the analogue of the second fundamental welfare theorem. Interestingly,

² An interesting exception is a variant of the ultimatum game in which the proposer receives ϵ and the responder receives 0 if the proposer's offer is rejected. In this game, ERC predicts that any positive offer will be accepted because the relative payoff position of the responder is always worse than if he accepts the offer. In contrast, the FS model predicts that unfair offers are rejected.

there is a close link between some of these results and the simple general equilibrium model of redistribution that we consider in Section 6.5, above.

Section 6.7 considers the evidence on models of fairness and reciprocity. Formal models of reciprocity require the apparatus of *psychological game theory* that is postponed to Part 4 of the book. In Section 6.7.1, we consider a range of evidence that incorporates a desire for efficiency and surplus maximization over inequity aversion (Charness and Rabin, 2002; Engelmann and Strobel, 2004). However, evidence has also accumulated for the opposite view, in favor of inequity aversion (Bolton and Ockenfels, 2006). On account of their training, economics students may bias results in favor of efficiency, while inequity aversion remains fairly salient for non-economics students (Fehr et al., 2006).³ The evidence against inequity aversion has typically come from dictator games. However, in games where each party has an economically meaningful action, there is support for models of inequity aversion. For instance, in a gift exchange game, when workers have a chance to maximize the surplus in the relationship at a personal cost, they typically do not do so (Fehr et al., 1993, 1998). If surplus maximization were the overriding concern then, in gift exchange experiments, for any wage offered by the firm, we should always observe the worker putting in the maximum level of effort. But, as we saw in Chapter 5, what we observe is a positive relation between effort and wages (gift exchange).

In models of social preferences, such as Fehr and Schmidt (1999) and Bolton and Ockenfels (2000), intentions of other players do not matter. By contrast, intentions seem important to humans, e.g., the law makes a distinction between murder and manslaughter. Section 6.7.2 considers the experimental evidence for the role of intentions (Falk et al., 2008). This evidence rejects models of pure social preferences and models of pure reciprocity. It seems that the evidence is best described by models that combine the role of intentions, social preferences, and reciprocity such as the model by Falk and Fischbacher (2006).

Section 6.7.2 considers some caveats to the evidence on intentions. First, the early literature did find it hard to discover clinching evidence of intentions that could not be explained by other-regarding preferences (Blount, 1995; Charness, 2004; Bolton et al., 1998). Second, more recent evidence shows that in some cases, other-regarding preferences explain the evidence better than intentions. However, in other cases, both intentions and other-regarding motives are unable to explain the evidence (Bolton and Ockenfels, 2005). In this case, one may need to explore other avenues, possibly *procedural theories of fairness*, fairness norms, or bring into consideration reference points for fairness.

Section 6.8 considers several *hybrid models* that combine elements of social preferences and reciprocal intentions. Some models combine intentions with social preferences and other social objectives such as efficiency (Charness and Rabin, 2002). In models of *interdependent preferences*, people favor others who are perceived to be similarly fairness-minded (Levine, 1998; Cox et al., 2007). These models analyze intentions in a reduced form manner and do not employ the tool of psychological game theory. The section concludes with a note on the axiomatic foundations of social preference models (Segal and Sobel, 2007; Neilson, 2006).

6.2 The Fehr–Schmidt model

Fehr–Schmidt (1999), henceforth FS, develop a model of social preferences that simultaneously explains both fairness induced behavior in small group interaction and self-regarding behavior

³ In Charness and Rabin (2002) and Engelmann and Strobel (2004), the predominant subject pool is economics students.

in market settings. They also find that the presence of a small fraction of self-regarding individuals within a group of individuals who have social preferences (and vice versa) has important implications for the observed outcome.

Formally, consider a set of n individuals, $N = \{1, 2, \dots, n\}$. For any player $j \in N$ define the set $N_{-j} = N - \{j\}$. N_{-j} is called the reference set with respect to individual j . The respective monetary payoffs of the individuals are $y_1, y_2, \dots, y_n$. The payoffs of all individuals other than i are denoted by the vector $\mathbf{y}_{-i}$. For any $i \in N$ the Fehr-Schmidt utility function of an individual (or FS utility function) is given by

$$U_i(y_i, \mathbf{y}_{-i}; \alpha_i, \beta_i) = y_i - \frac{\alpha_i}{n-1} \sum_{j \neq i} \max\{y_j - y_i, 0\} - \frac{\beta_i}{n-1} \sum_{j \neq i} \max\{y_i - y_j, 0\}, \quad (6.1)$$

where $\alpha_i \geq 0$ and $0 \leq \beta_i < 1$. Most experimental evidence indicates that $\beta_i \in [0, 1)$, $\beta_i < \alpha_i$. An individual is said to have *social preferences* or *other-regarding preferences*, if at least one of β_i and α_i is non-zero.

The FS utility function has several important features. First, inequity is modeled as *self-centered inequity aversion*, i.e., one's own income becomes the reference point with respect to which incomes of others are compared. Furthermore, inequity among other members of the reference group, i.e., among N_{-i} , is not taken into account by individual i . Second, the theory is silent on the composition and size of the reference group, N_{-i} . Third, the FS utility function reduces to that of a self-regarding individual when $\beta_i = \alpha_i = 0$. Fourth, although (6.1) restricts itself to the linear payoffs case, the theory allows U_i to be non-linear in the payoffs. Fifth, intentions, as in Rabin (1993), or average income, as in Bolton and Ockenfels (2000), do not play a role. Sixth, $\beta_i < 0$, on account of individuals who might derive pleasure from being better off than others (e.g., status-seeking individuals), is ruled out. However, in most cases this feature does not affect the results considered by FS. Seventh, while there is an upper bound on β_i there is none on α_i .⁴ Eighth, $\beta_i < \alpha_i$ implies that the player suffers more from *disadvantageous inequity* (the second term in (6.1)) rather than *advantageous inequity* (the third term in (6.1)). Ninth, payoffs arising from advantageous and disadvantageous inequality are normalized to take account of the number of players, hence, the second and third terms in (6.1) are divided by $n - 1$.

Unlike models that incorporate only altruism or envy (see Examples 6.1, 6.2 above), the FS model incorporates both concerns. Individuals with such preferences, therefore, care about how unequal their incomes are, relative to those who are richer and those who are poorer. For this reason, the class of such models is also sometimes known as *models of inequity aversion*. Table 6.1 gives the reader an indication of the magnitudes of α, β from various studies reported in Eckel and Gintis (2010). The estimates suggest that $\beta < 1$ and typically $\beta < \alpha$. Most tests of inequity averse preferences have been done on aggregate level data. Using a within-subjects design, Blanco et al. (2011) report differences in the estimated parameters of inequity aversion at the aggregate level and the individual level. Fisman et al. (2007) report substantial heterogeneity at the level of individuals in social preferences. For a debate on the calibration of the Fehr-Schmidt model, see the 2010 issue of the *Journal of Economic Behavior and Organization*, Volume 73.

⁴ To see why β_i is restricted to be less than one for all i , consider the following simple example. Suppose that the incomes are arranged in increasing order, so $y_n = \max\{y_1, y_2, \dots, y_n\}$ and $\beta_n \geq 1$. Then (6.1) implies that

$$U_n(y_n, \mathbf{y}_{-n}; \alpha_n, \beta_n) = (1 - \beta_n)y_n + \frac{\beta_n}{n-1} [y_1 + y_2 + \dots + y_{n-1}].$$

If $\beta_n \geq 1$ then the RHS is non-increasing in y_n , so (provided this does not change the income rankings) throwing away a dollar increases the payoff of player n , which seems unreasonable.

Table 6.1 Parameter estimates of the Fehr and Schmidt (1999) model of inequality aversion.

Study and population	Mean parameter values	
	α	β
Fehr and Schmidt (1999)		
Calibration (calculated from reported distributions)	0.85	0.315
Standard deviation	(1.119)	(0.252)
Bellemare, Kroger, and van Soesta (2008)		
Population sample, Dutch adults	1.892	0.801
	(0.640)	(0.921)
High education and below 35	1.018	−0.271
	(0.325)	(0.503)
Goeree and Holt (2000) ^b		
Student proposers	0.84	0.66
	(0.160)	(0.080)
Student responders	0.84	0.12
	(0.160)	(0.020)
Ellingsen and Johannessen (2004) ^c		
Student subjects	0.31	0.36
95% Confidence Interval	(0–1.2)	(0.25)–0.5
Daruvala (2009) ^d		
Student subjects	0.952	0.344
Standard deviation	(1.039)	(0.294)

Standard errors in parentheses, unless otherwise indicated. (a) Specification included variables to measure intentions, and is embedded in a logit decision framework; (b) Specification included logit decision framework; a common parameter was estimated; (c) Estimated from graphs; (d) Mean and standard deviation value calculated by the author from individual data. Source: Eckel and Gintis (2009).

The model described in (6.1) is the linear version of the FS model. The non-linear version has also been used to explain important phenomena such as societal redistribution. See for instance, Dhami and al-Nowaihi (2010a,b), which is discussed in Section 6.5 below.

6.2.1 Fairness and market competition

We now demonstrate how the FS model can explain the anomalies generated by the ultimatum game. Denote the proposer and the responder in the ultimatum game as, respectively, players 1 and 2. The share received by the responder is denoted by s , while the proposer's share is $1 - s$.

Lemma 6.1 *The responder accepts all offers $s \geq 0.5$. There is a critical share, $s_c < 0.5$, such that the responder (1) rejects all offers $s < s_c$, and (2) accepts all offers $s \geq s_c$.*

Proof of Lemma 6.1: Suppose that $s \geq 0.5$. From (6.1), the responder's utility is $U_2 = s - \beta_2 [s - (1 - s)] = \beta_2(1 - s) + s(1 - \beta_2)$, which is positive because $\beta_i \in [0, 1)$. Hence, the responder will always accept $s \geq 0.5$. Now suppose that $s < 0.5$. From (6.1), the respon-

der's utility is $U_2 = s - \alpha_2 [(1-s) - s] = s(1 + 2\alpha_2) - \alpha_2$. For this to be positive we require that

$$s \geq s_c = \frac{\alpha_2}{1 + 2\alpha_2}, s_c < 0.5. \quad (6.2)$$

Hence, the responder will accept all offers $s \geq s_c$ and reject all offers $s < s_c$. ■

The critical threshold, s_c , is increasing and strictly concave in α_2 . In the limiting case $\alpha_2 \rightarrow \infty$, $s_c = 0.5$.

Proposition 6.3 *The equilibrium share offered by the proposer to the responder is given by*

$$s^* = \begin{cases} s_c & \text{if } \beta_1 < 0.5 \\ 0.5 & \text{if } \beta_1 > 0.5 \\ s \in [s_c, 0.5] & \text{if } \beta_1 = 0.5 \end{cases}.$$

Proof of Proposition 6.3: It is simple to check that the proposer will not offer a share $s > 0.5$ to the responder. From Lemma 6.1 we know that the responder will accept any share $s : s_c \leq s \leq 0.5$. Let us consider such a share, s . From (6.1), the proposer's utility is $U_1 = (1-s) - \beta_1 [(1-s) - s]$, so $\frac{\partial U_1}{\partial s} = 2\beta_1 - 1$. Thus, if $\beta_1 < 0.5$, then $\frac{\partial U_1}{\partial s} < 0$, so the proposer should offer the minimum share that will be accepted by the responder, i.e., $s^* = s_c$. If $\beta_1 = 0.5$, then $\frac{\partial U_1}{\partial s} = 0$, so any feasible share $s^* \in [s_c, 0.5]$ that will be accepted is offered. If $\beta_1 > 0.5$, then $\frac{\partial U_1}{\partial s} > 0$, so we have a corner solution, $s^* = 0.5$. ■

The predictions of the theory seem to match the stylized facts reasonably well in Western societies. First, there are rarely any offers above 0.5. Second, offers of 0.5 are readily accepted and low offers (presumably below s_c) are rejected.

It is possible that one (or both) of the parties has asymmetric information about the preference parameters of the other party. For instance, suppose that the proposer has asymmetric information about the responder's disadvantageous inequity parameter, α_2 . In particular, suppose that the proposer knows that α_2 has support $[\underline{\alpha}, \bar{\alpha}]$ and a distribution function $F(\alpha_2)$. From Lemma 6.1, we know that the offer will only be accepted by the responder if $s \geq s_c$. So, using (6.2), $\Pr(s \geq s_c) \Leftrightarrow \Pr\left(\alpha_2 \leq \frac{s}{1-2s}\right) = F\left(\frac{s}{1-2s}\right)$, which is increasing in s . Hence, the probability that an offer is accepted by the responder is increasing in the offer, s . This is also consistent with the empirical evidence.

The results derived above change dramatically when there is one-sided market competition either among proposers or among respondents. This is well known in experimental economics. For instance, Davis and Holt (1993) make six concluding remarks to their book. In the very first, they state that *competitive market predictions have been observed in a rich variety of circumstances and in literally hundreds of instances*. Furthermore, they remark that experimental behavior approximates the competitive outcome even in the absence of the limiting cases (i.e., a large number of traders or complete information). These insights also apply to the predictions of the FS model for the ultimatum game as we now see.

MARKET GAMES WITH PROPOSER COMPETITION

Consider $n > 2$ players, where the first $n-1$ players are proposers and the n th player is the single responder in an ultimatum game. The proposers simultaneously make an offer to the responder. Each proposer's offer takes the form of a share $s \in [0, 1]$ that is offered to the responder. Denote

the maximum offer among the proposers by $\bar{s} = \max \{s_i\}$, $i \in N/n$. If $\bar{s}$ is unique, then the corresponding proposer gets to make that offer, with probability 1, to the responder. If, however, $k > 1$ proposers make the same offer, $\bar{s}$, then each of these proposers makes an offer to the responder with probability $1/k$. If the responder rejects the offer, everybody gets zero. If the responder accepts an offer s from the j th proposer, then their respective payoffs are $s, 1 - s$ and everybody else gets zero.

We first consider the benchmark result when all players have self-regarding preferences.

Lemma 6.2 *Suppose that all individuals are self-regarding. In any subgame perfect outcome, at least two proposers must offer $\bar{s} = 1$, which is accepted by the responder.*

Proof of Lemma 6.2: Suppose to the contrary that there is some $s : 0 < s < 1$ that is an equilibrium. Any of the proposers can now offer $s + \epsilon$, $\epsilon > 0$, which will be accepted by the responder and improves the payoff of the proposer. Furthermore, there cannot be a single proposer who offers $\bar{s} = 1$ in equilibrium. If this is the case, then let the next best offer by any proposer be $\tilde{s}$. The original proposer can now lower his offer to any number in $(\tilde{s}, 1)$. Such an offer is accepted by the responder and improves the payoff of the original proposer. Hence, the unique equilibrium offer is $\bar{s} = 1$, which is made by at least two proposers. ■

In the face of proposer competition, the responder reaps all gains from trade. These results are consistent with the experimental evidence in Roth et al. (1991).

What is the equilibrium outcome when all n players have social preferences? It turns out that the outcome is identical to that obtained under self-regarding individuals. We start by proving an intermediate result that is similar to Lemma 6.1.

Lemma 6.3 *The other-regarding responder will accept any offer $s \geq 0.5$.*

Proof of Lemma 6.3: For $s \geq 0.5$, the utility of the responder from accepting the offer is

$$U_n = s - \frac{\beta_n}{n-1} [s - (1-s)] - \frac{(n-2)\beta_n}{n-1} (s-0). \quad (6.3)$$

Since $0 < \beta < 1$, so $U_n \geq 0$. Since, $U_n \geq 0$ if an offer $s \geq 0.5$ is accepted and $U_n = 0$ if the offer is rejected, hence, the offer is accepted. ■

Proposition 6.4 *Suppose that all agents have other-regarding preferences given in (6.1). In the ultimatum game with proposer competition, the unique subgame perfect equilibrium outcome is for at least two proposers to offer $s = 1$, which is accepted by the responder. All gains from trade accrue to the responder.*

Proof of Proposition 6.4: Using Lemma 6.3, the responder will accept $s = 1$. We want to show that there cannot be an equilibrium in which $0.5 \leq s < 1$ is offered and accepted by the responder. Suppose that $\hat{n}$ proposers, such that $1 \leq \hat{n} < n-1$, offer $0.5 \leq \bar{s} < 1$ in equilibrium and their offer is accepted. Order, without loss of generality, the proposers as follows $1, 2, \dots, \hat{n}, \hat{n}+1, \dots, n-1$. The proposers $\hat{n}+1, \dots, n-1$ offer less than $\bar{s}$, say, 0 (this does not alter the proof). Now choose any proposer i from among $\hat{n}+1, \dots, n-1$. If he offers some $s_i < \bar{s}$, then proposer i does not win the contest and there are two other players who have a greater payoff: the proposer who wins and the responder, hence, using (6.1), his payoff is

$$U_i(s_i) = 0 - \frac{\alpha_i}{n-1} [(1-\bar{s}-0) + (\bar{s}-0)] = -\frac{\alpha_i}{n-1}. \quad (6.4)$$

If proposer i instead offers $\bar{s} + \epsilon$, $\epsilon > 0$, then he wins the contest for sure. In this case one player, the responder has a greater payoff than him (on account of Lemma 6.3), while $n - 2$ players have a lower payoff. Using (6.1), his payoff is

$$U_i(\bar{s} + \epsilon) = [1 - (\bar{s} + \epsilon)] - \frac{\alpha_i}{n-1} [(\bar{s} + \epsilon) - [1 - (\bar{s} + \epsilon)]] - \frac{\beta_i(n-2)}{n-1} [1 - (\bar{s} + \epsilon)]. \quad (6.5)$$

From (6.4), (6.5), it is easily checked that $U_i(\bar{s} + \epsilon) - U_i(s_i) > 0$. Hence, we derive a contradiction, so any $s < 1$ cannot be an equilibrium offer.

Now let $\hat{n} = n - 1$ and an offer $0.5 \leq \bar{s} < 1$ that is accepted (see Lemma 6.3). Each proposer then wins the contest with probability $\frac{1}{n-1}$. Any of the proposers can offer an amount $\bar{s} + \epsilon$, ϵ sufficiently small and be ensured to win with probability 1. This improves the payoff of the proposer and contradicts the original equilibrium.

Hence, any $s < 1$ cannot be an equilibrium offer. ■

MARKET GAMES WITH RESPONDER COMPETITION

Now consider the possibility of responder competition. Suppose that the first player is the proposer, while players $2, 3, \dots, n$ are responders. The proposer makes an offer, followed by the responders' decisions to accept or reject. Among the set of responders who accept the offer, one is randomly chosen with uniform probability. The payoffs of the proposer and the winning responder are $1 - s$, s , while all others get zero. If no responder accepts, then all players get zero.

Lemma 6.4 *Suppose that all players are self-regarding. In any subgame perfect outcome, the proposer offers $\bar{s} = 0$, which is accepted by at least one responder.*

Proof of Lemma 6.4: Analogous to the proof of Lemma 6.2, so it is omitted. ■

There is less experimental evidence for market games with responder competition, relative to those with proposer competition. In ultimatum games with two responders, if any responder rejects an offer, the other responder might accept it. Hence, punishing the proposer for low offers is more difficult. One should expect that the same offers are less likely to be rejected in the presence of responder competition.

These predictions are borne out in Fischbacher et al. (2009) where subjects were randomly rematched every period to ensure that the game remained a one-shot interaction. When a single proposer is matched with a single responder, the responder gets, on average, about 40%. In the case of two responders, the share is about 14% lower, and in the final period, it dips below 20%. When there are five responders, the responder's share is about 10% in the last 10 rounds of a 20 round experiment.

An important implication of these results is that if individuals have self-regarding preferences, then adding another responder (competition) should have no impact on the outcome and the proposer receives the entire surplus. Thus, the introduction of competition is neutral in terms of the outcome. However, the introduction of competition when agents have social preferences alters the outcomes significantly. Hence, the self-interest model understates the role of competition. The experimental results of Güth et al. (1997) also bear out these results.

The next proposition shows that in the presence of social preferences, the outcome under responder competition is identical to that under self-regarding preferences. We leave the proof as an exercise for the reader.

Proposition 6.5 *Let the advantageous inequality parameter of the proposer, β_1 , be bounded above by $(n-1)/n$. Suppose that we limit attention to $s < 0.5$. Then there exists a subgame perfect outcome in which each responder accepts any $0 \leq s < 0.5$ and the proposer offers $s = 0$.*

This is not the only equilibrium of the model; an exercise at the end of the chapter asks for other equilibria.

Propositions 6.4, 6.5 show that in the presence of competition, the preferences of individuals, self-regarding or other-regarding, do not matter for the results. Why? The reason is that each individual is too small to influence the degree of inequality in society. Hence, this makes the pursuit of individual self-interest more beneficial (through the first term in (6.1)), while at the same time turning the inequality to his advantage, which is preferable because $\beta_i < \alpha_i$. When individuals have a greater opportunity to influence the degree of inequality, e.g., by altering the level of effort in a principal-agent relationship, then other-regarding preferences have greater bite, as we shall see later.

An important implication is that if individuals are “small,” then fairness issues are less important in the market for goods (as Propositions 6.4, 6.5 exemplify) but of greater importance in labor markets where issues of endogenous effort in firms are rather more salient.

6.2.2 Fairness and cooperation

We now apply the FS preferences to an explanation of the results from public goods games. The implications of these results go beyond the public goods game. In multi-stage games with self-regarding players, cooperation is difficult to achieve under full information, because the tendency for opportunistic behavior in the last stage unravels all the way back to the initial stage. However, if players have other-regarding preferences, then they might not engage in opportunistic behavior in the last stage because of inequity aversion. Furthermore, if some players have initially managed to increase their relative payoffs by behaving in an opportunistic manner, they might be punished by the other-regarding players, who desire greater equality of payoffs. This imparts credibility to the punishment of opportunistic behavior. Contrast this with the non-credibility of punishments in a finite multi-stage full information game with self-regarding players.

To examine the issues involved, we use the public goods game, described in Section 5.4. The monetary payoffs in the game without punishment and with punishment are given by (5.1), (5.2). The utility of an other-regarding individual i is of the FS form,

$$U_i = m_i - \frac{\alpha_i}{n-1} \sum_{j \neq i} \max \{m_j - m_i, 0\} - \frac{\beta_i}{n-1} \sum_{j \neq i} \max \{m_i - m_j, 0\}. \quad (6.6)$$

PUBLIC GOODS GAME WITHOUT PUNISHMENTS

Using (5.1), in the public goods game without punishments, $m_j - m_i = g_i - g_j$. Thus, using (6.6), the utility of player $i \in N = \{1, 2, \dots, n\}$ is

$$U_i = y - g_i + rG - \frac{\alpha_i}{n-1} \sum_{j \neq i} \max \{g_i - g_j, 0\} - \frac{\beta_i}{n-1} \sum_{j \neq i} \max \{g_j - g_i, 0\}. \quad (6.7)$$

Lemma 6.5 *A sufficient condition for an other-regarding player with preferences given in (6.7) to choose $g_i^* = 0$ is that $r + \beta_i < 1$.*

Proof of Lemma 6.5: From (6.7), by contributing an extra unit to the public good, the monetary cost to an individual is $1 - r$. To see the maximum possible monetary benefit from this action, suppose that, among the set of contributors, this agent contributes the least. From (6.7), by contributing an extra unit, the benefit to this lowest contributor is a reduction in advantageous inequity by $\frac{\beta_i(n-1)}{n-1} = \beta_i$. Hence, the maximum possible benefit to player i is $\beta_i - (1 - r)$, which is negative if $r + \beta_i < 1$. ■

Proposition 6.6 *Let $\tilde{n} \in \{0, 1, \dots, n\}$ be the number of other-regarding players with $r + \beta_i < 1$. If $\tilde{n} \geq \frac{r}{2}(n - 1)$, then there is a unique equilibrium in Nash contributions such that $g_i^* = 0$ for all i .*

The proof of Proposition 6.6 is left to the reader as an exercise. From Lemma 6.5, $r + \beta_i < 1$ guarantees free-riding. Proposition 6.6 shows that if a lower bound is put on the number of free-riders, then all other-regarding players do not contribute. If the number of free-riders is large enough, in the sense specified in Proposition 6.6, then contributors to the public good would suffer too much disadvantageous inequity from contributing. Hence, they do not contribute. In this case, the result in Proposition 6.6 matches with the empirical evidence that there is little voluntary contribution in public goods games without punishments. There are, however, other equilibria in which a fraction of the agents contribute a positive amount; an exercise asks you to show this.

PUBLIC GOODS GAME WITH PUNISHMENTS

The possibility of punishment alters the results of public goods games with other-regarding individuals quite dramatically. Suppose that there are some free-riders in the first stage. Then first stage contributors, who are other-regarding, have an incentive to punish the free-riders in the second stage. Contributors in the second stage face disadvantageous inequity relative to the free-riders, because the latter get exactly the same level of benefit from the public good but have not paid for it. How much do the contributors care about this inequity? The answer depends on their inequity aversion parameter α_i . If this is large enough, then a contributor can make a credible threat to punish the free-riders. If the punishment has sufficient bite for the non-contributors, then it can reduce or eliminate free-riding, which is in line with the experimental evidence. These ideas are formalized below.

Definition 6.6 *A conditionally cooperative enforcer is characterized by two conditions: (1) $r + \beta_i > 1$, and (2) a willingness to cooperate if others cooperate.*

Proposition 6.7 *Suppose that there are $n_c \in \{0, 1, \dots, n\}$ conditionally cooperative enforcers. Order the players such that the first n_c are conditional cooperators, $1, 2, \dots, n_c, n_c + 1, \dots, n$. The remaining players are self-regarding, i.e., $\alpha_j = \beta_j = 0, j = n_c + 1, \dots, n$. Let*

$$\frac{\alpha_i}{(n-1)(1+\alpha_i) - (n_c-1)(\alpha_i + \beta_i)} > c, \quad (6.8)$$

where c is the cost per unit to punish another player; $\frac{1}{c}$ is also known as the efficiency of punishment (see Definition 5.4). The following constitutes a subgame perfect equilibrium. In the first stage, all individuals contribute an identical amount $g_i = \bar{g} \in [0, y]$. The second stage strategies are as follows. If all individuals contribute $\bar{g}$ in the first period, then there are no punishments. If player $j \in \{n_c + 1, \dots, n\}$ contributes $g_j < \bar{g}$, then each player $i \in \{1, 2, \dots, n_c\}$ inflicts a punishment $p_{ij} = (\bar{g} - g_j)/(n_c - c)$. There is no punishment inflicted by any players outside this set. If any of the conditionally cooperative enforcers $i \in \{1, 2, \dots, n_c\}$

contributes $g_i < \bar{g}$ or any player chooses $g > \bar{g}$, then a Nash equilibrium of the punishment game is played.

Proof of Proposition 6.7: Suppose that some player $j \in \{n_c + 1, \dots, n\}$ contributes $g_j < \bar{g}$ and all other players follow their equilibrium action. Then, given the punishment strategies outlined in the proposition, in conjunction with the monetary payoff given in (5.2), it can be checked that

$$m_j = y - g_j + r[(n-1)\bar{g} + g_j] - n_c \left(\frac{\bar{g} - g_j}{n_c - c} \right),$$

and for any player $i \in \{1, 2, \dots, n_c\}$,

$$\begin{aligned} m_i &= y - \bar{g} + r[(n-1)\bar{g} + g_j] - c \left(\frac{\bar{g} - g_j}{n_c - c} \right), \\ \Rightarrow m_i &= y - \left(\frac{n_c - c}{n_c - c} \right) (\bar{g} - g_j + g_j) + r[(n-1)\bar{g} + g_j] - c \left(\frac{\bar{g} - g_j}{n_c - c} \right), \\ \Rightarrow m_i &= y - g_j + r[(n-1)\bar{g} + g_j] - \left(\frac{\bar{g} - g_j}{n_c - c} \right) (n_c - c + c) = m_j. \end{aligned}$$

Thus, the deviator gets the same payoff as the enforcer. Because the non-deviating, non-enforcers, do not punish they get the highest payoff. Hence, it is better for a deviating non-enforcer to be a non-deviating, non-enforcer.

We now need to check two conditions.

I. Each of the enforcers finds it worthwhile to carry out the stated punishment: Suppose that an enforcer $i \in \{1, 2, \dots, n_c\}$ reduces punishment of some deviating player j by an amount $\epsilon > 0$ from the stated level $p_{ij} = (\bar{g} - g_j)/(n_c - c)$. This has three main effects.

1. Relative to the other non-deviating, non-enforcers, $n_c + 1, \dots, n$, enforcer i increases their own payoff by $c\epsilon$ (the effective cost of the punishment) for a benefit in disadvantageous inequity of

$$\frac{\alpha_i}{n-1} [n - (n_c - 1)] c\epsilon. \quad (6.9)$$

2. Relative to the other $n_c - 1$ enforcers who stick to the stated level of punishment, the payoff of a deviating enforcer goes up. This results in an additional cost to enforcer i in terms of advantageous inequity of

$$\frac{\beta_i (n_c - 1)}{n - 1} c\epsilon. \quad (6.10)$$

3. Relative to the defector, the payoff of player i has gone up by ϵ on account of the reduction in punishment. So player i finds that disadvantageous inequity has gone up by an amount

$$\frac{\alpha_i}{n-1} [\epsilon - c\epsilon]. \quad (6.11)$$

Subtracting (6.10), (6.11) from (6.9) and simplifying, we get the required condition (6.8).

II. Each of the enforcers finds it worthwhile to contribute $\bar{g}$: Suppose that one of the conditional enforcers $i \in \{1, 2, \dots, n_c\}$ reduces his contribution of the public good from the stated level $\bar{g}$ to $\bar{g} - \eta$, $\eta > 0$. Ignore punishments in the second stage for the moment. This increases the monetary payoff by $(1 - r)\eta$ but creates more advantageous inequality (a cost) with respect to all other players by an amount $\frac{\beta_i(n-1)}{n-1}\eta$. The net effect is $\eta(1 - r - \beta_i)$ which, using Definition 6.6, is negative for a conditionally cooperative enforcer. If we include second stage punishments, the net effect is even more negative. Hence, it does not pay to reduce contributions from the stated amount. ■

6.3 The ERC model

Bolton and Ockenfels (2000) propose a model of *equity, reciprocity, and competition* (ERC), a model of social preferences that is closely related to a model in Bolton (1991).⁵ ERC applies to both extensive form and normal form games. Suppose that there are n players, $i = 1, 2, \dots, n$. Payoffs to all players are non-negative monetary payoffs. Denote by $x_i \geq 0$, the payoff of player i . The sum of the total payoffs is $S = \sum_{i=1}^n x_i$. The relative payoff of player i , also called the *motivation factor*, is given by

$$r_i = r_i(x_i, S, n) = \begin{cases} x_i/S & \text{if } S > 0 \\ \frac{1}{n} & \text{if } S = 0 \end{cases} . \quad (6.12)$$

The utility function (the authors prefer the term, *motivation function*) of player i is $u^i = u^i(x_i, r_i)$. Using (6.12), it is given by

$$\begin{cases} u^i(x_i, \frac{x_i}{S}) & \text{if } S > 0 \\ u^i(x_i, \frac{1}{n}) & \text{if } S = 0 \end{cases} . \quad (6.13)$$

The assumptions on the utility function are as follows.

- A1: u^i is twice continuously differentiable in x_i, r_i .
- A2: u^i is increasing and concave in x_i ($u^i_1 \geq 0$, $u^i_{11} \leq 0$).
- A3: For any given monetary payoff, the utility function u^i is maximized at $r_i = 1/n$ and is strictly decreasing and strictly concave in r_i around this point.

Assumption A3 builds into the model the importance of an equal division of the social surplus (the authors term equal division as the *social reference point*). The crux of the model is the trade-off created by assumptions A2 and A3. From A2, the individual prefers to increase her monetary payoff. However, from A3, for any monetary payoff, the individual prefers equal division. Population heterogeneity is introduced by how different individuals view this trade-off.

ERC and FS preferences do not nest each other, and nor do their empirical predictions. ERC is able to explain several stylized facts from experimental results, among them, the following. Why do proposers give more in ultimatum, as compared to dictator games? Why is there a positive association between effort and rents in the gift exchange game? In dictator games, ERC predicts that the proposer will keep between half and all of the money. It also predicts that zero offers will be rejected in ultimatum games and the rejection rates will increase as the proposer's offer

⁵ For an understanding of the evolution of this model see the section on bargaining games in Chapter 12. For an earlier published motivation behind the ERC model, see Bolton and Ockenfels (1998).

falls. The ERC model also explains the competitive outcomes under proposer competition in an ultimatum game. The ERC model allows for asymmetric information, and it can also explain the data in three-player ultimatum games where a passive third party appears to be treated unfairly, relative to the predictions of the Fehr–Schmidt model (Bolton and Ockenfels, 1998).

By contrast, there are other predictions of the ERC theory that are at variance with the evidence. In the ERC model, individuals compare their payoff with the aggregate payoff of others. In this sense, ERC has a less nuanced notion of inequality, relative to the FS model. In particular, individuals using ERC preferences do not care about interpersonal comparisons of their own payoff and the payoffs of others. The following experimental games show that this gives rise to different predictions relative to the FS model.

Suppose that a player has the mean payoff. Now consider a mean preserving spread of the payoff distribution. ERC predicts that the player's utility should be unaffected. However, the experimental results of Charness and Rabin (2002) do not support this prediction; utility falls with the mean preserving spread. By contrast, the FS utility of a player will be affected by such mean preserving changes.

The ERC model is not able to explain the outcomes in public goods games with punishments. In the second stage of this game, individuals with ERC preferences will not punish the non-contributors because they are not concerned with individual payoffs. Sensing this, potential non-contributors will not contribute. Hence, the prediction of the ERC model is that the outcome of the public goods games with and without punishments should be identical. This is strongly refuted (see Section 5.4.3). Antisocial punishments (see Section 5.4.6) are inconsistent not just with the ERC model but also many other models of social preferences. However, observed punishments in most experiments conducted in Western societies are prosocial. When the institution of punishment is enshrined via an upfront referendum, then punishments are not antisocial, nor are punishments antisocial when there is a norm for punishing non-cooperative behavior, or punishments elicit feelings of guilt and shame (Tyran and Feld, 2006; Herrmann et al. 2008; Carpenter et al., 2009). These kinds of punishments have relatively greater external validity, and the FS model explains them well.

The ERC model also finds it difficult to explain the results of games that demonstrate indirect reciprocity (see Section 5.6). Consider, for instance, the following variant of the third party punishment game. There are three players A, B, and C with respective endowments S , 0 , $S/2$. Players A and B play a dictator game with A being the first mover (so B has no meaningful actions to take). Player C can punish player A if the offer of player A to B is considered inadequate. For each unit of money that player C spends on the punishment, the monetary payoff of player A is reduced by three units. The total surplus in the game is $S + S/2 = 3S/2$. The relative share of player C, in the absence of punishment, is $1/3$, thus, given assumption A3, the Bolton–Ockenfels model predicts no punishment. In most experiments, the majority of C players undertake some form of punishment, which is at variance with the ERC model.

Charness and Rabin (2002) consider a variant of the third party punishment game. Suppose that three players A, B, and C have respective shares given by (x_1, x_2, x_3) . Player C is now given a choice between the following two allocations: $(575, 575, 575)$ and $(900, 300, 600)$. Player C gets the fair share (a third) in both allocations. The Bolton–Ockenfels model predicts that player C will choose the second allocation because, among the two, it gives player C a higher material payoff (Assumption A2). The experimental evidence is that 54% of the subjects prefer the first allocation. Self-interest alone predicts that C will choose the second allocation, hence, we cannot conclude that the choice of the remaining 46% supports the Bolton–Ockenfels model.

Although the experimental evidence from the simple games considered above does not always support the Bolton–Ockenfels model, as no model of other-regarding preferences is supported in

all possible experiments, it is nevertheless an important milestone in the development of models of social preferences. It formalizes an important intuitive insight about how people may behave. The model may perhaps perform even better in more cognitively challenging games with a very large number of players, where one might only have a rough notion of one's position relative to the aggregate.

6.4 Fairness and stochastic dominance

Given two probability distributions over outcomes, it is critical to be able to answer the question of which distribution an individual prefers. Under self-regarding preferences, this question has been the cornerstone of many important advances in economic theory. The relevant machinery to deal with these issues is *first and second order stochastic dominance* and *generalized Lorenz dominance*.⁶ What can we say when individuals have social preferences?

Dhami and al-Nowaihi (2013) show that the standard concepts of *first order and second order stochastic dominance* as well as *generalized Lorenz dominance* are no longer valid when one has social preferences of the FS form. They propose alternative dominance concepts that are suitable for such social preferences. In this section we consider their work.

Consider n possible income classes. Income, y , of any member belongs to the set of possible incomes $Y = \{y_1 < y_2 < \dots < y_n\}$. Under any distribution function, P , over this set of incomes, the proportion of individuals with income y_i is denoted by $p_i \geq 0$, $\sum_{i=1}^n p_i = 1$. The *cumulative probability distribution* is given by $P_0 = 0$, $P_j = \sum_{i=1}^j p_i$, $j = 1, 2, \dots, n$. Denote by Π , the set of all such distributions over Y .

The *cumulative of the cumulative distribution* is given by $\tilde{P}_0 = 0$, $\tilde{P}_j = \sum_{i=1}^j P_i$, $j = 1, 2, \dots, n$. Let $\tilde{\Pi}$ be the set of all such distributions. The average, or mean, of $y_1, y_2, \dots, y_n$ under the distribution $P \in \Pi$ is

$$\mu_P = \sum_{i=1}^n p_i y_i.$$

Let $u(y_i)$ be the utility function of an individual who has self-regarding preferences and has the income y_i . Denote by $\mathbf{u}$ the class of all non-decreasing utility functions.

In this section we will work with the general form of Fehr–Schmidt preferences, sometimes known as the *non-linear form of Fehr–Schmidt preferences*. These preferences modify the linear form of FS preferences given in (6.1) as shown in the next definition. For axiomatic foundations of such preferences, see Section 6.8.4 below.

Definition 6.7 (*General form of FS utility*): Consider the income distribution, P and $u \in \mathbf{u}$. The general form of the Fehr–Schmidt utility function (FS) for an individual with income $y_j \in Y$ is given by

$$U(y_j, P) = u(y_j) - \beta \sum_{i=1}^{j-1} p_i [u(y_j) - u(y_i)] - \alpha \sum_{k=j+1}^n p_k [u(y_k) - u(y_j)], \alpha \geq 0, 0 \leq \beta < 1.$$

(6.14)

⁶ See Rothschild and Stiglitz (1970, 1971).

$$U(y_j, P) - U(y_j, Q) = \beta \sum_{i=1}^{j-1} (q_i - p_i) [u(y_j) - u(y_i)] + \alpha \sum_{k=j+1}^n (q_k - p_k) [u(y_k) - u(y_j)]. \quad (6.15)$$

For an individual with self-regarding preferences, $\alpha = \beta = 0$, so

$$U(y, P) = u(y). \quad (6.16)$$

The approach in the literature is to begin with a purely statistical definition of a dominance concept and then to establish its economic relevance (or *soundness*) by relating it to the relevant decision theory. In the discrete case, first and second order dominance are defined as follows.

Definition 6.8 (First order stochastic dominance): Let $P, Q \in \Pi$. Then P first order stochastically dominates Q ($P \succsim_1 Q$) if $P_j \leq Q_j$ for $j = 1, 2, \dots, n$. If, in addition, the inequality is strict for some j , then P strictly first order stochastically dominates Q ($P \succ_1 Q$).

Definition 6.9 We say that the income levels are equally spaced if for some positive real number, δ , $y_{i+1} - y_i = \delta$ for $i = 1, 2, \dots, n-1$.⁷

Definition 6.10 (Second order stochastic dominance): Suppose that incomes are equally spaced (Definition 6.9).⁸ Let $P, Q \in \Pi$. Then P second order stochastically dominates Q ($P \succsim_2 Q$) if $\tilde{P}_j \leq \tilde{Q}_j$ for $j = 1, 2, \dots, n$ and if $\mu_P \geq \mu_Q$.⁹ If, in addition, one of these inequalities is strict, then P strictly second order stochastically dominates Q ($P \succ_2 Q$).

Propositions 6.8 and 6.9 consider the classical case in which self-regarding individuals maximize their expected utility with respect to a given distribution, behind a veil of ignorance (Mas-Colell et al., 1995, ch. 6).

Proposition 6.8 Let $P, Q \in \Pi$. Then $P \succsim_1 Q$ if, and only if, for any $u \in \mathbf{u}$,

$$\sum_{i=1}^{i=n} p_i u(y_i) \geq \sum_{i=1}^{i=n} q_i u(y_i), \quad (6.17)$$

and $P \succ_1 Q$ if, and only if, the inequality in (6.17) is strict for all strictly increasing $u \in \mathbf{u}$.

Choosing $u(y) = y$ as a candidate function in Proposition 6.8 we get that if $P \succsim_1 Q$ ($P \succ_1 Q$), then the restrictions on the means under P and Q is $\mu_P \geq \mu_Q$ ($\mu_P > \mu_Q$).

⁷ This is not restrictive because we can always introduce extra income levels, each with probability zero, to achieve equal spacing. For example, suppose $y_1 = 5, y_2 = 7, y_3 = 11$, which is not equally spaced, and $P_1 = \frac{1}{3}, P_2 = \frac{2}{3}$. Consider $y_1 = 5, y_2 = 7, y_3 = 9, y_4 = 11$, which are equally spaced, and $Q_1 = \frac{1}{3}, Q_2 = \frac{2}{3}, Q_3 = \frac{2}{3}$. Both P and Q describe the same reality.

⁸ Often "equal spacing" is not stated as part of the definition of second order stochastic dominance. But then it has to be stated in the propositions. We find it more convenient to state this in the definition. This assumption is not needed in the continuous case.

⁹ Often the condition $\mu_P \geq \mu_Q$ is not included as part of the definition of second order stochastic dominance. But then it needs to be included in the statement of the relevant propositions. It is slightly simpler to include $\mu_P \geq \mu_Q$ as part of the definition of second order stochastic dominance.

Proposition 6.9 Let $P, Q \in \Pi$. Then $P \succsim_2 Q$ if, and only if, for any concave $u \in \mathbf{u}$,

$$\sum_{i=1}^{i=n} p_i u(y_i) \geq \sum_{i=1}^{i=n} q_i u(y_i), \quad (6.18)$$

and $P \succ_2 Q$ if, and only if, the inequality in (6.18) is strict for all strictly increasing strictly concave $u \in \mathbf{u}$.

Propositions 6.8 and 6.9 no longer hold under FS preferences once individuals know their income levels and the income levels of others.

Example 6.3 Consider three income levels, 0, 25, 50 and two distributions, P, Q defined in the table below, where $\frac{1}{6} \leq \varepsilon \leq \frac{5}{6}$.

y	$y_1 = 0$	$y_2 = 25$	$y_3 = 50$
$p(y)$	0	1/3	2/3
$q(y)$	1/6	ε	$(5/6) - \varepsilon$
$P(y)$	0	1/3	1
$Q(y)$	1/6	$(1/6) + \varepsilon$	1
$\tilde{P}(y)$	0	1/3	4/3
$\tilde{Q}(y)$	1/6	$(1/3) + \varepsilon$	$(4/3) + \varepsilon$

From the above table it can be easily seen that $P \succ_1 Q$. Also from the above table it is clear that $P \succ_2 Q$ ($\mu_P = \frac{125}{3} > \frac{125}{3} - 25\varepsilon = \mu_Q$). For the individual with income $y_2 = 25$, we can find from (6.15)

$$U(25, Q) - U(25, P) = \alpha \left(\varepsilon - \frac{1}{6} \right) [u(50) - u(25)] + \frac{\beta}{6} [u(0) - u(25)]. \quad (6.19)$$

It is easy to construct examples where an individual with Fehr–Schmidt preferences strictly prefers Q to P . For example, let $\varepsilon = \frac{1}{3}$ and take u to be any strictly increasing own utility function. Take $\alpha = 1$ and $\beta = 0$. Consider an individual with income 25. Then, from (6.19) we get $U(25, Q) - U(25, P) = \frac{u(50) - u(25)}{6} > 0$.

Consider now some implications of pride ($\beta < 0$) rather than altruism ($\beta > 0$). Let u be any strictly increasing own utility function. Take $\alpha = 1$ and $\varepsilon = \frac{1}{6}$. Then $U(25, Q) - U(25, P) = -\frac{\beta}{6} [u(25) - u(0)]$. Clearly the choice between P, Q now hinges solely on the sign of β . An individual with altruism prefers P to Q but an individual with pride prefers Q to P . This example makes the more general point that an individual with pride will always prefer the dominated distribution.

Example 6.4 In FS preferences, see (6.14), $\beta > 0$ captures altruism or compassion towards individuals who are poorer. Hopkins (2008) differentiates between the sign of β in two different literatures. In behavioral economics, where models of social preferences have gained increasing acceptance, $\beta > 0$. But in happiness economics, $\beta < 0$ is allowed for. Hopkins (2008) refers to $\beta < 0$ as “pride,” “competitiveness,” or “downward envy.” Consider now some implications of pride ($\beta < 0$) rather than altruism ($\beta > 0$) in the data given in Example 6.3. Let u be any strictly increasing own utility function. Take $\alpha = 1$ and $\varepsilon = \frac{1}{6}$. Then, from (6.19) we get that $U(25, Q) - U(25, P) = -\frac{\beta}{6} (u(50) - u(25))$. Clearly the choice between P, Q now hinges solely on the sign of β . An individual with altruism prefers P to Q but an

individual with pride prefers Q to P . Dhami and al-Nowaihi (2013) show formally that an individual with pride will always prefer the dominated distribution.

These examples beg the question: what are the appropriate dominance concepts for FS preferences? The next remark provides a clue to the answer.

Remark 6.1 Begin with an income distribution $P \in \Pi$, under which an individual has the Fehr–Schmidt preferences given in (6.14) and has income y_j . Consider two income levels, y_s, y_t , such that $y_s < y_t < y_j, p_s < 1$ and $p_t > 0$. Suppose that we obtain the distribution Q from P by transferring a fraction Δ of individuals from the income class y_t to y_s , where $0 < \Delta < \min \{1 - p_s, p_t\}$, then from (6.15) $U(y_j, P) - U(y_j, Q) = \beta \Delta [u(y_t) - u(y_s)] > 0$. Thus individuals with FS preferences dislike “rich to poor transfers” among people poorer than them. Now consider two income levels, y_l, y_m , such that $y_j < y_l < y_m, p_s < 1$ and $p_m > 0$. Now obtain a new distribution Q by transferring a fraction Δ of individuals from the income class y_m to y_l , where $0 < \Delta < \min \{1 - p_l, p_m\}$, then from (6.15) $U(y_j, Q) - U(y_j, P) = \alpha \Delta (u(y_m) - u(y_l)) > 0$. Thus, individuals with FS preferences like “rich to poor transfers” among people richer than them.

Consider the following example that illustrates some of the main ideas in Dhami and al-Nowaihi (2013). Consider the data given in the table below, where $0 < \epsilon < \frac{1}{6}$. We are interested in the individual with income of $y_3 = 50$ who has FS preferences.

y	$y_1 = 0$	$y_2 = 25$	$y_3 = 50$	$y_4 = 75$	$y_5 = 100$
$p(y)$	0	$1/3$	$(1/3) + 2\epsilon$	$(1/3) - 2\epsilon$	0
$q(y)$	ϵ	$(1/3) - \epsilon$	$1/3$	$(1/3) - \epsilon$	ϵ
$P(y)$	0	$1/3$	$(2/3) + 2\epsilon$	1	1
$Q(y)$	ϵ	$1/3$	$2/3$	$1 - \epsilon$	1
$\tilde{P}(y)$	0	$1/3$	$1 + 2\epsilon$	$2 + 2\epsilon$	$3 + 2\epsilon$
$\tilde{Q}(y)$	ϵ	$(1/3) + \epsilon$	$1 + \epsilon$	2	3

Below income $y_3 = 50$, distribution Q moves some richer individuals to a poorer income level and above income $y_3 = 50$ it move some poorer individuals to a richer income level. Thus, using Remark 6.1, we would expect such an individual to prefer the distribution P to the distribution Q . What inferences can we draw from this unambiguous preference for P ? An examination of the table reveals that the following restrictions apply

$$P_i \leq Q_i, P_k \geq Q_k; i < 3, k \geq 3. \quad (6.20)$$

$$\tilde{P}_i \leq \tilde{Q}_i, \tilde{P}_k \geq \tilde{Q}_k; i < 3, k \geq 3. \quad (6.21)$$

From (6.20), neither distribution first order stochastically dominates the other, and, from (6.21), neither distribution second order stochastically dominates the other. Thus, a priori, one cannot predict which distribution would be preferred by an individual with self-regarding preferences. But we shall see that these restrictions are ideal to determine which of the distributions is preferred by an individual with FS preferences whose income is $y_3 = 50$. One can make the following simple calculations using (6.15), for any $u \in \mathbf{u}$,

$$U(50, P) - U(50, Q) = \beta \epsilon [u(25) - u(0)] + \epsilon \alpha [u(75) - u(50) + u(100) - u(50)] > 0,$$

as claimed. We did not need to invoke the concavity of the utility function for this example because it satisfies (6.20), in which case it is sufficient that $u \in \mathbf{u}$.

In the case of first and second stochastic dominance (Definitions 6.8, 6.9) the dominant distribution, P , has a mean that is no lower than the dominated distribution, Q . This need not be the case under FS dominance. From the data given to us,

$$\mu_Q - \mu_P = 50\epsilon > 0.$$

Thus, the preferred distribution in this case, P , has a lower mean. The reason is that, for $\alpha > 0$, a reduction in the mean can be associated with a reduction in disadvantageous inequality for an individual with income y_j without affecting that individual's income. For instance, this could happen by destroying some of the wealth of individuals with income higher than y_j (but still leaving them with incomes no less than y_j). This is supported by the evidence where subjects in experiments prefer a smaller cake which is more equitably distributed.¹⁰

The conditions in (6.20), (6.21) are examples of the more general conditions that are developed in Dhami and al-Nowaihi (2013). They call these conditions, respectively, *first order FS dominance* and *second order FS dominance*. We state these conditions and some of the relevant results below that the reader is asked to prove in the exercises.

Definition 6.11 (*First order Fehr–Schmidt dominance*): Consider an individual with income $y_j \in \mathbf{Y}$. Let $P, Q \in \Pi$. Then P first order FS_{y_j} dominates Q ($P \succeq_{FS1_{y_j}} Q$) if

- (a) $P_i \leq Q_i$ for each $i = 1, 2, \dots, j-1$, and
- (b) $P_k \geq Q_k$ for each $k = j, j+1, \dots, n-1$.

If, in addition, one of these inequalities is strict, then we say that P strictly first order FS_{y_j} dominates Q ($P \succ_{FS1_{y_j}} Q$).

Thus, distribution P first order FS dominates Q for an individual of income level $y_j \in \mathbf{Y}$ if it first order stochastically dominates below y_j and is first order stochastically dominated above y_j . Second order FS dominance that we now define relies on a similar interpretation.

Definition 6.12 (*Second order Fehr–Schmidt dominance*): Let $P, Q \in \Pi$, $y_j \in \mathbf{Y}$. Then P second order FS_{y_j} dominates Q ($P \succeq_{FS2_{y_j}} Q$) if

- (a) Incomes are equally spaced (Definition 6.9),
- (b) $\tilde{P}_i \leq \tilde{Q}_i$ for each $i = 1, 2, \dots, j-1$, and
- (c) $\tilde{P}_k \geq \tilde{Q}_k$ for each $k = j, j+1, \dots, n-1$.

If, in addition, one of the inequalities in (b) or (c) is strict, then we say that P strictly second order FS_{y_j} dominates Q ($P \succ_{FS2_{y_j}} Q$).

Proposition 6.10 (Dhami and al-Nowaihi, 2013): Consider an individual with income $y_j \in \mathbf{Y}$. Let $P, Q \in \Pi$. Then P first order FS_{y_j} dominates Q ($P \succeq_{FS1_{y_j}} Q$) if, and only if, $U(y_j, P) \geq U(y_j, Q)$ for all $u \in \mathbf{u}$, $\alpha \geq 0$ and $\beta \in [0, 1)$.

Proposition 6.11 (Dhami and al-Nowaihi, 2013): Let $P, Q \in \Pi$, $y_j \in \mathbf{Y}$. Let P second order FS_{y_j} dominate Q ($P \succeq_{FS2_{y_j}} Q$). Then $U(y_j, P) \geq U(y_j, Q)$ for all $\alpha \geq 0$, $\beta \in [0, 1)$ and all concave $u \in \mathbf{u}$.

The next proposition establishes a partial converse of Proposition 6.11.

¹⁰ See, for instance, Ackert et al. (2007), Bolton and Ockenfels (2006), and Tyran and Sausgruber (2006).

Proposition 6.12 (Dhami and al-Nowaihi, 2013): Let $P, Q \in \Pi$ and $y_j \in Y$. Suppose incomes are equally spaced (Definition 6.9). Suppose $\tilde{P}_{j-1} = \tilde{Q}_{j-1}$. If $U(y_j, P) \geq U(y_j, Q)$ for all $\alpha \geq 0$, all $\beta \in [0, 1)$ and all concave $u \in \mathbf{u}$, then P second order FS_{y_j} dominates Q ($P \succeq_{FS_{y_j}} Q$).

In the model with self-regarding preferences, the relation between second order stochastic dominance and generalized Lorenz dominance is a crucial one.¹¹ Indeed, such a relation is often used in a justification of cross-country ranking of income distributions. Dhami and al-Nowaihi (2013) show that an analogous result holds for the case of FS preferences. It turns out that a slight modification of second order FS dominance is equivalent to a particular modification of generalized Lorenz dominance. This makes available the full toolkit of dominance concepts for social preferences of the FS form.

6.5 Behavioral political economy

The explanation of societal redistribution is one of the central problems in political economy. Existing explanations rely on voters who have self-regarding preferences.¹² In the literature, the simplest political institution is *direct democracy* in which voters vote directly to choose between alternative policies. This is not only an important benchmark institution, its role in actual policy choice has been expanding.¹³ For these reasons, we restrict ourselves to direct democracy in this section.

The recent empirical literature strongly supports the role of social preferences in voting models (Ackert et al., 2007; Bolton and Ockenfels, 2006; Tyran and Sausgruber, 2006). These papers establish that voters often choose policies that promote equity/fairness over purely self-regarding considerations. Bolton and Ockenfels (2006), for instance, examine the preference for equity versus efficiency in a voting game. Groups of three subjects are formed and are presented with two alternative policies. One policy promotes equity, while the other promotes efficiency. The final outcome is chosen by a majority vote. About twice as many subjects in experiments preferred equity as compared to efficiency.¹⁴

Experimental results on voting under direct democracy lend support to the FS model. Tyran and Sausgruber (2006) explicitly test for the importance of the FS framework in explaining the experimental results on voting. They conclude that the FS model predicts much better than the standard model with self-regarding preferences. In addition, for the three income classes in their experiments, the FS model provides, in their words, “strikingly accurate predictions for individual voting in all three income classes.”

The econometric results of Ackert et al. (2007), based on their experimental data, lend further support to the FS model in the context of redistributive taxation. The estimated coefficients of altruism and envy in the FS model are statistically significant and, as expected, negative in sign. In the context of voting experiments, Bolton and Ockenfels (2006) conclude that “while not everyone measures fairness the same way, the simple measures offered by [Bolton and Ockenfels (2000)] or FS provide a pretty good approximation to population behavior over a wide range of scenarios that economists care about.”

¹¹ See Lambert (2001) for a statement and proofs of the relevant results as well as a good discussion of the issues.

¹² For a wide-ranging review, see Persson and Tabellini (2000).

¹³ See, for instance, Matsusaka (2005).

¹⁴ Furthermore, even those willing to change the status quo for efficiency are willing to pay, on average, less than half relative to those who wish to alter the status quo for equity.

Economy-wide voting under direct democracy is impersonal and anonymous, thereby making it unlikely that intentions play any important role. Hence, FS preferences are a particularly good vehicle through which to model social preferences in a voting context. In the discussion below we use the term *self-regarding voter model* when all voters have self-regarding preferences and the term *other-regarding voter model* when at least some voters have social preferences.

Dhami and al-Nowaihi (2010a, 2010b) provide the relevant theoretical results that are motivated by the experimental results and open the way for further applications of the other-regarding voter model. In particular, the neoclassical results on the existence of Condorcet winner derived for self-regarding voters no longer apply for fair voters. Dhami and al-Nowaihi (2010a) derive conditions for the existence of a Condorcet winner when all voters have a social preference of the FS form. Dhami and al-Nowaihi (2010b) consider the implications for societal redistribution when there is a mixture of other-regarding and self-regarding voters. We consider these two contributions below.

6.5.1 Existence of a Condorcet winner when voters have social preferences

We consider a standard, general equilibrium model with endogenous labor supply that gives the simplest trade-off between equity and efficiency.¹⁵ Let there be $n = 2m - 1$ voters who combine the roles of producer-consumers and voters. Let the skill level of voter j be s_j , $j = 1, 2, \dots, n$, where $0 < s_i < s_j < 1$, for $i < j$. Denote the skill vector by $\mathbf{s} = (s_1, s_2, \dots, s_n)$ and the median skill level by s_m . Each voter has a fixed time endowment of one unit, so supplies l_j units of labor and enjoys $1 - l_j$ units of leisure, where $0 \leq l_j \leq 1$. Labor markets are competitive and each firm has access to a linear production technology. Thus, firm j , represented by voter j , produces an amount y_j such $y_j = s_j l_j$. Hence, the wage rate offered to each voter coincides with the marginal product, i.e., the skill level, s_j . Thus, the before-tax income of voter j is given by $y_j = s_j l_j$. Denote the average before-tax income by $\bar{y} = \frac{1}{n} \sum_{j=1}^n y_j$.

The government operates a linear progressive income tax system. This is characterized by a constant marginal tax rate, t , $t \in [0, 1]$, and a uniform transfer, b , to each voter that equals the average tax proceeds, thus, $b = t\bar{y}$. The budget constraint of voter j is given by $0 \leq c_j \leq (1 - t)y_j + b$, which can be written as

$$0 \leq c_j \leq (1 - t)s_j l_j + b. \quad (6.22)$$

We define a voter's preferences in two stages. First, let voter j have a continuous *own-utility* function, $\tilde{u}_j(c_j, 1 - l_j)$, defined over own-consumption, $c_j > 0$, and own-leisure, $1 - l_j$, $0 \leq l_j \leq 1$. Each voter has a twice continuously differentiable, and identical, utility function

$$\tilde{u}_j = \tilde{u}(c_j, 1 - l_j). \quad (6.23)$$

Differences in c_j and l_j across voters arise only due to differences in skill levels but not due to differences in the utility functions. We assume that own-consumption is always desirable, so $\tilde{u}_j$ is a strictly increasing function of c_j . Second, voters have *other-regarding preferences* as in

¹⁵ This is also known as the Roberts–Romer–Meltzer–Richards model (or RRMR model) based on the names of economists who formally articulated it in seminal contributions to political economy. For the relevant references, see, for instance, Persson and Tabellini (2000) and Dhami and al-Nowaihi (2010 a,b).

Fehr–Schmidt (1999). Let c_{-j} and l_{-j} be the vectors of consumption and labor supplies, respectively, of voters other than voter j . Then the FS-utility of voter j is¹⁶

$$\tilde{U}_j(c_j, l_j; c_{-j}, l_{-j}, t, b, \alpha, \beta, s) = \tilde{u}_j - \frac{\alpha}{n-1} \sum_{k \neq j} \max\{0, \tilde{u}_k - \tilde{u}_j\} - \frac{\beta}{n-1} \sum_{i \neq j} \max\{0, \tilde{u}_j - \tilde{u}_i\},$$

where, for fair voters, $0 < \beta < 1$ and $\alpha > 0$ with the typical experimental finding suggesting that $\alpha > \beta$. For *self-regarding* voters $\alpha = \beta = 0$, so $\tilde{U}_j = \tilde{u}_j$. Thus, $\tilde{u}$ is also the utility function of a self-regarding voter, as in the standard model. The other-regarding voter cares about own payoff (first term), payoff relative to those where inequality is disadvantageous (second term) and payoff relative to those where inequality is advantageous (third term).

Since $\tilde{u}_j$ is a strictly increasing function of c_j , and since $\alpha \geq 0$ and $0 \leq \beta < 1$, it follows that $\tilde{U}_j$ is also a strictly increasing function of c_j . Hence, the budget constraint (6.22) holds with equality. Substituting $c_j = (1-t)s_j l_j + b$, from (6.22), into the utility function of a self-regarding voter, $\tilde{u}_j(c_j, 1-l_j)$, we get

$$u_j = u(l_j; t, b, s_j) = \tilde{u}((1-t)s_j l_j + b, 1-l_j) \quad (6.24)$$

Correspondingly, the FS-utilities take the form

$$U_j(l_j; l_{-j}, t, b, \alpha, \beta, s) = u_j - \frac{\alpha}{n-1} \sum_{k \neq j} \max\{0, u_k - u_j\} - \frac{\beta}{n-1} \sum_{i \neq j} \max\{0, u_j - u_i\}. \quad (6.25)$$

We consider a two-stage game. In the first stage, all voters vote directly for their most preferred redistributive tax rate. If a median voter equilibrium exists, then the tax rate preferred by the median voter is implemented—the outcome in this case is known as a *Condorcet winner*. In the second stage, all voters make their labor supply decision, conditional on the tax rate chosen by the median voter in the first stage. On choosing their labor supplies in the second stage, the announced first period tax rate is implemented. In the second stage, the voters play a one-shot Nash game: each voter, j , chooses his/her labor supply, l_j , given the vector, l_{-j} , of labor supplies of the other voters, so as to maximize his/her FS-utility (6.25).

The solution is by backward induction. So consider first, the second stage. Given the tax rate, t , and the transfer, b , both determined in the first stage, the voters play a one-shot Nash game to determine their labor supplies (in the subgame determined by t and b).

Proposition 6.13 (Dhami and al-Nowaihi, 2010a): *In the second stage of the game, voter j , whether other-regarding or self-regarding, chooses own labor supply, l_j , so as to maximize own-utility, u_j defined in (6.24), given t , b and s_j .*

The proof of Proposition 6.13 is left as an exercise for the reader. From Proposition 6.13, an other-regarding voter, despite having social preferences, chooses labor supply exactly like a self-regarding voter. However, when making a decision on the redistributive tax rate in the first stage (see below), the two kinds of voters make different choices. In other words, in two separate domains, labor supply and redistributive choice, the other-regarding voter behaves *as if* he had self-regarding preferences in the first domain and social preferences in the second. In this respect, these results are similar to the contrast between the results in the ultimatum game with

¹⁶ Here, we consider the non-linear version of FS preferences; see Fehr and Schmidt (1999) and Definition 6.7.

and without proposer (or responder) competition in Section 6.2.1.¹⁷ Note that Proposition 6.13 would not hold if in the FS utility function of voter j in (6.25), the u_i , $i \neq j$, entered non-additively. However, the empirical evidence supports the adopted form for the FS-utility function in the voting domain.

For the generic utility function $\tilde{u}(c, 1-l)$, defined in (6.23), we use the standard notation $\tilde{u}_1 = \frac{\partial \tilde{u}(c, 1-l)}{\partial c}$ and $\tilde{u}_2 = \frac{\partial \tilde{u}(c, 1-l)}{\partial (1-l)}$. The utility function has the following, plausible, properties:

$$(a) \tilde{u}_1 > 0, (b) l > 0 \Rightarrow \tilde{u}_2(c, 1-l) > 0, (c) \tilde{u}_2(c, 1) = 0, (d) \tilde{u}_1(c, 0) \leq \tilde{u}_2(c, 0), \quad (6.26)$$

$$(a) \tilde{u}_{11} \leq 0, (b) \tilde{u}_{12} \geq 0, (c) l > 0 \Rightarrow \tilde{u}_{22}(c, 1-l) < 0, (d) (\tilde{u}_{12})^2 \leq \tilde{u}_{11}\tilde{u}_{22}. \quad (6.27)$$

From (6.26a), the marginal utility of consumption is positive, while (6.26b) implies that marginal utility of leisure is positive, unless $l = 0$ in which case (6.26c) says that the consumer is satiated with leisure. From (6.26d), when a consumer has no leisure, she always (weakly) prefers one extra unit of leisure to one extra unit of consumption. (6.27a) says that marginal utility of consumption is non-increasing. From (6.27b), consumption and leisure are complements, while (6.27c) implies that the marginal utility of leisure is strictly declining unless, possibly, the consumer is satiated with leisure (in which case $\tilde{u}_{22}(c, 1) = 0$). Conditions (6.26) and (6.27) guarantee that a maximum exists, that it is unique and that it is an interior point ($0 < l_i < 1$), unless $t = 1$ in which case the maximum will lie at $l_i = 0$. Conditions (6.27a,c,d) guarantee that $\tilde{u}$ is concave.

We list, in Lemmas 6.6 and 6.7 below, some useful results; the proofs are left as an exercise for the reader. Lemma 6.6 shows that there is a unique solution to the labor supply of an individual and characterizes other useful properties of labor supply. For instance, the optimal labor supply is decreasing in the redistributive tax rate (conflict between efficiency and equity).

Lemma 6.6 (*Properties of labor supply; Dhami and al-Nowaihi, 2010a*):

- (a) Given t, b , and s_j , there is a unique labor supply for voter j , $l_j = l(t, b, s_j)$, that maximizes own-utility $u(l_j; t, b, s_j)$ defined in (6.24),
- (b) $t \in [0, 1] \Rightarrow 0 < l_j < 1$,
- (c) $l_j = 0$ at $t = 1$,
- (d) $t \in [0, 1] \Rightarrow \left[\frac{\partial u}{\partial l_j}(l_j; t, b, s_j) \right]_{l_j=l(t, b, s_j)} = 0$,
- (e) $l(t, b, s_j)$ is continuously differentiable,
- (f) $\frac{\partial l(t, b, s)}{\partial b} \leq 0$,
- (g) for each $t \in [0, 1]$, the equation $b = \frac{1}{n} t \sum_{i=1}^n s_i l(t, b, s_i)$ has a unique solution $b(t, s) \geq 0$; and $b(t, s)$ is continuously differentiable.

Substituting optimal labor supply, given by Lemma 6.6(a), in $u(l_j; t, b, s_j)$ we get the indirect own-utility function of voter j ,

$$v_j = v(t, b, s_j) = u(l(t, b, s_j); t, b, s_j). \quad (6.28)$$

Lemma 6.7 lists some useful properties of the indirect utility function, v . In particular, v is increasing in the lump-sum transfer, b . If the tax rate is bounded away from 1, then v is increasing in the skill level of the individual, s , and decreasing in the redistributive tax rate, t (holding fixed, b).

¹⁷ In other contexts, individuals might, for instance, send their own children to private schools (self-interest) but could at the same time vote for more funding to government run schools in local or national elections (public interest). Thus, individuals can put on different hats in different situations.

In Lemma 6.7, the distinction between the results at $t = 1$ and $t \in [0, 1)$ follow from Lemma 6.6(b), (c)

Lemma 6.7 (*Properties of the indirect own utility function; Dhimi and al-Nowaihi, 2010a*):

$$\begin{aligned} (a) \quad & \frac{\partial v}{\partial b} > 0, \\ (bi) \quad & \left[\frac{\partial v}{\partial s} \right]_{t=1} = 0, & (bii) \quad t \in [0, 1) \Rightarrow \frac{\partial v}{\partial s} > 0, \\ (ci) \quad & \left[\frac{\partial v}{\partial t} \right]_{t=1} = 0, & (cii) \quad t \in [0, 1) \Rightarrow \frac{\partial v}{\partial t} < 0, \\ (di) \quad & \frac{\partial^2 v}{\partial s \partial b} \leq 0, & (dii) \quad \left[\frac{\partial^2 v}{\partial s \partial b} \right]_{t=1} = 0. \end{aligned}$$

Substituting labor supply, $l(t, b, s_j)$, into $y_j = s_j l_j$ gives before-tax income $y_j(t, b, s_j) = s_j l(t, b, s_j)$.

We now solve the first stage problem of the choice of the optimal tax rate. Substituting labor supply, $l(t, b, s_j)$, into the utility function, $U_j(l_j; 1-j, t, b, \alpha, \beta, s)$, and using (6.25) and (6.28), the indirect utility function, $V_j(t, b, \alpha, \beta, s)$, of voter j is given by

$$V_j(t, b, \alpha, \beta, s) = v_j - \frac{\alpha}{n-1} \sum_{k \neq j} \max\{0, v_k - v_j\} - \frac{\beta}{n-1} \sum_{i \neq j} \max\{0, v_j - v_i\}, \quad (6.29)$$

where v_j is defined in (6.28). In the light of Lemmas 6.7(bi) and 6.7(bii), (6.29) becomes

$$V_j(t, b, \alpha, \beta, s) = v_j - \frac{\alpha}{n-1} \sum_{k > j} [v_k - v_j] - \frac{\beta}{n-1} \sum_{i < j} [v_j - v_i]. \quad (6.30)$$

We now introduce two further assumptions, A1 and A2, followed by the main result.

A1: $t \in [0, 1) \Rightarrow \frac{\partial^2 v}{\partial s \partial t}(t, b, s) < 0$, where v is defined in (6.28).

Recall, from Lemma 6.7(cii), that $t \in [0, 1) \Rightarrow \frac{\partial v}{\partial t} < 0$. Hence, $\frac{\partial}{\partial s} \left(t \frac{\partial v}{\partial t} \right) = t \frac{\partial^2 v}{\partial s \partial t} < 0$ can be interpreted as saying that an extra 1% on the redistributive tax rate hurts a poor person less than a rich person. Thus Assumption A1 roughly says that redistributive taxes hurt the poor less than the rich. This is the basic foundation of the modern welfare state.¹⁸

A2: $\frac{\partial V_j}{\partial t}(t, b, \alpha, \beta, s) \leq 0$.

Since $\frac{\partial v}{\partial t} < 0$, for $t < 1$ (Lemma 6.7(cii)), an increase in tax (benefit, b , remaining fixed), is undesirable for a self-regarding-voter, which is entirely reasonable. Assumption A2 extends this to fair-voters as well. It implies that envy is not so great as to make a fair-voter prefer an increase in tax, even if it has no gain for anyone in terms of an increase in the benefits, b (b is held fixed in computing $\frac{\partial V_j}{\partial t}$ in A2).

Thus, conditions A1 and A2 bear an intuitive interpretation. Dhimi and al-Nowaihi (2010a) show that the class of utility functions that satisfy assumptions A1 and A2 is very large. They also discuss the interpretation of assumptions A1, A2 in terms of a voter's own-utility $\tilde{u}_j(c_j, 1 - l_j)$ but there are no obvious insights to be gained from this.

Proposition 6.14 (*Dhimi and al-Nowaihi, 2010a*): *Under assumptions A1 and A2 a majority of the other-regarding voters prefer the tax rate that is optimal for the median-skill voter.*

¹⁸ Assumption A1 is satisfied, for instance, in the important case of quasi-linear preferences, which is widely used in this literature.

Proposition 6.14 gives the conditions for the existence of a Condorcet winner when voters have social preferences. The modern literature on political economy when voters are self-regarding really took off after the publication of the conditions for the existence of a Condorcet winner. Proposition 6.14 gives the corresponding results for the case of other-regarding voters.

So far, we have assumed that all voters have social preferences. However, the evidence typically shows up a mixture of self-regarding and other-regarding individuals. The next section considers the implications of such a mixture for optimal redistribution. In order to prepare the ground for the next section we give some intermediate results.

Substituting labor supply, $l(t, b, s_j)$, into $y_j = s_j l_j$ gives before-tax income $y_j(t, b, s_j) = s_j l(t, b, s_j)$. Substituting $b(t, s)$, given by Lemma 6.6(g), into the indirect own-utility function (6.28), gives

$$w_j(t, s) = v(t, b(t, s), s_j). \quad (6.31)$$

Let

$$W_j(t, \alpha, \beta, s) = V_j(t, b(t, s), \alpha, \beta, s), \quad (6.32)$$

where $b(t, s)$ is given by Lemma 6.6(g). Then (10.18) can be written as

$$W_j(t, \alpha, \beta, s) = w_j(t, s) - \frac{\alpha}{n-1} \sum_{k>j} [w_k(t, s) - w_j(t, s)] - \frac{\beta}{n-1} \sum_{i<j} [w_j(t, s) - w_i(t, s)], \quad (6.33)$$

where w_j is defined in (10.22). This gives the Fehr–Schmidt utility of an individual with skill level s_j that depends only on the redistributive tax rate, t .

6.5.2 Heterogeneity in voter preferences and redistribution

We now examine some of the insights from Dhami and al-Nowaihi (2010b) who consider the implications for societal redistribution arising from a mixture of other-regarding and self-regarding voters. Consider the model in Section 6.5.1 but assume, for pedagogical ease, that voters belong to one of three categories: poor, middle class, and rich, whose skills are ordered as $s_1 < s_2 < s_3$. We also assume that voters have identical α, β parameters.¹⁹ Each type of voter can have either self-regarding preferences ($\alpha = \beta = 0$) or other-regarding preferences ($\alpha, \beta \neq 0$). One can assume either that there is a large number of voters who are uniformly distributed over the three skill levels or just three voters, one of each type. We use the latter interpretation below.

We assume that the own-utility function is quasi linear, with constant elasticity of labor supply, which is the most commonly used functional form in various applications of the median voter theorems.

$$\tilde{u}(c, 1-l) = c - \frac{\epsilon}{1+\epsilon} l^{\frac{1+\epsilon}{\epsilon}}, \quad (6.34)$$

¹⁹ For the case where the fairness parameters of voters can differ, so that the fairness parameters of voter $j = 1, 2, 3$ are summarized α_j, β_j , see Dhami and al-Nowaihi (2010b).

where ϵ is the constant elasticity of labor supply, assumed positive.²⁰ One can check that the more general case considered in Lemmas 6.6 and 6.7 continues to hold. Using (6.33), (6.34), and Lemmas 6.6 and 6.7, we can prove the next proposition.

Proposition 6.15 (Dhami and al-Nowaihi, 2010b) (Existence of optimal tax rates): Consider $W_j(t, \alpha, \beta, s)$ defined in (6.33).

- (a) Given α, β and s , $W_j(t, \alpha, \beta, s)$ attains a maximum at some $t_j \in [0, 1)$.
- (b) Whenever there is an interior solution to the tax rate, the tax rate, t_j , preferred by voter j , is unique.
- (c) $W_j(t, \alpha, \beta, s)$ is strictly increasing on $[0, t_j)$ and strictly decreasing on $(t_j, 1]$.
- (d) t_j is non-decreasing in α, β .

The final result in Proposition 6.15(d) deserves to be commented upon. The intuition is that an increase in α_j increases disutility arising from disadvantageous inequity. By increasing the redistributive tax rate, the voter reduces, relatively, the utility of anyone who is richer, hence, reducing disadvantageous inequity. On the other hand, an increase in β_j increases disutility arising from advantageous inequity. An increase in the redistributive tax benefits everyone poorer than the voter relatively more, thus, reducing advantageous inequity.

For the three-voter model, poor, middle class, and rich, there are $2^3 = 8$ possible combinations of the preferences of voters. Denoting by S and F , respectively, a self-regarding and a fair, or other-regarding voter, the eight possible combinations of the voters, each combination arranged in order of increasing skill level from left to right, are: $SSS, SSF, SFS, SFF, FSS, FSF, FFS, FFF$. For instance, SFF denotes an economy in which the lowest skill voter is self-regarding and the middle and high skill voters are both fair. Since the experimental evidence is that 1/3 to 2/3 of the participants are self-regarding, the most important cases are $SSF, SFS, SFF, FSS, FSF, FFS$.

Proposition 6.15 tells us about the features of the most preferred tax rate by any voter. Does a Condorcet winner exist in this economy? We cannot directly apply the results of Proposition 6.14 because we now have a mixture of self-regarding and other-regarding voters. The main result is summarized in Proposition 6.16.

Proposition 6.16 (Dhami and al-Nowaihi, 2010b): Let

$$\theta_1 = \frac{s_2^{1+\epsilon} - s_1^{1+\epsilon}}{s_3^{1+\epsilon} - s_2^{1+\epsilon}}; \theta_2 = \frac{2}{2 + \theta_1},$$

then,²¹ for the three-voter model:

- (a) The median-skill voter with skill s_2 is the Condorcet winner. Let $t_j, j = 1, 2, 3$ be the most preferred tax rate of voter j . Then the most preferred tax rates are ordered by magnitude as $t_3 < t_2 < t_1$, in the following cases:
 - (i) SSS, FSS, FFS , and FFF ,
 - (ii) SFS and SFF , if $\frac{\alpha}{2-\beta} < \theta_1$,
 - (iii) SSF and FSF , if $\beta < \theta_2$,

²⁰ The case $\epsilon = 1$ has special significance in the literature. In this case, $\tilde{u}(c, 1-l) = c - \frac{1}{2}l^2$.

²¹ It can be proved that the constant θ_1 is, for the three-voter case, directly proportional to the ratio of advantageous to disadvantageous inequality for voter 2; see Dhami and al-Nowaihi (2010b). θ_2 depends inversely on θ_1 .

- (b) *The low-skill voter with skill s_1 is the Condorcet winner, and $t_3 < t_1 < t_2$, in cases SFS and SFF, if $\theta_1 < \frac{\alpha}{2-\beta}$,*
- (c) *The high-skill voter with skill s_3 is the Condorcet winner, and $t_2 < t_3 < t_1$, in cases SSF and FSF, if $\theta_2 < \beta$.*

We now offer a discussion of Proposition 6.16.

In cases SSS, FSS, FFS, and FFF, the median-skill voter is decisive in the redistributive tax choice (Proposition 6.16a(i)). Note that in the cases SSS, FSS, the median-skill voter is self-regarding, while in cases FFS and FFF the median-skill voter is other-regarding. Hence, in these four cases, the redistributive outcome in the case of a mixture of voter types is identical to a hypothetical economy in which all voters are of the same type as the median-skill voter. So, for instance, the redistributive outcome in the FSS economy is the same as that in a SSS economy, while the outcome in the FFS economy is identical to the FFF economy.

The same result is also true for the other four cases SFS, SFF, SSF, and FSF provided that the relevant upper bounds on α and/or β are satisfied (Proposition 6.16a(ii) and a(iii)). In each of these eight cases, the tax rate implemented is the tax rate, t_2 , preferred by the median-skill voter.

If the upper bounds in Proposition 6.16a(ii), a(iii) are not satisfied, then the median-skill voter is no more decisive. From Proposition 6.16b, in cases SFS and SFF, the low-skill (and self-regarding) voter becomes decisive, and the implemented tax rate is t_1 . From Proposition 6.16c, in cases SSF and FSF, the high-skill (and other-regarding) voter becomes decisive, and the implemented tax rate is t_3 .

Note that in all such cases, *the implemented tax rate is the median tax rate and not, necessarily, the tax rate preferred by the median-skill voter.*

Consider the four striking cases in the following corollary to Proposition 6.16.

Corollary 6.1

- (a) *For SFS, the majority of voters are self-regarding yet, if $\frac{\alpha}{2-\beta} < \theta_1$, then a majority votes for the tax rate preferred by the median-skill other-regarding voter (Proposition 6.16a(ii)).*
- (b) *For SFF, the majority of voters are fair yet, if $\theta_1 < \frac{\alpha}{2-\beta}$, then a majority votes for the tax rate preferred by the low-skill self-regarding voter (Proposition 6.16b).*
- (c) *For SSF, the majority of voters are self-regarding yet, if $\theta_2 < \beta$, then a majority votes for the tax rate preferred by the high-skill other-regarding voter (Proposition 6.16c).*
- (d) *For FSF, the majority of voters are fair yet, if $\theta_2 < \beta$, then a majority votes for the tax rate preferred by the median-skill self-regarding voter (Proposition 6.16a(iii)).*

Denote by t_j^S and t_j^F to be the most preferred tax rates of voter j if he is, respectively, self-regarding and fair or other-regarding. From Proposition 6.16a(i), in an SSS economy, the median-skill voter is decisive, so $t_3^S < t_2^S < t_1^S$ because richer self-regarding voters prefer a lower redistributive tax. This is shown in Figure 6.1 by the curve labeled *abc*. Now suppose that we replace the *self-regarding* median-skill voter in an SSS economy by an other-regarding voter to generate the SFS economy. There are two cases:

- (a) If the inequality $\theta_1 < \frac{\alpha}{2-\beta}$ holds, then from Proposition 6.16b, for the SFS economy, we get the ranking of most preferred tax rates to be $t_3^S < t_1^S < t_2^F$. This is shown in Figure 6.1 by the curve *adc*. The tax rate of the low-skill self-regarding voter, t_1^S , is now able to defeat any of the other two tax rates in a pairwise contest (either of the other two classes of voters find it optimal to align with the low-skill voter against the tax proposal of the third voter). The optimal tax rate




Figure 6.1 The SSS and SFS economies.

of the lowest-skill voter becomes the decisive redistributive choice. Hence, in moving from the SSS to the SFS economy, there is a jump in the tax rate from point b to point a (or from t_2^S to t_1^S).

(b) If the restriction $\frac{\alpha}{2-\beta} < \theta_1$ holds, then we get the case in Proposition 6.16a(ii). In this case, the ranking of tax rates, for the SFS economy, is $t_3^S < t_2^F < t_1^S$, so the median skill voter is also the decisive median voter. In Figure 6.1, this ranking of tax rates is shown by the curve aec . In this case, as we transit from the economy SSS to the economy SFS, there is a smaller upward jump from point b to point e .

In Proposition 6.16b, the crucial inequality is $\theta_1 < \frac{\alpha}{2-\beta}$; for the SFS economy, the relevant curve is adc in Figure 6.1. To see why this inequality is important, suppose that there is pairwise voting between the two tax rates, t_1^S, t_2^F . For the high-skill voter (whose skill level is s_3) to prefer the tax rate of the low-skill voter, t_1^S , the optimal tax rate of the middle-skill voter, t_2^F , should be “too high” in the sense that $t_1^S < t_2^F$. In this case, a majority (the rich and the poor) will prefer the tax rate t_1^S to t_2^F . The following factors help to satisfy the inequality $\theta_1 < \frac{\alpha}{2-\beta}$, and, hence, are conducive to t_2^F being “too high.”

(a) High inequity aversion, as captured by the magnitudes of α, β : From Proposition 6.15d, higher magnitudes of the inequity aversion parameters, α, β , increase the optimal tax rate for other-regarding voters, so increase t_2^F .

(b) High skill inequality at the upper end of the skill distribution ($s_3 - s_2$) and low inequality at the lower end of the income distribution ($s_2 - s_1$) reduce θ_1 and increase t_2^F : To see this, suppose that $s_3 - s_2$ increases because of an increase in s_3 . This leads to an increase in disadvantageous inequity faced by a voter of skill level s_2 . Self-regarding voters would like to redistribute more when the rich get richer because average incomes increase, so the lump-sum available for redistribution is higher. Other-regarding voters have an additional motive to redistribute more, namely, that it reduces disadvantageous inequity. This facilitates an increase in t_2^F .

For self-regarding voters, a decrease in the low-skill level reduces the redistributive tax rate. The intuition is that reduction in low-skill income reduces average income available for redistribution, hence, reducing the marginal benefits of increasing the tax rate. For other-regarding voters, however, the results can go either way. The reason is that on the one hand, the other-regarding voter is influenced by very similar considerations to the self-regarding voter (because the other-regarding voter also cares about “own” payoff). However, on the other hand, the empathy/concern for poorer voters, on account of the disutility arising from advantageous inequity induces the other-regarding voter in the opposite direction, i.e., greater redistribution. The interplay between these two opposing factors determines if the other-regarding voter will respond,




Figure 6.2 The SSS, SSF, and FSF economies.

unlike the self-regarding voter, by redistributing more in response to poverty. The restriction $\theta_1 < \frac{\alpha}{2-\beta}$ ensures that overall, one gets the ranking $t_3^S < t_1^S < t_2^F$ for the optimal tax rates of the three voting classes.

In Figure 6.2, we illustrate the second generic case when $\beta > \theta_2$ (see Proposition 6.16c). For the benchmark SSS economy, we know from Proposition 6.16a(i), that the median-skill voter, of skill s_2 , is decisive, so $t_3^S < t_2^S < t_1^S$. Thus, the SSS economy is represented by the line abc . Now replace the highest-skill self-regarding voter by an other-regarding voter to generate the economy SSF. Under the restriction specified in Proposition 6.16c ($\beta > \theta_2$), the economy SSF is shown in Figure 6.2 by the curve abe and one gets the ordering of tax rates to be $t_2^S < t_3^F < t_1^S$. In this case, the rich but other-regarding voter is decisive because the proposed tax rate, t_3^F , can defeat any of the other two proposed tax rates in a pairwise contest. Under the restriction $\beta > \theta_2$, the FSF economy is also shown in Figure 6.2, by the curve dbc . The decisive median voter is again the rich other-regarding voter, who is not the median-skill voter.

Why does the restriction $\beta > \theta_2$ not involve α ? The reason is that for the cases in Figure 6.2 to arise, we need the most preferred tax rate of the rich other-regarding voter to be high enough, in the sense that $t_2^S < t_3^F$. The rich other-regarding voter faces no disadvantageous inequality, which explains the absence of the parameter α . However, if β is high enough, then the rich other-regarding voter suffers sufficiently highly from advantageous inequality with respect to the other two class of voters. In particular, if $\beta > \theta_2$, then $t_2^S < t_3^F$, which leads to the surprising outcome in Proposition 6.16c.

6.6 Fairness, general equilibrium, and welfare

In the presence of other-regarding preferences, can we invoke the standard machinery of general equilibrium and welfare economics? For instance, we know from undergraduate microeconomics that in the presence of consumption externalities, the fundamental welfare theorems need not hold. Such consumption externalities are also reflected in social preferences; see, for instance, the Fehr–Schmidt preferences in (6.1). The relevant issues are addressed by Dufwenberg et al. (2011). In this section we give readers a flavor of their results.²²

Consider a production economy with goods (inputs and outputs) indexed by $l \in \{1, 2, \dots, L\}$, profit maximizing firms indexed by $j \in \{1, 2, \dots, J\}$, and utility maximizing consumers indexed

²² We omit the proofs. However, the proofs use standard arguments used in graduate texts in microeconomics.

by $i \in \{1, 2, \dots, I\}$. The price of good l is $p_l \geq 0$, and the price vector of all goods is given by p . The production set of firm j , $Y_j \subseteq \mathbb{R}^L$, is closed and bounded; negative and positive entries in the production set are, respectively, inputs and outputs. The profit function of firm j is given by $\pi_j(p)$; its existence for all p is guaranteed by the assumptions. The set of all possible production profiles is given by $Y = \prod_{j=1}^I Y_j$, and the typical element is denoted by y . Consumer i has an initial endowment of $e_i \in \mathbb{R}_+^L$, and consumption vector given by $x_i \in \mathbb{R}_+^L$. The consumption vector of all I consumers is $x \in \mathbb{R}_+^{LI}$. All firms are owned by consumers. The share of consumer i in firm j is θ_{ij} , hence, the income of the consumer is $w_i = pe_i + \sum_{j=1}^I \theta_{ij} \pi_j(p)$.

Denote by $B_i = \{x_i \in \mathbb{R}_+^L : px_i \leq w_i\} \subseteq \mathbb{R}_+^L$, the non-empty and compact budget set of consumer i . Let $B = (B_1, B_2, \dots, B_I)$ be a profile of budget sets of all consumers and $\mathcal{B}$ be the set of all budget profiles. Other-regarding preferences are captured by defining, for any consumer i , a binary preference relation, $\succsim_i$, over objects in the set $\mathbb{R}_+^L \times \mathcal{B}$; each object (x, B) is a pair of the consumption profile of all consumers and of the budget sets of all consumers. It will sometimes be convenient to write (x, B) as (x_i, x_{-i}, B) where x_{-i} is the consumption profile of all consumers other than consumer i .

Definition 6.13 *The preferences of the individuals satisfy the following properties:*

- (a) *Strictly monotone in own consumption: Fix x_{-i}, B and let $x_i \geq x'_i$ but $x_i \neq x'_i$. Then strict monotonicity requires $(x_i, x_{-i}, B) \succ_i (x'_i, x_{-i}, B)$, where $\succ_i$ is the strict preference relation.*
- (b) *Continuous: Fix any x', B . The set $P = \{(x, B) : (x, B) \succsim_i (x', B)\}$ is closed.*
- (c) *Strictly convex over own consumption: Fix x_{-i}, B . Let $x_i \neq x'_i$ and $(x_i, x_{-i}, B) \succsim_i (x'_i, x_{-i}, B)$. Then strict convexity requires that for $\alpha \in (0, 1)$, $(\alpha x_i + (1 - \alpha)x'_i, x_{-i}, B) \succ_i (x'_i, x_{-i}, B)$.*

Under these assumptions, suppose that each preference relation, $\succsim_i$, can be represented by a continuous utility function $U(x, B) : \mathbb{R}_+^{LI} \times \mathcal{B} \rightarrow \mathbb{R}$. Clearly, this utility function is more general than the standard one because utility also depends on the budget sets. For instance, an individual may be envious of another who has the ability to buy a larger basket of goods.

Definition 6.14 *(Economy): An economy $\mathcal{E}$ is described by the vector $(I, e, \{U_i\}, J, Y, \theta)$, i.e., the set of players, their endowments, their utility functions, the set of firms, their production sets, and the shares of consumers in each firm.*

We have assumed that the budget set is compact, the utility function of each consumer is continuous, and preferences are strictly convex, so the demand function exists and for consumer i it is given by $d_i(x_{-i}, B) \in \arg \max_{x_i \in B_i} U_i(x_i, x_{-i}, B)$. Thus, the optimal demand of any consumer, d_i , depends not only on the consumption levels but also the budget sets of others. This contrasts with the neoclassical case of self-regarding preferences in which these effects of the consumption and budget sets of others are absent. Can one impose conditions on $U_i(x_i, x_{-i}, B)$ such that these effects are also absent when we allow for other-regarding preferences? If so (see below), consumer i is said to behave as-if-classical.

Definition 6.15 *(As-if-classical consumers): Consumer i behaves as-if-classical if d_i does not depend on x_{-i}, B_{-i} , where B_{-i} gives the profile of budget sets of all consumers other than i .²³*

²³ The actions of others do determine prices, p , hence, also w_i . These factors also determine d_i , hence, even as-if-classical consumers are "indirectly" influenced by the actions of others.

The actions of others do determine prices, p , hence, also w_i . These factors also determine d_i , hence, even as-if-classical consumers are “indirectly” influenced by the actions of others.

The demand functions of consumer i , $d_i(x_{-i}, B)$, for the L goods can be written as $d_i(x_{-i}, B_i, B_{-i})$. Substituting the optimal demands in the definition of the budget sets we can write $B_i = B_i(p, w_i)$. Hence, we may write the demand functions of consumer i as $d_i(p, w_i, x_{-i}, B_{-i})$.

The next definition is helpful to determine the conditions on preferences that ensure consumer i behaves as-if-classical.

Definition 6.16 (*Separable preferences*): Given any two consumption profiles, x, x' , and two budgets sets, B, B' , the preferences of consumer i are said to be separable if

$$(x_i, x_{-i}, B) \succeq_i (x'_i, x_{-i}, B) \Leftrightarrow (x_i, x'_{-i}, B') \succeq_i (x'_i, x'_{-i}, B').$$

In Definition 6.16, preferences over own consumption bundles are independent of the consumption bundles and budget sets of others.

Proposition 6.17 (*Dufwenberg et al., 2011*): Separable preferences have a utility representation given by $V_i(m(x_i), x_{-i}, B)$, where (i) V_i is increasing in m , and (ii) the function $m: \mathbb{R}_+^L \rightarrow \mathbb{R}$, known as the consumer’s “internal utility function,” is continuous, strictly monotone, and strictly quasi-concave.

The importance of separable preferences is shown in the next proposition.

Proposition 6.18 (*Dufwenberg et al., 2011*):

- (i) If the preferences of consumers can be written as $V_i(m(x_i), x_{-i}, B)$ (Proposition 6.17), then the consumer behaves as-if-classical (Definition 6.15).
- (ii) If $d_i(p, w_i, x_{-i}, B_{-i})$ is continuously differentiable in (p, w_i) , then consumers behave as-if-classical if the utility function can be represented as $V_i(m(x_i), x_{-i}, B)$ (Proposition 6.17).

The conditions in Proposition 6.18 hold for Fehr–Schmidt preferences, as we now show in the next example.

Example 6.5 (*Weighted utilitarian preferences; Dhami and al-Nowaihi, 2010a,b*) Consider the Fehr–Schmidt preferences given in 6.1. Define the sets A_i and D_i as the set of consumers with respect to whom consumer i has respectively, advantageous and disadvantageous inequity. So

$$A_i = \{j : i \neq j \text{ and } x_j \leq x_i\}, \quad (6.35)$$

$$D_i = \{k : k \neq i \text{ and } x_k > x_i\}. \quad (6.36)$$

Denote the respective cardinalities of these sets by $|A_i|$ and $|D_i|$. Then FS-utility can be written in a way that is reminiscent of the weighted utilitarian form:

$$U_i = \omega_{ii}x_i + \sum_{j \neq i}^I \omega_{ij}x_j, \quad (6.37)$$

where,

$$\begin{aligned} j \in A_i &\Rightarrow \omega_{ij} = \frac{\beta}{n-1} > 0, \\ i = j &\Rightarrow \omega_{ii} = 1 - \frac{|A_i|\beta}{n-1} + \frac{|D_i|\alpha}{n-1} > 0, \\ k \in D_i &\Rightarrow \omega_{ik} = -\frac{\alpha}{n-1} < 0. \end{aligned} \quad (6.38)$$

Furthermore, the weights sum up to one, i.e., $\sum_{j=1}^I \omega_{ij} = 1$. For self-regarding consumers, $\omega_{ii} = 1$ and $\omega_{ij} = 0$ ($i \neq j$). Clearly (6.37) is of the same form as V_i (Proposition 6.17), hence, the consumer must behave as-if-classical.

We are now interested in exploring the link between the Walrasian equilibrium of an economy with other-regarding preferences, and the Walrasian equilibrium with self-regarding consumers.

Definition 6.17 A Walrasian equilibrium is given by a triple of prices, consumption profiles, and production plans of firms, i.e., (p^*, x^*, y^*) , such that: (i) Each x_i^* maximizes the utility of the consumer U_i , subject to the budget constraint B_i , (ii) Each y_j^* maximizes the profit function of the firm, π_j , subject to the production set Y_j . (iii) All markets for the L goods clear.

Definition 6.18 Consider any economy $\mathcal{E} = (I, e, \{U_i\}, J, Y, \theta)$ (Definition 6.14). The corresponding ‘internal economy’ is defined by $\mathcal{E}^{Int} = (I, e, \{m_i\}, J, Y, \theta)$, where m_i is defined in Proposition 6.17.

In Definition 6.18, the only difference between an economy and its internal economy is that in the latter, all consumers are self-regarding. The next proposition gives the link between the Walrasian equilibria of the two economies.

Proposition 6.19 (Dufwenberg et al., 2011): Suppose that the preferences of all consumers are separable (Definition 6.16) and strictly monotone in own-consumption. Then the set of Walrasian equilibria of the two economies, $\mathcal{E}$ and $\mathcal{E}^{Int}$ are identical.

Example 6.6 (Dhami and al-Nowaihi, 2010a,b): Section 6.5, above, gives a behavioral political economy application in a simple general equilibrium model of endogenous labor supply. Lemma 6.6 shows that there is a unique labor supply for voter j , $l_j = l(t, b, s_j)$, that maximizes own-utility $u(l_j; t, b, s_j)$ despite the fact that voter j has other-regarding preferences. In other words, the voter behaves as-if-classical. The same individual with Fehr–Schmidt preferences behaves as if he had self-regarding preferences when he chooses labor supply but other-regarding preferences when he chooses the redistributive tax rate. This result was puzzling to the authors and they used it to argue that preferences may be context dependent.²⁴ The second stage general equilibrium outcome in this model is identical for voters who have self-regarding or other-regarding preferences. This result was given before the more general insights of Dufwenberg et al. (2011) were known, albeit in the specific context of voting.²⁵

We now consider issues of welfare economics. Here, the analysis requires a more restricted version of $V_i(m(x_i), x_{-i}, B)$.

Definition 6.19 (Well-being externalities): Well-being externalities arise if the utility function of any consumer i is given by $W_i(m_i, m_{-i})$, where m_i is given in Proposition 6.17. W_i is increasing in each argument and $m_k(x_k)$ is increasing for each k .

Example 6.7 Consider two examples of well-being externalities. The first one is due to Edgeworth (see citation in Dufwenberg et al. 2011).

²⁴ The authors argued that: “For example, individuals, when making a private consumption decision might act so as to maximize their selfish interest. But in a separate role as part of the government, as a school governor or as a voter, could act so as to maximize some notion of public well being. Individuals might, for instance, send their own children to private schools (self interest) but could at the same time vote for more funding to government run schools in local or national elections (public interest).”

²⁵ The original results were presented in the Public Economics Weekend, held at Leicester in 2006.

$$W_i(m_i, m_{-i}) = m_i + \frac{\eta_i}{I-1} \left(\sum_{k \neq i} m_k \right),$$

where $\eta_i > 0$ captures altruism, while $\eta_i < 0$ captures spite. The second example is a version of the Bolton–Ockenfels ERC model.

$$W_i(m_i, m_{-i}) = m_i - \eta_i \left| m_i - \frac{\sum_k m_k}{I} \right|; 0 \leq \eta_i < 1.$$

Definition 6.20 (Feasibility and Pareto optimality): An allocation $x = (x_1, x_2, \dots, x_I)$ is feasible if there exists some production plan $y = (y_1, y_2, \dots, y_J) \in Y$ such that $\sum_{i=1}^I x_i \leq \sum_{i=1}^I e_i + \sum_{j=1}^J y_j$. A feasible allocation, x , is Pareto optimal if there does not exist another feasible allocation x' such that at least one consumer is better off and no consumer is worse off.

We first show by an example that the second fundamental welfare theorem need not hold.

Example 6.8 The object of this example is to show that Pareto efficient allocations need not be Walrasian equilibria in the presence of other-regarding preferences. Suppose that there are two consumers ($I = 2$), two goods ($L = 2$), and the (identical) utility functions of the consumers are given by

$$W_i = m_i - 2m_k, i \neq k \text{ and } m_i(x_i) = h(x_{i1}) + h(x_{i2}), h' > 0, h'' < 0. \quad (6.39)$$

The aggregate endowment of each of the two goods is 1. The allocation $((0, 0), (0, 0))$ is not a Walrasian equilibrium of the internal economy' $\mathcal{E}^{Int}$ (Definition 6.18) because preferences are strictly monotone, yet no consumer consumes anything. However, given (6.39), $((0, 0), (0, 0))$ is Pareto optimal.

Example 6.8 indicates that additional restrictions on preferences are required for the standard welfare theorems to hold in the presence of other-regarding preferences. In particular, we would like to rule out the pathological case of Example 6.8.

Definition 6.21 (Social monotonicity): Consider an allocation, x . Then there exists another allocation z such that $z \geq 0$ and²⁶

$$W_i(m_i(x_i + z_i), m_{-i}(x_{-i} + z_{-i})) > W_i(m_i(x_i), m_{-i}(x_{-i})),$$

where $\sum_{k=1}^I z_k = z$ and at least one element of the vector z is strictly positive.

Definition 6.18 requires that redistributing an extra amount of resources to all consumers increases everyone's utility. This rules out the extreme case of spite in Example 6.8.

Proposition 6.20 (Dufwenberg et al., 2011): Suppose that the utility function exhibits well-being externalities (Definition 6.19) and the social monotonicity condition (Definition 6.21) holds.

- (i) The set of Pareto optimal allocations of the economy $\mathcal{E}$ is a subset of the Pareto optimal allocations of the economy $\mathcal{E}^{Int}$.
- (ii) Every Pareto optimal allocation of the economy $\mathcal{E}$ can be obtained as a Walrasian equilibrium if lump-sum transfers are allowed.

Proposition 6.20(i) shows the relation between the Pareto optimal equilibria of economies with other-regarding and self-regarding preferences. Proposition 6.20(ii) is the analogue of the second

²⁶ We use here the notation $\geq$, which applies the *greater than or equal to* operation to elements of vectors.

welfare theorem but for other-regarding preferences. The social monotonicity condition by itself does not ensure that Pareto efficient equilibria can be implemented as a Walrasian equilibrium unless lump-sum transfers are allowed; the next example discusses this point.

Example 6.9 *Suppose that there are three consumers and one good. The endowment of consumer 1 and consumer 3 is one unit each of the good, while consumer 2 has zero units (so $e = (1, 0, 1)$). Consumer 2 has purely self-regarding preferences. The utility of consumer 1 is $W_1 = x_1 + \frac{2}{3}x_2$ and the utility of consumer 3 is $W_3 = x_3 + \frac{2}{3}x_2$. Suppose that we allow for voluntary lump-sum transfers of goods between the three consumers and the social monotonicity condition (Definition 6.21) holds. Consumers 1 and 3 do not want to transfer any of the good to consumer 2 because giving up a unit costs 1 but gives a marginal benefit $\frac{2}{3} < 1$. Hence, the economy is stuck with the initial allocation $e = (1, 0, 1)$, yet it is Pareto dominated by the allocation $(0, 2, 0)$. Thus, no Walrasian equilibrium can implement the Pareto optimal allocation $(0, 2, 0)$.*

6.7 Evidence on models of social preference and reciprocity

In this section we consider first the evidence on models of social preferences (Section 6.7.1) followed by the evidence on the role intentions (Section 6.7.2).

6.7.1 The evidence on models of social preferences

One criticism of the inequity aversion models is that they do not specify the reference group over which an individual compares his payoff (Sobel, 2005). However, the reference group can be elicited from subjects in experiments and alternative assumptions on the reference group can be experimentally tested. Given individual heterogeneity, it may be undesirable for the theory to specify a reference group. Some individuals may simply compare their consumption with those on the street that they live in, others with colleagues in the workplace, still others with the population of subjects in similar professions, and so on.

Some models introduce the possibility that individuals care for the total economic surplus (Charness and Rabin, 2002). A natural way to test the relative concerns for surplus maximization and inequity aversion is to construct experiments that set these two opposing concerns against each other.²⁷ The main result of these papers is that in dictator games, a non-negligible fraction of proposers are willing to take an action that increases the total surplus, even at a cost in terms of personal payoff, and even if this increases disadvantageous inequity for them.

Andreoni and Miller (2002), and Andreoni and Vesterlund (2001) vary the price at which the proposer can transfer a unit of money to the responder in dictator games. Normally, in dictator games, by giving up a unit of money, the dictator can transfer a unit of money to the receiver (the 1:1 case). In the 1:1 case, experimental evidence shows that the dictator never proposes to transfer more than half the amount. Now suppose that we reduce the price of giving by using other ratios such as 1:2 (i.e., for every unit given up by the dictator, 2 units are transferred to the responder), and 1:3. In these cases, Andreoni and Miller (2002) and Andreoni and Vesterlund (2001) find that in a non-negligible fraction of cases, the dictators end up with less money relative to the

²⁷ See, for instance, Andreoni and Miller (2002), Andreoni and Vesterlund (2001), Bolle and Kritikos (1998), Charness and Rabin (2002), Cox (2000).

Table 6.2 Payoffs of three players in three different allocations.

Allocation	A	B	C
Person 1	9	8	11
Person 2	8	8	10
Person 3	4	8	9
Total	21	24	30

Source: Engelman and Strobel (2004).

receiver, but the total surplus between the two parties is higher. Such results are not consistent with models of inequity aversion models such as the ERC model, the Fehr–Schmidt model, and the composite model of Falk and Fischbacher (2006). There is substantial heterogeneity at the level of the individual. Andreoni and Vesterlund find that 44% of the subjects are completely self-regarding, 35% seem to have egalitarian preferences (they try to equalize payoffs), and only 21% are surplus maximizers.

Charness and Rabin (2002) consider the following two payoff-pairs in a dictator game: (400, 400) and (400, 750), where the first figure represents the dictator's payoff, while the second figure is the responder's payoff. Of the subjects 31% are found to be egalitarian, which is similar to the finding in Andreoni and Vesterlund (2001). To check for the presence of individuals who might not be willing to make the surplus maximizing choice at a cost to their own payoff, consider another choice condition: (400, 400) and (375, 750). In this case, only 49% make the surplus maximizing choice, because it reduces own-payoff by 25. When income disparities are increased further, e.g., in the choices (400, 400) and (400, 2000), then 62% make the egalitarian choice in preference to the surplus maximizing choice.

Iriberry and Rey-Biel (2013) find that in dictator games, there is heterogeneity in preference types (self-regarding, social welfare maximizing, inequity averse, and competitive); almost half the subjects are self-regarding. We have noted above, reservations about the use of the dictator game results to draw robust conclusions about social preferences (Fehr and Schmidt, 2006). However, an interesting element of the results in this paper is that people alter their preference types depending on the available social information (e.g., about the actions of others); this is suggestive of interdependent preferences. There is also substantial heterogeneity in beliefs about others that is dependent on one's preference type.

Engelmann and Strobel (2004) present the following example that illustrates concerns for efficiency and inequity aversion. There are three individuals and three possible allocations, A, B, and C; Table 6.2 shows the payoffs of the three players in the three different allocations. Suppose that individual 2 is asked to choose between allocations A and B, which give the individual an identical payoff of 8. Inequity aversion theories predict that individual 2 will choose allocation B over A because it has a more equitable payoff. However, there is another potential explanation. Allocation B is more efficient as compared to A (joint surplus being 24 rather than 21). So it might well be the case that what seems to be a choice based on inequity aversion is really a concern for efficiency. This is an extremely stringent test for inequity aversion theories because the more equitable allocation is also the surplus maximizing one.

Now suppose that individual 2 has a choice between all three allocations, A, B, and C. Suppose that allocation C is chosen. This will provide support for selfishness (2's payoff is highest

in allocation C), efficiency (joint surplus is highest in allocation C), inequity aversion,²⁸ and maximin preferences (the poorest individual is best off in allocation C). Engelmann and Strobel (2004) perform experiments to disentangle these three effects, namely, efficiency, inequity aversion, and maximin preferences. They conclude that a concern for efficiency and Rawlsian preferences, jointly, is more important, relative to a concern for inequity aversion. This result supports the Charness and Rabin (2002) model (see Section 6.8), over models of inequity aversion.

Preferences for efficiency and Rawlsian maximin, relative to inequity aversion were discussed in a series of papers in the December 2006 issue of the *American Economic Review*. Bolton and Ockenfels (2006) report fresh experiments to show that twice as many individuals sacrifice an outcome where they get a higher payoff for an equitable outcome as compared to an efficient one. Engelmann and Strobel (2006) argue that a potential confounding factor in these experiments is that the gains from efficiency are lower than the gains from equity.

Fehr et al. (2006), however, show that the results of Engelmann and Strobel (2004) are driven by the presence of economics students who learn about the importance of economic efficiency in their first year. After replicating several experiments in Engelmann and Strobel, they find that the dominance of the efficiency motive holds true only for economics and business administration students. Non-economist subjects find inequity aversion more attractive relative to efficiency.

The evidence on the surplus maximizing motive comes largely from dictator games in which only one player has a non-trivial action. There are several issues with using the dictator game for detecting the presence of a surplus maximization motive (Fehr and Schmidt, 2006). First, other than charitable donations, the dictator game is not a particularly important analogy for most other real-world economic situations. Second, as we have seen above, the results of the dictator games, unlike those of the ultimatum games, are not very robust, hence, this may not be a particularly suitable game to test surplus maximization. Third, there are end game effects (as in Andreoni and Miller, 2002), where the predictions of a surplus maximization preference are not borne out in the final periods of the experiment. These final period effects can be accounted for by models of inequity aversion and intention-based reciprocity. Fourth, what if we introduced the possibility of an element of strategic consideration in which more than one player has a non-trivial choice? Is there a reduction in the surplus maximizing choice in this case? The experiments of Fehr et al. (1993, 1998) that we consider next are important in this regard.

Fehr et al. (1993, 1998) embed a gift exchange game in a strategic environment. A proposer (firm) makes a wage offer that is observed by the responder (worker) who then chooses a non-binding effort level. Since the effort level is non-binding, the responder has a non-trivial choice that can affect the total surplus in the relation. An increase in the effort level by the responder increases the total surplus by five units on average. If surplus maximization is the overriding objective, certainly more important than outcome inequality and reciprocity, then we should observe the responder putting in the maximum possible effort level. Furthermore, we should not observe any correlation between wages and effort levels. Both these implications are rejected by the evidence. In no more than 2% of the cases does the responder undertake an effort level such that the proposer's payoff is relatively greater. There is also a clear correlation between the wage offered and the effort level undertaken. The actual effort level undertaken by the responder is significantly below the maximum feasible effort level.

When subjects punish others in public goods games in the lab, there are two possible motives. First, subjects may wish to act in a *reciprocal manner*. For instance, contributors may punish

²⁸ The reader can easily check that, for person 2, the Fehr-Schmidt utility is higher for distribution C as compared to distribution A.

non-contributors. A second motive may be to equalize payoffs. In a nice experimental design, Dawes et al. (2007) offer a way of separating the two motives. They replace the first (cooperation) stage in a Fehr–Gächter type cooperation and punishment game with a random allocation of incomes to the players that mimics the income distribution in Fehr and Gächter (2002) after the first stage. Then they give the players in stage 2, a possibility to reward or punish others in the group by assigning reward or punishment tokens to them. Interestingly, players with below-average incomes are often rewarded while players with above-average incomes often get punished. In this set up, reciprocity can play no role at all but FS or ERC-type inequity aversion can. However, ERC cannot explain this pattern of targeting rewards and punishments. This result suggests powerfully that inequity aversion (apart from reciprocity) as in the Fehr and Schmidt (2007) model is a powerful force in sanctioning behavior in public goods games.

6.7.2 *The evidence on intentions and reciprocity*

The *state of mind* of the accused and his/her *intentions* are of crucial importance in most legal systems when dealing with murder. The law, in most countries, defines *murder* as the *intentional* killing of another human being with full premeditation. On the other hand, the killing of a human carried out under some form of *diminished responsibility* that reflects *less than full intentionality* is treated differently.²⁹ In particular, the killing of a human where there is a clear lack of *intent* to kill is termed as *manslaughter* and it carries a different penalty from premeditated murder.

Murder and manslaughter are two extreme categories with several others in between. For instance, in the Canadian legal system, murder is further classified into first and second degrees based on *intent* of the perpetrator. First degree murder is the intentional killing of a human being with premeditation, while second degree murder does not involve premeditation, e.g., killing in the heat of the moment, but nevertheless, with intent to kill, which distinguishes it from manslaughter.³⁰

The *inequity aversion* models of Fehr and Schmidt (1999) and Bolton and Ockenfels (2000) are purely outcome oriented, and ignore the role played by intentions. In many contexts, one's behavior is determined by beliefs about the intentions of others. For instance, charitable giving in which one does not know the identity of the receiver of a gift might well be motivated by issues of inequity aversion. On the other hand, a firm contemplating a wage cut might be concerned that workers may form negative inferences about its intentions and may respond with negative reciprocity in the form of low effort.

An early experiment that considered the role of intentions in the ultimatum game was conducted by Blount (1995). In ultimatum games, if the initial offer is made by a computer generated random number (random offer condition) or a third party (third party condition), then the offer does not reveal any information about the intentions of the proposer. Hence, if intentions-based motives are the only explanation behind the results of the ultimatum game, we should observe no rejections in these two cases. Relative to the benchmark case, where a human proposer makes an offer, Blount (1995) finds that the rejection rates of responders decline significantly in the random offer condition. However, there is no significant change in the third party condition. The first result provides support for intentions-based models. But the second result (no significant

²⁹ Several factors may indicate diminished responsibility. These include, insanity, depression, side effects of certain medications, and even post-natal depression (in the case of murder of a child by the mother), etc.

³⁰ Intentions can also be transferrable, within the law, in the following sense. In English law, transferred malice occurs when an individual, say A, is attacked with an intention to kill but in fact another individual, say B, is killed. The perpetrator, say individual C, can be charged with the murder of individual B.




Figure 6.3 Modified ultimatum game.

Source: "On the nature of fair behavior," Falk, A., Fehr, E., and Fischbacher, U., *Economic Inquiry*, 41(1): 20–26. Copyright © 2007, John Wiley and Sons.

change in the third party condition) contradicts intentions-based explanations and is consistent with models of inequity aversion.³¹ Using the ultimatum game, Offerman (2002) finds that intentional low offers are 67% more likely to be followed by actions of opponents that are payoff reducing. Kagel et al. (1996) also find that the rejection rates in ultimatum games are higher when intentionally unequal proposals are made.

Falk et al. (2003) consider the modified ultimatum game in Figure 6.3 to isolate the role of intentions. The proposer, labeled by P, has two choices, 1 and 2, in offering a split of 10 dollars. The responder, labeled by R, has a choice to either accept (action a) or reject (action r). The payoffs are shown at each end node. Two variations of the game are considered. In the first variant, game-A, the payoffs from the strategy pair (2, a) are (5, 5), while in the second variant, game-B, the payoffs from the strategy pair (2, a) are (10, 0). In all other respects, the payoffs from the two games are identical.

Consider game-A and let the proposer choose strategy 1. In the absence of reciprocity, and based purely on self-interest, the best response of player 2 is to accept, resulting in the payoff pair (8, 2). However, in the presence of reciprocity and intentions, while deciding to accept or reject the offer, the responder brings into consideration the fact that the proposer has not chosen strategy 2, a relatively fairer strategy, which would have ensured payoffs of (5, 5). This reveals to the responder that the proposer's intentions *could have been fairer*.

Now consider game-B. If the proposer chooses strategy 1 over strategy 2, the responder knows that the proposer gave up a likely payoff of 10 and chose a payoff of 8 instead. This reveals to the responder that the intentions of the proposer *could have been less fair*. Hence, if intentions are important, then the responders in game-B should be more likely to accept the proposer's offer of an 8–2 split, despite the fact that the outcome (8, 2) is identical in both games. The essential feature of reciprocity is not necessarily to alter the distribution of payoffs but to respond to kindness and unkindness with like behavior.

Falk et al. (2003) find that the 8–2 split is rejected in 44.4% of the cases in game-A but it is rejected in only 8.9% of the cases in game-B. This runs counter to the predictions of the inequity

³¹ There were some peculiarities in these experiments that may or may not account for these conflicting results. For instance, in two out of the three treatments, subjects made decisions in the roles of proposers and responders but the payment they received was for one of these two roles, which was only subsequently revealed to them. Furthermore, in one treatment, deception was involved—subjects thought they were making the decisions as proposers but, in fact, the experimenter made the proposals.

aversion theories, which would predict that the rejection rates should not be different in the two games because the outcome (8, 2) is identical in both. However, the fact that there should be any rejection rate at all for game-B suggests that the reciprocity approach based on intentions does not provide a full explanation either. Clearly more empirical work in this area will provide a fuller picture.

One possible implication of intentions driven behavior for public goods games is the following. The extent of punishment meted out by cooperators to non-cooperators should depend on the difference in their contributions, which reveals the extent of unkindness of non-cooperators. Fehr and Gächter (2002) found this to be the case in their public goods experiments. They found that the single most important predictor of the punishment meted by contributors to non-contributors was the difference in contributions between the two players. Furthermore, 75% of all punishments were carried out by contributors whose contributions were above average.

In order to isolate the role of intentions, Falk et al. (2008) use the moonlighting game, which is due to Abbink et al. (2000). There were 112 subjects in the experiment. The moonlighting game allows one player, say *A*, to be kind or unkind in the first stage and then, in the second stage, allows another player, say *B*, to reciprocate with rewards or punishments. In particular, each of the players starts with an initial endowment of 12 monetary units of experimental money (MU). Player *A* chooses some action $a \in \{-6, -5, \dots, 0, \dots, 5, 6\}$ in the first stage. The interpretation of *A*'s action is that if $a > 0$, then player *B* receives a tokens (kind intentions), while if $a < 0$, then player *B* must surrender a tokens to *A* (unkind intentions). In the second stage, conditional on *A*'s action, player *B* takes some action $b \in \{-6, -5, \dots, 17, 18\}$. If *B* chooses $b > 0$, then *B* gives an amount b to *A* (reward) and if $b < 0$, then *B* loses b , while *A* loses $3b$ (punishment).

The experiment allows for two treatments. In the *intentions treatment* (I-treatment), players choose their actions freely, so intentions can be inferred from actions. In the *no-intentions treatment* (NI-treatment), a random device is used to choose *A*'s action, so intentions cannot be inferred. The ex-ante predictions are as follows.


1. If both *A* and *B* have self-regarding preferences, then the solution, using backward induction, is simple. In the second stage, when it is *B*'s turn to move, bygones are bygones. *B* maximizes her payoff by choosing $b = 0$. Anticipating *B*'s actions, *A* takes maximum advantage by choosing $a = -6$.
2. Suppose individuals care only about outcomes and ignore intentions, as in the inequity aversion models of Fehr and Schmidt (1999) and Bolton and Ockenfels (2000). Then, a higher action by player *A*, reduces *A*'s payoff. Since player *B* dislikes inequality, she should respond by contributing more. Thus b should be increasing in a . In particular, when $a = 0$, then $b = 0$. However, since there is no role for intentions, we do not expect any difference among the I-treatment and the NI-treatment.
3. Suppose now that only intentions (but not outcomes) matter.³² One can construct plausible equilibria such that a larger contribution, a , by player *A*, signals better intentions. In this case, in the I-treatment, we expect player *B* to reciprocate *A*'s action, so b is predicted to be increasing in a . In the NI-treatment, however, intentions cannot be inferred, so the model predicts $b = 0$ for all actions taken by player *A*. Finally, in models that take account of both intentions and outcomes, e.g., Falk and Fischbacher (2006), one expects b to be increasing in a . However, when intentions are revealed, as in the I-treatment, then we expect b to respond even more strongly to a on account of reciprocity. The predictions are summarized in Table 6.3.

³² Dufwenberg and Kirchsteiger (2004) and Falk and Fischbacher (2006) consider the role of intentions and outcomes; see the material on psychological game theory in Part 4 of the book.

Table 6.3 Summary of predictions for player B.

Model	I-treatment	NI-treatment
Standard prediction	$b = 0, \forall a$	$b = 0, \forall a$
Only payoff consequences matter (Fehr–Schmidt and Bolton–Ockenfels)	b increases in a	exactly the same behavior as in I-treatment
Only fairness intentions matter (Dufwenberg–Kirchsteiger)	b increases in a^0	$b = 0, \forall a$
Payoff consequences and fairness intentions matter (Falk–Fischbacher)	b increases in a	b increases in a but less than in the I-treatment

Source: Falk et al. (2008).

**Figure 6.4** Effect on b as a changes.

Source: Reprinted from *Games and Economic Behavior*, 62(1): 287–303. Falk, A., Fehr, E., and Fischbacher, U., “Testing theories of fairness-intentions matter.” Copyright © 2008, with permission from Elsevier.

The experimental results are shown in Figure 6.4. The horizontal axis is player A’s action, a , while the vertical axis is player B’s action, b , in response to A’s action. In the I-treatment, average and median rewards by player B are increasing in the action of player A. This contradicts the prediction of self-regarding preferences, $b = 0$ for all a . The median response in the NI-treatment is exactly in line with the prediction of the self-regarding preferences ($a = b = 0$). The difference in the intentions and the no-intentions treatment is statistically significant and intentions matter

in the domain of both positive and negative reciprocity (recall that a can be negative or positive). Thus, the results also reject the “outcomes only” models.

Regression results also show that reciprocity is not completely wiped out in the NI-treatment. This is seen by running a regression of b on a constant term, α_0 , the action a of player A, a dummy variable d that equals 1 for the I-treatment and zero otherwise, and, finally, the interaction term, $a \times d$. The sum of the coefficients on the constant term and on the action “ a ,” give the additional amount of b that is elicited in the NI-treatment. This is positive and statistically significant, which is not consistent with the prediction in Dufwenberg and Kirchsteiger (2004). In conclusion, the aggregate evidence is that outcomes and intentions are both important, hence, more general models such as the one by Falk and Fischbacher (2006) have a better chance of succeeding.

The experiment also generated data at the disaggregated level, i.e., on individual choices. The percentage of subjects in the role of player B whose behavior corresponds to self-regarding preferences ($b = 0$ for all a), increases from 0 to 30 as one goes from the NI to the I treatments. Hence, a substantial part of the behavior is explained by fairness intentions. The percentage of subjects who exhibit positive or negatively reciprocal behavior, reduces from 76 to 30 as one moves from the I to the NI treatment. But importantly, reciprocity is not completely wiped out even in the NI treatment. About 40% of the subjects continue to exhibit reciprocal behavior. This supports the assertion that intentions alone do not drive human behavior; the concern for fair outcomes also plays an important role.

One set of experiments finds little evidence that fairness intentions play an important role in the domain of positively reciprocal behavior (Bolton et al., 1998; Cox, 2000; Charness, 2004), while another set of results suggests that intentions are important in the domain of negatively reciprocal behavior (Blount, 1995; Offerman, 2002). By contrast, the experimental results of Falk et al. (2008), described above, find clear evidence that reciprocal intentions are important in the domain of positively and negatively reciprocal behavior. What accounts for the differences in results?

Falk et al. (2008) raise four main concerns about earlier experimental work in explaining the differences from their results. First, in some of the experiments, a significant number of the subjects had an efficiency motive in the sense that they were willing to tolerate a reduction in their own payoffs to increase the total size of the cake. This may have confounded the reciprocity motive. Second, in some experiments, the repeated trials could have introduced noise and spillovers across the periods, making it difficult to isolate intentions. Third, in some experiments, the fairness or unfairness of actions is not salient enough to signal intentions of subjects. Fourth, the intentions and the no-intentions treatments in some of the experiments might not have been well differentiated.

Cox (2004) considers variants of the trust game. In treatment 1, which is similar to the moonlighting game above, Player A is given \$10 to pass on to player B, via the experimenter. The experimenter triples the amount that A has offered, before passing it on to B. Player B can then return any amount to player A. In treatment 2, Player B does not have the opportunity to return any money to player A. Finally in treatment 3, player A does not move, but the experimenter makes a move on behalf of player A, passing on the same amount as in treatment 1. The main difference between treatments 1 and 3 is that player A’s intentions can be inferred in treatment 1 but not from treatment 3. So, if players were purely inequity averse, then one should not observe any difference between the results of the two treatments. The experimental results show that player B returns more money in treatment 1 and, in particular, returns more when player A offers more money, hence, demonstrating the role of intentions and reciprocity.

Bartling and Fischbacher (2012) consider a four-player dictator game. There are two possible dictators, players A and B, and two possible receivers, players C and D. There are only two possible allocations. The *fair allocation* in which each player gets a payoff of 5, and the *unfair allocation* in which players A and B get a payoff of 9 each and players C and D get a payoff of 1 each. In the no-delegation treatment, player A has to directly choose among the fair and the unfair allocations. In the delegation treatment, player A delegates player B to make the allocation decision (player B cannot refuse). In the *punishment option*, player C observes the chosen allocation and can then engage in punishing players A and B. In the *no-punishment option*, player C cannot punish.

Under pure inequity aversion, the choice of delegation by A should not matter to the punishment decision of player C. However, the choice of allocation and the choice of delegation both convey important information about the intentions of each of the players A and B. While such intentions can be difficult to pin down, the paper makes progress in this direction. For instance, in one treatment, called *asymmetric*, A can either choose the fair allocation (but not the unfair allocation), or delegate the choice to B, who can then choose among the fair and the unfair allocations. In this case, delegation by A clearly reveals unkind intentions because there is some likelihood that B may choose the unfair allocation, so player A gets 9 instead of 5, while players C and D get the low payoff of 1 each.

The main findings are as follows. When the fair option is chosen by either A or B there is no punishment. The dictator (A or B) who chooses an unfair allocation is punished more and there is very little punishment of the other player. Thus, one observes more delegation by player A to player B when the punishment option by player C is available. Thus, by delegating to B, player A also shifts most of the punishment to B if an unfair allocation is chosen.

When player A delegates to player B who then chooses an unfair offer then player A is punished relatively more as compared to the case when he directly chooses the unfair allocation. In this case, C may believe that A could have chosen the fair allocation but chose instead to delegate to B who chose an unfair allocation, hence, revealing possibly unkind intentions on the part of A. As compared to the symmetric case (where A can choose among the fair and unfair allocations), punishment of player A by player C is significantly higher in the asymmetric treatment if B chooses the unfair allocation. The reason is that this clearly reveals, to player C, the unkind intentions of player A.

INTENTIONS, INEQUITY AVERSION, AND FAIRNESS REFERENCE POINTS

Most people would concede the importance of the *intentions motive* in influencing the actions of players. However, the relative importance of the *intentions motive* (the *attribution hypothesis*) and other-regarding preferences (the *distribution hypothesis*) has been an area of some conflicting views; see Bolton et al. (2008). See also the comments by Falk et al. (2008) on the possible reasons for the differences in results between the early and the later work.

The early literature found it difficult to produce evidence on the sole importance of the attribution hypothesis. The evidence is probably better described by a combination of the attribution hypothesis and the distribution hypothesis. However, in this section, we shall see that some of the evidence begs for the introduction of even newer theoretical considerations such as a *procedural approach to fairness* and/or a better understanding of a *fairness reference point*.

We have considered issues of a fairness reference point in relation to the work of Gneezy and List (2006) and Fehr et al. (2009), elsewhere in this chapter. Bolton et al. (2005) consider experimental evidence based on the ultimatum game that may form the basis of theoretical work on a procedural approach to fairness. The main idea of these experiments is that individuals are willing to accept an unfair outcome, provided it is the outcome of a procedure that is agreed to be fair. The authors give two insightful examples of the importance of their approach.

Example 6.10 In a court case in 1842 (*US v. Holmes*), crew members of a ship were found guilty of throwing 14 non-crewmen overboard from a leaky and crowded lifeboat. However, the guilty verdict centered not on throwing the 14 non-crewmen overboard (perhaps this was unavoidable) but rather on the unfairness of the “procedure” that excluded everyone from the crew to be thrown overboard. The judge ruled that the victims should have been chosen by a fair lottery that included both crew and non-crew members.

Example 6.11 A consortium of eight US states (the Southeast Compact) choose North Carolina as a destination for a radioactive waste management facility, using a lottery method. The North Carolina legislature accepted the decision on the grounds that they were chosen based on a fair procedure.

Existing theories of social preferences and intentions are neutral towards the perceived fairness of procedures that govern outcomes. Rather, their interest centers around the fairness of outcomes and the kindness of intentions. Clearly, if the outcomes are sensitive to the perceived fairness of the procedure, then we need a broader model; the experimental evidence in Bolton et al. (2005) makes such a case.

Recall that Blount (1995) produced mixed evidence for the role of intentions—rejection rates in ultimatum games were very similar in the baseline and the third party treatments, which actually contradict the intentions motive. Another set of experiments, conducted around the mid 1990s and reported in Charness (2004), applied Blount’s idea to a gift exchange game.³³ Consistent with the attribution hypothesis, he found a positive correlation between the actions of two players (gift exchange). However, when the first mover’s action are allocated to a third party or chosen randomly, so intentions are impossible to infer, the contributions by the second player are even higher. This contradicts the attribution hypothesis.

Similar difficulties in finding support for the intentions hypothesis were encountered by Bolton et al. (1998), which is an early, published, experimental contribution. They consider a two-player sequential game; players observe a 3×6 normal form game (see Figure 6.5). However, as explained below, they actually play a 2×6 normal form game in three different treatments. The row player always has the choice of the middle row (m), in each treatment. But in three different treatments, he can choose one additional row from among the top (t) or bottom (b) rows. The choice of the additional row reveals important information about the intentions of the row player. The column player chooses between columns, c1 to c6, and has a dominant action, c1, irrespective of the choice of the row player. However, by successively choosing actions c2, c3, . . . , c6 he can, at some personal cost to his payoffs, increase the payoff of the row player, such that the total surplus

	c1	c2	c3	c4	c5	c6
t	800, 2050	1000, 2000	1200, 1950	1400, 1900	1600, 1850	1800, 1800
m	900, 1650	1100, 1600	1300, 1550	1500, 1500	1700, 1450	1900, 1400
b	1000, 1250	1200, 1200	1400, 1150	1600, 1100	1800, 1050	2000, 1000

Figure 6.5 Normal form representation of a two-player, sequential, social dilemma game.

Source: *Experimental Economics*, “Measuring motivations for the reciprocal responses observed in a simple dilemma game,” 1(3), 1998, 207–19, Gary Bolton, Jordi Brandts, and Axel Ockenfels. © Economic Science Association. With permission of Springer.

³³ Earlier versions of these results circulated in a working paper, circa 1996.

increases (social dilemma game). Once each player chooses his actions, payoffs are determined as in Figure 6.5.

In the baseline treatment, treatment II, the row player picks his second row to be *m*, thus, the choice by the row player is irrelevant. This is the *no free choice treatment* because the intentions of the row player cannot be inferred.

In treatment I, the row player picks his second row to be *b*. This is the *positively reciprocal treatment* because by choosing *m* over *b*, the row player can improve the payoff of the column player, should the latter play his dominant action, *c1*. The column player can now, if he wishes, reward the kind intentions of the row player by a choice that is different from *c1*. The higher the chosen column number, the greater is the reward for kind intentions; we shall refer to this as the degree of *cooperative behavior* of the column player; this is the horizontal axis in Figure 6.6. One can then compare the results in treatments I and II, say, when player 1 chooses *m*, to determine the extra contribution of intentions to the outcome. Finally, in the *negatively reciprocal treatment*, treatment III, the row player chooses row *t* as his second row (so the choice is between the two rows *t* and *m*). In this case, by choosing *m*, the row player signals to the column player that he has unkind intentions. Again, one may compare the choices with treatment II to infer the marginal effect of intentions.

The results are shown in Figure 6.6 in terms of the cumulative frequencies of the cooperative choices of the column player in each treatment. Thus, a higher graph indicates a less cooperative choice (this just follows from first order stochastic dominance). The choices by the column player, particularly departures from action *c1*, can be used to infer reciprocal behavior in both reciprocity treatments, I and III. Notice, from Figure 6.6, that in the no free choice treatment (treatment II), about half the column players choose an action higher than *c1*. Thus, there is a significant element of other-regarding preferences that can be noiselessly isolated from any reciprocity considerations. However, when the three distributions in Figure 6.6 are compared using the relevant statistical tests, one cannot reject the null hypothesis that the contributions of the column player belong to the same underlying distributions. Nor can we reject similar null hypotheses when the contributions under any two treatments are compared in a pairwise manner. The authors conclude (p. 213): “Hence our statistical tests imply that all of the observed reciprocal action is attributable to the distribution hypothesis.”

Bolton and Ockenfels (2005) observe variations in the results from ultimatum games, say, the probability of responder rejections, as one alters the parameters in the game; they refer to this as




Figure 6.6 Cumulative distribution of the second player's cooperative decision (*c1* to *c6*).

Source: Reprinted from Bolton, G. E., Brandts, J., Katok, E., Ockenfels, A., and Zwick, R., "Testing theories of other-regarding behavior: A sequence of four laboratory studies," in C.R. Plott and V. L. Smith, eds., *Handbook of Experimental Economics Results*. Copyright © 2008, with permission from Elsevier.




Figure 6.7 Two variations of the ultimatum game, indexed by the parameters x and y . The UG^*x game is in the left panel and the UG^*+y game is in the right panel.

Source: *Economic Theory*, 25(4): 957–82. “A stress test of fairness measures in models of social utility,” Bolton, G. and Ockenfels, A. © Springer-Verlag 2005. With kind permission from Springer Science and Business Media.

stress testing the results. The two variations are shown in their multiplicative and additive forms, respectively, in the left and right panels of Figure 6.7. A proposer has a cake of size \$4 that he can split with the responder. The proposer has two actions; action *l* is the more equitable choice, while action *r* is less equitable. The responder observes the action chosen by the proposer, then chooses to accept, *a*, or reject, *r*.

Consider first the multiplicative game in the left panel of Figure 6.7. In this game, stress testing involves varying the values of $x = 0.2, 0.6, 1$ and 1.2 ; each of these games is known as UG^*x . The payoffs of both players are not required to sum up to \$4. When $x = 1$ this is a standard ultimatum game. Notice, in particular, that the parameter x does not influence the payoff of the responder.

Under self-regarding preferences, the subgame perfect outcome of any UG^*x game is (r, a) . Given the *kindness functions* used by the authors,³⁴ purely intentions driven models predict that the size of x should not influence the rejection rates of the responder. However, the magnitude of x does influence the relative payoffs of the two players, hence, models of other-regarding preferences predict variability of rejection rates with respect to x .

The results on the rejection rates of responders in the UG^*x game are shown in Figure 6.8. There are no rejections when the proposer plays the strategy *l*. However, when the proposer plays the strategy *r*, rejection rates rise sharply from zero in the $UG^*0.2$ game to 43% in the $UG^*0.6$ game but, thereafter, these rates stabilize. This is inconsistent with the attribution hypothesis.

What explains the sharp rise in rejections between the $UG^*0.2$ game and the $UG^*0.6$ game? Consider the following explanation. When $x = 0.2$, both players get an equal payoff of 0.6 when the proposer chooses *r*. However, when $x > 0.2$, we get highly unequal payoffs from the action profile (r, a) in any UG^*x game. The responder gets 23% of the cake in $UG^*0.6$ and 13% in $UG^*1.2$. One may conjecture that starting from a 50% share of the cake in $UG^*0.2$, a 23% share of the cake appears highly unattractive, hence, rejections increase dramatically between $UG^*0.2$ and $UG^*0.6$. However, the difference between a 23% and 13% share of the cake is less dramatic, hence, it elicits a smaller difference in rejection rates between $UG^*0.6$ and $UG^*1.2$.

In the UG^*x game we varied *distributive measure of fairness* and held the *kindness measure* fixed. We now flip these roles in the UG^*+y game in the right panel of Figure 6.7, where $y \in$

³⁴ For a fuller discussion, see the discussion on psychological game theory in Part 4.




Figure 6.8 Average rejection rates in the UG^*x game.

Source: *Economic Theory*, 25(4): 957–82. “A stress test of fairness measures in models of social utility,” Bolton, G. and Ockenfels, A. © Springer-Verlag 2005. With kind permission from Springer Science and Business Media.




Figure 6.9 Average rejection rates in the UG^*+y game.

Source: *Economic Theory*, 25(4): 957–82. “A stress test of fairness measures in models of social utility,” Bolton, G. and Ockenfels, A. © Springer-Verlag 2005. With kind permission from Springer Science and Business Media.

$\{-1.6, 0, 0.8, 1.6\}$. The subgame perfect outcome is (r, a) . Since the payoffs following the choice of r by the proposer are independent of y , the distribution hypothesis predicts that the rejection rates must be identical in any UG^*+y game. The rejection rates following a choice of l by the proposer are predicted to fall between the games $UG^*-1.6$ and UG^*+0 . However, the distribution hypothesis cannot predict the rejection rates following a choice of l by the proposer as y increases beyond zero. The reason is that the absolute payoff increases but the relative payoff falls (think, e.g., of the Fehr–Schmidt model).

By contrast, the intentions-based hypothesis (attribution hypothesis) predicts that as y increases, the rejection rates following a choice of r by the proposer should increase. The reason is that the proposer leaves only 0.6 for the responder by playing r when he could have left a lot more by playing l . This is the sense in which the *kindness measure* is varied in the UG^*+y game.

The average rejection rates are shown in Figure 6.9. As predicted by the distribution hypothesis, the rejection rates following a choice of l by the proposer fall sharply between the games $UG^*-1.6$

and UG^*+0 . However, the rejection rates following the choice of r by the proposer are different in the various UG^*+y games. Hence, there is mixed support for the distributional hypothesis. Following a choice of r by the proposer, as y increases from -1.6 to 0 , the rejection rates increase; this supports the attribution hypothesis. However, as y increases beyond 0 , the rejection rates fall, contradicting the attribution hypothesis. In a strict sense, these results reject both distribution-only and attribution-only hypotheses.

What accounts for the rejection rates in the UG^*+y game? One plausible conjecture is that the responder's notion of what is *fair* depends on y . When $y = 0$, the payoffs from the strategy pair (l, a) are $(2, 2)$, which may possibly serve as a *reference point for fairness*. In this case, the rejections following a choice of r by the proposer that give rise to unequal payoffs are maximal (Figure 6.9). As y increases, the payoffs from the strategy pair (l, a) , given by $(2 - y, 2 + y)$, become increasingly unequal. Thus, the new reference point for fairness may have become more unequal, prompting a greater acceptance rate following the choice of r by the proposer. Indeed, when $y = -1.6$, the rejection rate following a choice of r by the proposer is the lowest. This suggests an alternative approach to the distributional and the attribution hypotheses in which there is an explicit consideration of the reference point for fairness.

6.8 A discussion of other promising models

In this section, we consider other models of fairness and reciprocity such as the hybrid model of Charness and Rabin (2002), and models of interdependent preferences. We describe the relevant theory and also pit the models against the evidence.

6.8.1 The Charness and Rabin (2002) hybrid model

Charness and Rabin (2002) argue that experiments that test for social preferences have been unduly restrictive in the sense of giving too few options to the subjects in experiments. For instance, in the context of experiments in the ultimatum game, they argue that (1) a responder only has the option to accept or reject, which rules out other interesting choices such as “social objectives,” and (2) the structure of the game ensures that the only retaliatory instrument (rejection of the responder's offer) reduces inequality between the two players. They design experiments that allow for a larger range of choices for subjects in experiments. One of their main findings is that very few subjects are willing to sacrifice their own money in order to reduce inequality by lowering the payoffs of others. Half of the subjects are willing to reduce their own payoffs in order to make inequality-increasing sacrifices.

This does, however, run counter to the experimental results from public goods games with punishments as well as those from experiments conducted in a voting context (see Section 6.5 above). Charness and Rabin (2002) write (p. 5): “Indeed, while our analysis stresses that our data contradict difference-aversion models, we do not think we have conclusively disproved these models. Fehr and Schmidt [1999], for instance, have argued that only 40 percent of subjects need be difference-averse to explain the phenomena they explain. This is arguably consistent with our data.”

Charness and Rabin (2002) present two versions of a theoretical model. The models combine inequity aversion and the role of intentions, hence, they may be termed as *hybrid models*. The first version is a model of social preferences. Suppose that the set of individuals is $N = \{1, 2, \dots, n\}$,

and individuals are given respective allocations or consumption levels, $x_1, x_2, \dots, x_n$. The social welfare function is defined as

$$W = \delta \min \{x_1, x_2, \dots, x_n\} + (1 - \delta) (x_1 + x_2 + \dots + x_n); \delta \in (0, 1). \quad (6.40)$$

Social welfare is a convex combination of a Rawlsian maximin criterion, reflecting Rawlsian inequity aversion, (first term) and the sum of individual payoffs reflecting altruism and/or efficiency (second term). The utility of any individual $i \in N$ is defined as follows,

$$U_i = \lambda x_i + (1 - \lambda) W,$$

where $\lambda \in (0, 1)$ and W is defined in (6.40). When $n = 2$, denoting the two players by A, B , and $A \neq B$, the utility function for player A is given by:

$$U_A = \begin{cases} x_A + (1 - \lambda)(1 - \delta)x_B & \text{if } x_A < x_B \\ (1 - \delta + \lambda\delta)x_A + (1 - \lambda)x_B & \text{if } x_A \geq x_B \end{cases}. \quad (6.41)$$

For the preferences given in (6.41), the indifference curves in (x_A, x_B) space are kinked and downward sloping. To see this, we can use (6.41) to calculate the marginal rates of substitution (MRS) in each of the two outcome domains:

$$\frac{dx_B}{dx_A} = \begin{cases} MRS_{AB}^- = \frac{-1}{(1-\lambda)(1-\delta)} & \text{if } x_A < x_B \\ MRS_{AB}^+ = \frac{-(1-\delta+\lambda\delta)}{(1-\lambda)} & \text{if } x_A \geq x_B \end{cases}.$$

The absolute value of the slope is relatively greater below the 45° line ($|MRS_{AB}^-| > |MRS_{AB}^+|$). Hence, when $x_A < x_B$, the opportunity cost of a unit of x_i for player i , expressed in terms of x_j is relatively higher. Thus, individuals are more reluctant to accept a reduction in their payoffs when they are behind others as compared to the case where they are ahead of others.

This model is able to explain why proposers in dictator games give a strictly positive amount to responders and why trustees in trust games return a strictly positive amount to the investor. However, it is not able to explain why subjects in experiments retaliate and reduce the material payoffs of others at a cost to themselves as, for instance, in public goods games with punishments.

For this reason, Charness and Rabin propose a second version of their model, in which players are affected by outcome differences as well as the intentions of other players. Like models of interdependent preferences that we consider below, this model ignores the more rigorous framework of psychological game theory in which beliefs are modeled endogenously; see Part 4 of the book. By contrast, the Charness and Rabin (2002) and Levine (1998) models opt for a reduced-form framework that may be more tractable.

Suppose that there are two individuals, A and B , with respective material payoffs, x_A and x_B . Then, the preferences of player B are represented by the following utility function (the utility function of player A is analogously written by swapping the indices A and B).

$$U_B(x_A, x_B) = \begin{cases} (1 - \rho - \theta q)x_B + (\rho + \theta q)x_A & \text{if } x_A \leq x_B \\ (1 - \sigma - \theta q)x_B + (\sigma + \theta q)x_A & \text{if } x_A > x_B \end{cases}. \quad (6.42)$$

The payoff of any player is a weighted average of one's own payoff and the payoff of the other player. The weights depend on one's relative payoff with respect to the other player and also the intentions of the other player. In (6.42), $q = -1$ if player A has revealed negative intentions and $q = 0$ otherwise. The parameters of social preference are ρ, σ, θ . The first two, ρ, σ , relate to inequity aversion, while θ captures the role of intentions.

Suppose first that $\theta = 0$, so ignore the role of intentions for the time being.

1. An individual is said to have *competitive preferences* if $\sigma \leq \rho \leq 0$. In this case

$$U_B(x_A, x_B) = \begin{cases} x_B + |\rho|(x_B - x_A) & \text{if } x_A \leq x_B \\ x_B - |\sigma|(x_A - x_B) & \text{if } x_A > x_B \end{cases}.$$

Hence, player B derives utility from a higher material payoff and disutility from a lower payoff. Furthermore, since $|\sigma| > |\rho|$, player B exhibits a form of loss aversion; it is more painful to fall behind in material payoffs. Similar comments apply to player A .

2. The *difference aversion* (or *inequity aversion*) models of Bolton and Ockenfels (2000), Fehr and Schmidt (1999), and Loewenstein et al. (1989) can be represented by $\sigma < 0 < \rho < 1$. In this case, using (6.42) and continuing to assume that $\theta = 0$, we get

$$U_B(x_A, x_B) = \begin{cases} x_B - \rho(x_B - x_A) & \text{if } x_A \leq x_B \quad (\text{Altruism}) \\ x_B - |\sigma|(x_A - x_B) & \text{if } x_A > x_B \quad (\text{Envy}) \end{cases}.$$

Difference aversion and reciprocity can be combined, as in Falk and Fischbacher (2006), but without the apparatus of psychological games, by allowing $\theta \neq 0$. Suppose now that $\theta > 0$. Let $q = -1$, i.e., player B believes that player A has revealed negative intentions. Then, from (6.42), player B reduces the weight that she places on the payoff of player A , relative to her own payoff. This is quantitatively similar to a reduction in σ or ρ .

Charness and Rabin then define a *reciprocal fairness equilibrium* (RFE). The RFE is defined in the context of a richer model in which there are n players and there is a *demerit profile* $\rho = (\rho_1, \rho_2, \dots, \rho_n)$. The i th element of the demerit profile, $\rho_i \in [0, 1]$, indicates how much player i "deserves" in the view of the other players. The weight placed on player i 's payoff in the utility function of other players is decreasing in ρ_i . The demerit of any player is unanimously agreed upon by all other players by comparing the strategy of player i to some selfless standard of behavior. One can extend (6.40) to write the preferences of player i in this case as

$$U_i = (1 - \lambda)x_i + \lambda \left(\delta \min \left\{ x_i, \min_{j \neq i} (x_j + d\rho_j) \right\} \right) + \lambda(1 - \delta) \left(x_i + \sum_{j \neq i} \max \{1 - k\rho_j, 0\} x_j - f \sum_{j \neq i} \rho_j x_j \right). \quad (6.43)$$

where $d, k, f \geq 0$ are the model parameters. When $d = k = f = 0$, then (6.40), (6.43) are identical. Player i places less weight on the payoff of player j when d, k are large. When f is large, and ρ_j is large, then player i might derive utility from a reduction in the payoff of player j .

An RFE is a profile of strategies and a demerit profile such that (1) each player maximizes his payoff given in (6.43), and (2) the demerit profile is consistent with the profile of strategies. It is difficult to use the RFE concept except for simple games and the large number of parameters creates difficulties in testing the model. The demerit profile is defined only in equilibrium, which

makes it a difficult concept to use in the presence of multiple equilibria—a problem that is not uncommon in many models in economics.

Does the experimental data reveal that subjects have maximin preferences? The results of Engelmann and Strobel (2004), reviewed above in Section 6.7.1, reveal some support for such preferences but also persuasive evidence to the contrary. In the three player ultimatum game, the experimental evidence is that the passive receiver typically gets nothing (Riedl and Vyrastekova, 2002). This is clearly evidence against maximin preferences. Frechette et al. (2003) also find evidence against maximin preferences. In a group of five players, they ask one player to make a proposal to allocate a fixed sum of money among the five players. The support of at least three players (majority rule) is needed to implement the proposal. In 65% of the cases, the winning proposal made no allocation to two of the five players. Hence, all gains are mopped up by the winning coalition.

Okada and Riedl (2005) perform another experiment with a strategic element. There is one proposer and two responders. The proposer allocates a given amount to himself and the other two players; denote this allocation by (x_1, x_2, x_3) . If the proposer decides to form a three-person coalition, the total amount to be allocated is 3000. If he decides to form a two-person coalition, the total amount available is an element of the set $\{1200, 2100, 2500, 2800\}$. For the proposer's offer to be successful, both responders must accept the offer in the three-person coalition and the single responder must accept the offer in the two-person coalition. The prediction of maximin preferences is that the proposer chooses a three-player coalition and $x_1 = x_2 = x_3$. When the amounts available for the two-player coalition are high enough (2500, 2800), then 90% of the players choose the two-player coalition. The three-player coalition is chosen by most proposers only when the amount available to the two-player coalition shrinks sufficiently, i.e., to 1200.

These results suggest that in strategic situations, the maximin motive does not seem to be always salient. On the other hand, in dictator games, there is a strictly positive level of giving to the responder that is consistent with the presence of such a motive. It is conceivable that the maximin motive is probably more important in non-strategic settings such as charity, or in the context of anonymous voting for national referenda. The reader may also wonder if these results negate models of inequity aversion, such as the Fehr–Schmidt model. The answer depends on what is the appropriate reference group. From models of social identity (see Chapter 7), we know that individuals treat *ingroup* and *outgroup* members differently, and passive outsiders, in the experiments reported above, may be perceived as outgroup members. In this case, there is no violation of models of inequity aversion. Indeed, using these ideas one might better endogenize the issue of reference groups in models of inequity aversion.

6.8.2 *Models of interdependent preferences*

In models of *interdependent preferences*, actions taken by a player depend on the perceived type of the other player. For instance, altruistic players may wish to be kind to other altruistic players but not to other non-altruistic players. Let us consider the model of interdependent preferences in Levine (1998).³⁵

Suppose that there are n players and the game has an extensive form representation. At any end node, player $i = 1, 2, \dots, n$ has *direct utility* u_i . Each player has a type that is characterized by some altruism coefficient, $a_i \in (-1, 1)$. We say that the individual is spiteful if $a_i < 0$, self-regarding

³⁵ For the Ellingsen and Johannesson (2008) model of type-based reciprocity, see Section 8.5.2.

if $a_i = 0$, and altruistic if $a_i > 0$. A higher value of $a_i > 0$ denotes a more altruistic player. The altruism-adjusted or composite utility of player i is

$$v_i = u_i + \sum_{j \neq i} \beta_j u_j; \beta_j = \frac{a_i + \lambda a_j}{1 + \lambda}; \lambda \in [0, 1]. \quad (6.44)$$

In the special case $\lambda = 0$, we have $v_i = u_i + \sum_{j \neq i} a_i u_j$. So, depending on the sign of a_i , we have a model of *pure altruism* or *pure spite*, as in Ledyard (1995). When $a_i = 0$ and $\lambda = 0$ for all i , then we get the self-regarding case, $v_i = u_i$.

Players' concern for fairness is captured in the following manner.

1. If $\lambda > 0$ and $a_i > 0$, $a_j > 0$, then $\beta_j > 0$, so the payoff of altruistic players is increasing in the payoff of other altruistic players. If $\lambda > 0$, then β_j is increasing in a_j , i.e., the weight put on player j 's payoff is increasing in the altruism of player j .
2. When $\lambda > 0$ and $a_i > 0$, $a_j < 0$ such that $\frac{a_i}{\lambda} < |a_j|$ we have $\beta_j < 0$. In this case, the utility of an altruistic player i is decreasing in the spiteful player j 's payoff. So essentially, individuals want to be nice with altruistic individuals and spiteful with spiteful ones.

Population heterogeneity is introduced by assuming that a_i is drawn from some cumulative distribution function, $F(a)$, that is common knowledge to the players. However, a_i is privately known to player i only. Hence, the underlying game is one of private information about the types of others, where a type is the degree of altruism of a player. Players, through their actions, can reveal their private information about their altruistic coefficient, a_i . As in the Fehr-Schmidt model, the payoffs of other players are taken into account, but in addition, the types of the other players are also taken into account. An objective of the model is to explain the results of small stake experiments, so, typically $u_i(x_i) = x_i$.

The data from the ultimatum game is used to estimate the value of λ and to calibrate the distribution $F(a)$. Using these parameters, the model is able to explain why some individuals contribute in public goods games or why strictly positive offers are rejected in ultimatum games. However, it cannot explain why there is positive giving in the dictator game. Another potential issue, common with other signaling games, is that there might be a multiplicity of equilibria.

Rotemberg (2008) proposes another model of interdependent preferences to account for the results of the ultimatum and dictator games. One of the motivations of this paper is that the model by Levine (1998) is unable to explain the results of the modified dictator game of Charness and Rabin (2002); it predicts far too much spite relative to what is observed. In an ultimatum game, the responder's utility in Rotemberg's model is given by

$$U_R = x_R + (a_R - f(a_p^e, a_{\min})) x_P,$$

where x_R and x_P are the respective payoffs of the responder and the proposer, a_R is the responder's altruism, a_p^e is an estimate made by the responder of the proposer's altruism, $a_{\min}$ is some minimum level of altruism, and f is a function such that

$$f = \begin{cases} 0 & \text{if } a_p^e > a_{\min} \\ 1 + a_R & \text{if } a_p^e \leq a_{\min} \end{cases}.$$

Thus, if the perceived altruism of the proposer falls below a certain level ($a_p^e \leq a_{\min}$), then U_R is decreasing in the proposer's payoff. As a_p^e changes and $a_p^e - a_{\min}$ alters signs, we are likely to observe discontinuous jumps in the actions of the responder from acceptance to rejection; this is consistent with the experimental findings. The responder's decision to accept or reject makes the proposer's payoff risky. This is incorporated by specifying the proposer's utility function to be

$$U_P = E[x_P + a_R x_R]^\gamma,$$

where $\gamma > 0$ captures the degree of risk aversion of the proposer. In order to explain the data, it is assumed that the proposer is risk loving, i.e., $\gamma > 1$. Under the assumptions made above, the model explains well the data from the ultimatum game. However, it is not clear how well the model explains results from other experimental games on human sociality.

6.8.3 The model of Cox, Friedman, and Gjerstad (2007)

Cox, Friedman, and Gjerstad (2007) propose a flexible utility function with interdependent preferences that applies to sequential two-person games of perfect information. The utility function of player i is

$$U_i = \begin{cases} \frac{1}{\alpha} (x_i^\alpha + \beta x_j^\alpha) & \text{if } \alpha \neq 0 \\ (x_i x_j)^\beta & \text{if } \alpha = 0 \end{cases}, \quad (6.45)$$

where $\alpha \in (-\infty, 1]$ captures the ease of substituting x_j for x_i , and $\beta \neq 0$ is the *emotional state* of the player. In (x_i, x_j) space, the slope of an indifference curve for the preferences given in (6.45) is

$$MRS_{i,j} = \beta^{-1} \left(\frac{x_j}{x_i} \right)^{1-\alpha}, \quad (6.46)$$

where $\frac{x_j}{x_i}$ gives the relative payoffs of the two players. When $\alpha = 1$, $MRS_{i,j} = \beta^{-1}$, i.e., the allocations of the two individuals are substitutes. For $\alpha < 1$, one recovers the standard case of convex indifference curves. When $\alpha = 0$, $MRS_{i,j}$ is constant along any ray from the origin, so we get homothetic preferences. Leontief preferences correspond to $\alpha \rightarrow -\infty$.

Let $\pi_i^{\max}(s_j)$ be the maximum payoff that player i can get, conditional on the strategy, s_j , of player j and let $\bar{\pi}_i$ be some reference payoff of player i . The *reciprocity motive* of player j , denoted by ρ_j , is defined as $\rho_j = \pi_i^{\max}(s_j) - \bar{\pi}_i$.

The parameter β depends positively on the reciprocity motive, i.e., $\beta = \beta(\rho_j)$, which is an increasing function.³⁶ If $\rho_j < 0$, then player j is believed to have unkind intentions towards player i . In this case, $\beta(\rho_j) < 0$, so U_i is decreasing in x_j . Hence, player i would want to be unkind to player j .

The authors estimate the parameters of their model using data from ultimatum games and the Stackelberg duopoly game in Huck et al. (2001). The model fits well the data from these games. It remains to see if it also fits the data from other experimental games on human sociality.

³⁶ In the more general model $\beta = \beta(\rho_j, \phi)$, where ϕ is a variable that measures status. β is assumed to be increasing in both arguments.

6.8.4 A brief note on axiomatic foundations

We have considered three main kinds of theoretical models above. Models of social preferences (e.g., Fehr and Schmidt, 1999), models of interdependent preferences (e.g., Levine, 1998), and models of intentions-based reciprocity (e.g., Charness and Rabin, 2002).³⁷ The aim of these models is to propose psychologically plausible utility functions that can explain a range of experimental results. It is not clear, however, what restrictions on preferences are required for these utility functions.

Segal and Sobel (2007) make progress on this issue. We give a purely heuristic outline of the relevant conditions. They ask: what axioms need to be imposed in order to generate preferences for fairness and reciprocity? Segal and Sobel specify preferences directly over the profile of mixed strategies. Suppose that the generic mixed strategies of two players i, j in a two-player game are denoted by σ_i and σ_j , respectively. They define a preference relation for the i th player as $\succeq_{i\sigma_j}$ that is defined directly over the mixed strategies of player i , but conditional on the mixed strategy followed by player j , σ_j . More generally, the preference relation can depend on the entire mixed strategy profile of the players, which the authors call the *context* of the game.

Suppose that the following conditions hold: (1) Preferences over outcomes satisfy the axioms of expected utility. (2) Preferences over mixed strategies, $\succeq_{i\sigma_j}$, satisfy the continuity and the independence axioms. These two axioms are as stated under expected utility (see Part 1), but the objects over which they are stated are mixed strategies. (3) The self-interest axiom holds in the sense that if all other players are indifferent between two mixed strategies, then one chooses the mixed strategy that gives a higher own-payoff. Under these conditions, Segal and Sobel show that the preferences $\succeq_{i\sigma_j}$ can be represented by a utility function of the following form.

$$U_i = v_i(\sigma_i, \sigma_j) + \alpha_{i,\sigma_j} v_j(\sigma_i, \sigma_j), \quad (6.47)$$

where $v_i(\sigma_i, \sigma_j)$ is the utility of the i th player from material payoffs that arise from the strategy profile (σ_i, σ_j) . The parameter α_{i,σ_j} , which can be either positive or negative, depends on the strategy of the opponent. Although the authors do not use these terms, the sign of the parameter α_{i,σ_j} captures strategy-dependent altruism ($\alpha_{i,\sigma_j} > 0$), or strategy-dependent spite ($\alpha_{i,\sigma_j} < 0$). Several models that are discussed in Part 2, on altruism, spite, and inequity aversion, are special cases of (6.47), although in these models, altruism/spite are not dependent on the strategy of the opponent.

The following example demonstrates the implications of the preferences in (6.47).

	AM	PM
AM	$10(1 + i\alpha), 10$	$0, 0$
PM	$0, 0$	$10(1 + i\alpha), 10$
ALL	$7 + 10i\alpha, 10$	$7 + 10i\alpha, 10$

The column player is a plumber and the row player is a homeowner who has a leaky faucet. The plumber can either come in the morning (AM), or in the evening (PM). The homeowner needs

³⁷ For a formal consideration of models of intentions-based reciprocity, based on the more satisfactory framework of psychological game theory, see Part 4 of the book.

to take time off from work to facilitate the plumber's access to the faucet. If the homeowner takes time off in one session (either AM or PM), then it is not as damaging to his material payoff as taking the whole day off (the strategy, ALL). The plumber has purely self-regarding preferences, while the homeowner has social preferences. The homeowner places a weight $\alpha > 0$ on the payoff of the plumber. In this sense the homeowner is altruistic.

The joint payoffs are obviously maximized if the plumber and the homeowner can coordinate on the time (AM or PM); the parameter i , in the payoff matrix, captures feelings of reciprocity of the homeowner towards the plumber. Suppose that the plumber plays the two strategies, AM and PM, with respective probabilities $p, 1 - p$. If the plumber plays a pure strategy, i.e., $p = 0$ or $p = 1$, then the homeowner assigns "kind" intentions to the plumber because the latter is viewed as someone who likes to coordinate. In this case, $i = 1$, which implies positive reciprocity from the homeowner. However, if the plumber plays a purely mixed strategy, i.e., $p > 0$, then the plumber does not seem interested in coordinating and may cause the homeowner to take a full day off. The homeowner then assigns unkind intentions, $i = -1$, i.e., the homeowner places a negative weight on the payoffs of the plumber (negative reciprocity).

1. Suppose first that the plumber plays a pure strategy (so $p = 0$ or $p = 1$). If $p = 1$, then AM is strictly dominant for the homeowner, and if $p = 0$, then PM is strictly dominant for the homeowner. Hence, the two pure strategy Nash equilibria are (AM, AM) and (PM, PM).
2. Suppose that the homeowner has purely inequity averse preferences and does not care about intentions, so that $i = 1$ for all values of x . In this case, it is easy to check that (AM, AM) and (PM, PM) are still the Nash equilibria of the game.
3. Now suppose that intentions too matter, so $i = -1$ if the plumber plays a strictly mixed strategy, and $i = 1$ if he plays a pure strategy. We want to show that, in this case, the strategy ALL might be a part of the Nash equilibrium strategy profile. Define $g(\alpha) = \frac{0.7-\alpha}{1-\alpha}$. For ALL to dominate AM we require that $g(\alpha) > p$ and for ALL to dominate PM we require that $p > 1 - g(\alpha)$. Hence, the following is a mixed strategy Nash equilibrium. The homeowner plays ALL and the plumber plays AM with a probability p such that $g(\alpha) > p > 1 - g(\alpha)$. For instance, if $\alpha = 0.1$, then $g(0.1) = 2/3$. Hence, any p such that $\frac{2}{3} > p > 1/3$ is sufficient to induce the homeowner to play ALL in the Nash equilibrium.

Neilson (2006) provides axiomatic foundations for the Fehr and Schmidt (1999) preferences. The key axiom is *self-referent separability* that we consider heuristically. Suppose that we have three players. Consider ordered tuples of the allocations of the three players (x_1, x_2, x_3) , where x_i is the allocation of player $i = 1, 2, 3$. Suppose that the first player reveals the preference

$$(60, 60, 60) \succeq (60, 80, 40). \quad (6.48)$$

The payoff of the first player is common to both choices. The standard separability axiom requires that for any x_1 , it must be the case that $(x_1, 60, 60) \succeq (x_1, 80, 40)$. In particular, it must be the case that $(50, 60, 60) \succeq (50, 80, 40)$. But this has the implication that player 1 prefers an allocation in which he is ranked third in the income distribution rather than second. In other words, the separability axiom is not sensitive to one's rank in the income distribution.

The main thrust of a modified separability axiom, the *self-referent separability axiom*, is to take account of the rank of players making choices. Suppose that the set of feasible allocations of players belong to some set X . The following two conditions are required for this axiom to be satisfied in the context of this three person game.

1. Starting from any symmetric allocation, if a player prefers to add $y \in X$ to his allocation rather than $y' \in X$, then he should do so for any other symmetric allocation also. Formally, for any $x, y, z, y' \in X$,

$$\text{if } (x + y, x, x) \succeq (x + y', x, x), \text{ then } (z + y, z, z) \succeq (z + y', z, z).$$

2. Starting from any two symmetric allocations, if a player prefers adding some amount $y \in X$ to the outcomes of the other players for one allocation, he should also prefer to do so for any other amount $y' \in X$. Formally, for any $x, y, z, y' \in X$,

$$\text{if } (x, x + y, x + y) \succeq (z, z + y, z + y), \text{ then } (x, x + y', x + y') \succeq (z, z + y', z + y').$$

Proposition 6.21 (Neilson, 2006): Suppose that there are $n \geq 2$ players. Consider any allocation $x = (x_1, x_2, \dots, x_n)$ of payoffs to the n players. The preference ordering satisfies self-referent separability if and only if there exist functions $u_1, \dots, u_n$ such that preferences can be represented by a function U of the form

$$U(x) = u_1(x_1) + \sum_{i=2}^n u_i(x_i - x_1). \quad (6.49)$$

Several inequity averse utility functions, for instance, the Fehr–Schmidt (1999) utility function, are of the same form as (6.49). In (6.49), one's own consumption bundle is the reference point against which the consumption bundles of others are compared. As the reader may have guessed, one can also get a representation for prospect theory type preferences using similar axioms; see Neilson (2006).

CHAPTER 7

Human Morality and Social Identity

7.1 Introduction

In neoclassical economics, there is no explicit role for *character virtues*, such as the desire to be honest, to keep promises, to avoid deception and lies, and to be fair to one's fellow human beings. Neoclassical economics assumes that people are disposed to tell *selfish black lies* (Erat and Gneezy, 2012), i.e., lies that benefit the individual but also typically harm others. By contrast, we do observe individuals whose behavior accords well with “good” character virtues. The explanation for such behavior in neoclassical economics is *instrumental* because individuals are *consequentialists*. In this view, individuals strategically exhibit the relevant character virtues, based on a cost–benefit analysis. Neoclassical economics, thus, takes a fairly dim view of human nature, relative to many of the other social sciences in which humans *intrinsically* display character virtues.

The purpose of the first part of this chapter (Section 7.2) is to give experimental evidence on the source of human character virtues—instrumental, or intrinsic. Furthermore, we wish to apply these concepts to a theoretical model of promises and threats. We begin, in Section 7.2.1, with some of the relevant evidence. The evidence shows that character virtues may directly provide utility to humans, say, on account of enhanced self-esteem. Such character virtues are likely to dominate behavior, particularly when the gains from non-virtuous behavior are small. Humans are not likely to lie for small gains, but as the gains become large, the morality of a fraction of the subjects is loosened (Gneezy, 2005; Gneezy et al., 2013; Gibson et al., 2013). Character virtues may also depend on the prevalence of such character virtues in the population. For instance, bribing a traffic policeman, or evading taxes may be considered less non-virtuous in societies where such activity is already prevalent.

In an interesting experiment, Fischbacher and Föllmi-Heusi (2013) asked subjects to privately roll a six-sided dice, observe and then self-report the outcome. The distribution of self-reported outcomes can now be compared with the predicted statistical distribution where each of the six outcomes, 1 to 6, occur with probability 1/6. Roughly 40% of the subjects are *completely honest*, even when lying could have improved their payoffs. A significant fraction of the subjects tell *partial lies*, even when they could have lied more strongly (see also Gneezy et al., 2013). About 22% behave in a manner that is consistent with the selfish black liars assumed in neoclassical economics. We also report the results of a similar elicitation by Pruckner and Sausgruber (2013), who examine non-conclusive evidence on gender differences in lying, as well as the importance of many non-economic factors in the lying decision.

In neoclassical economic theory, *promises* and *threats* are simply cheap talk. By contrast, a significant fraction of subjects in experiments honor their promises (Ellingsen and Johannesson, 2004; Charness and Dufwenberg, 2006; and Vanberg, 2008). We consider some of the relevant evidence in Section 7.2.2. We also consider here, a theoretical model of promises and threats in a buyer–seller problem in which there is the possibility of the hold up of the seller’s investment (Ellingsen and Johannesson, 2004). When the buyer and the seller have other-regarding preferences, and there is a cost of breaking a promise, or not carrying out a threat, two things happen. First, promises and threats become effective, and alter the eventual outcome; this is confirmed by the experimental results. Second, the effect of other-regarding preferences can magnify the cost of breaking a promise and the cost of not carrying out a threat.

One may also carry out promises and threats for reasons other than the costs of lying. For instance, promises may induce others to form certain beliefs. If one believes this to be the case, then one may suffer from *guilt* by breaking the original promise (Charness and Dufwenberg, 2006). Issues of guilt and the relevant empirical evidence is considered in Part 4 of the book. However, the results of Vanberg (2008) support the *cost of lying* explanation of breaking promises.

Section 7.3 considers the topic of *social identity*, which traditionally is a part of the subject matter of social psychology. Our focus, however, is on exploring the recent advances from the rapidly expanding literature in economics. In *social identity theory*, one’s *social identity* refers to the *social category* one belongs to, such as black or white, blue collar worker or white collar worker, and soldier or civilian. Different identities may be triggered by different social contexts. For instance, a national identity may apply to a football supporter when his country plays against another, and a club identity when his preferred club plays against another. Once the social identity of a group is taken on, group members prefer to conform to the group norm.

In most classic experiments on social identity, individuals are *primed* for a group identity. There is evidence that even a *minimal group identity*, based on even trivial characterizations, say, blue and red groups, may alter the behavior of group members. The typical finding is that one favors one’s own group members, or the *ingroup*, relative to non-group members, or the *outgroup* (Billig and Tajfel, 1973; Tajfel and Turner, 1986).

In Section 7.3.1, we consider the evidence from recent experiments in economics on social identity. The Asian American identity is associated with the norms of hard work and patience. When Asian Americans are primed for their identity, they make more patient choices than whites. By contrast, in the unprimed condition, their choices are more impatient than whites (Benjamin et al., 2010). Altruism and envy are found to be stronger towards ingroup members and players are more likely to forgive unfavorable behavior by an ingroup member relative to an outgroup member (Chen and Li, 2009).

Charness et al. (2007) play a two-player battle of sexes (BOS) game and a two-player prisoner’s dilemma (PD) game. They find that subjects make more aggressive choices (e.g., defect in the PD game) when group membership is highly salient (e.g., choices are made in front of other group members who observe the outcomes immediately). When a player observes that the opponent is making choices under highly salient group membership, he responds by making more accommodating choices (e.g., cooperate in the PD game). When group membership is less salient, the percentage of aggressive choices drops significantly.

Goette et al. (2012) argue that real-life social groups (SG) exhibit far more social ties and friendships between group members relative to the minimal groups (MG) in experiments. Based on field studies on Swiss army trainees, they find that there are qualitative and quantitative differences between the behavior of SG and MG group members that can be used to reconcile existing results. There is significantly greater cooperation between SG members and they punish defecting

ingroup members as much as defecting outgroup members. However, while MG members punish defecting outgroup members as much as SG members do, they punish ingroup members far less.

Since ingroup members engage in unfavorable behavior towards outgroup members, one might wonder if the existence of groups reduces overall welfare. Hargreaves Heap and Zizzo (2009) argue and empirically show that group membership confers private hedonic benefits. Taking account of these hedonic benefits produces an overall positive effect of group formation.

Section 7.3.2 considers a simple principal–agent model of the interaction between identity and incentives (Akerlof and Kranton, 2000, 2005, 2008). The principal has the choice of appointing a supervisor to monitor the non-contractible effort level of the agent. However, the very act of appointing the supervisor alters the identity of the agent, who now considers himself an outsider. When the principal has the option of appointing a supervisor, yet does not appoint one, then the agent feels trusted, and dons the identity of an insider.

Once the agent considers himself as an insider he internalizes a *higher effort norm*. Thus, choosing not to appoint a supervisor has a cost and a benefit for the principal. The cost is the inability to engage in ex-post punishments, such as fines, if the agent shirks; this reduces deterrence. The benefit is that since the agent internalizes a higher effort norm, the individual rationality constraint of the agent can be satisfied at a lower expected wage cost. These sorts of trade-offs then determine the optimal organizational design.

7.2 Human virtues

We separate the discussion in this section into two parts. In the first part, we discuss the empirical evidence on *lying behavior*. In the second part, we consider formal models of *promises* and *threats*.

7.2.1 Evidence on human virtues

While the interest in human attitudes to lying in economics is a relatively recent and growing phenomenon, psychologists have been interested in these issues for much longer (Vrij, 2000, 2008). In neoclassical economics, one takes a fairly dim view of human nature. Using the terminology of Erat and Gneezy (2012), neoclassical economics assumes that humans feel no remorse in telling *selfish black lies*. Typically, such lies benefit the person telling the lie but harm the other party. In particular, an individual telling selfish black lies suffers no decrease in utility on account of the lies. Neoclassical economics then proceeds to determine incentive compatible mechanisms that induce truth-telling, purely in response to the underlying economic incentives. The dismal nature of this view of human nature is not a reason to reject it. Our interest, below, is in empirically pitting this view against alternative views of human nature.

In contrast to selfish black lies, *white lies* do not harm the other party, and may indeed contribute to increasing the utility of others. Erat and Gneezy (2012) give two insightful examples. Consider a supervisor giving feedback to a failing employee. If the feedback is entirely truthful, it may seriously undermine the employee's confidence in undertaking future tasks. Shading the truth may be efficacious in meeting longer-term objectives. In other contexts, doctors may sometimes give placebos to patients when they fully know that the placebo has no pharmacological effect.

Two further categories of white lies have also been used in empirical work. A player who tells *altruistic white lies* may be harmed by the lies, but the intention is to increase the payoff of the other party. A player tells *Pareto white lies* when such lies improve the payoffs of both players.

Table 7.1 Payoffs of the players in each treatment in Gneezy (2005).

Treatment	Option	Payoff of	
		Player 1	Player 2
1	A	5	6
	B	6	5
2	A	5	15
	B	6	5
3	A	5	15
	B	15	5

It is also useful to consider the following terminology in Gneezy (2005). Suppose that action A involves a lie but action B does not, yet both actions lead to the same monetary payoff for a player. If the player is indifferent between the actions A and B, we say that the player takes a *consequentialistic approach*. Players in neoclassical economics are consequentialists. Otherwise, if there is a preference for one action over the other, the player takes a *non-consequentialistic approach*. The bulk of the empirical evidence indicates that the actions of players are non-consequentialistic.

In cheap talk games, unlike signaling games, the messages sent by players do not affect the payoffs of players. Rather, the messages influence the beliefs of the other player about the state of the world, hence, they indirectly influence payoffs through their effect on actions. Gneezy (2005) considers cheap talk games between two players; the three treatments in his experiments are shown in Table 7.1. In each treatment, player 1 (the sender) has two possible choices: A and B. Player 2 (the receiver) does not see the payoffs, however, player 1 can see the payoffs.

Player 1 sends one of two possible messages to player 2. Message A says: “Option A will earn you more money than option B.” Message B says: “Option B will earn you more money than option A.” Message B is a *lie* because for all three treatments, player 2 gets a lower payoff from option B. The sender always gets a higher payoff from option B. Player 2 observes the message and then chooses either option A or B. Player 2’s chosen option is then implemented. The structure of the game is not rich enough for a signaling equilibrium to be sustained; player 2 does not even know if the incentives of the two players are aligned.

Out of the 50 participants who were assigned the role of a sender, in an incentive compatible elicitation, 82% expected the receiver to follow their message. This was close to the actual percentage of receivers who did follow the messages, 78%. Thus, conditional on their messages being largely followed, if senders are assumed to tell selfish black lies, then they should always choose the option that maximizes their payoff (message B in all treatments).

Notice that the incentive to lie for player 1 (choosing message B) is different in each of the three treatments on two different counts. First, the gain in own-payoff is 1 for treatments 1 and 2 but it is 10 for treatment 3. Second, the loss to player 2 from being lied to is 1 in treatment 1 but 10 in treatments 2 and 3. The percentage of player 1s who lied in treatments 1, 2, and 3, respectively, is 36, 17, and 52 and these percentages are statistically different. In treatments 1 and 2, the personal gain from lying is 1, however, the loss to player 2 increases from 1 to 10. This reduces the fraction of player 1s who lie by more than half. However, in treatments 2 and 3, the loss to player 2 is 10 in each case but the personal gain to player 1 increases from 1 to 10. This significantly increases the

percentage of player 1s who lie to 52. This experiment shows that the extent of lying responds to the gain from lying and the cost to the other player. However, it cannot cleanly separate the *role of other-regarding preferences from aversion to lying*.

In a separate questionnaire, Gneezy (2005) asks a different set of players to judge the fairness of lies with different consequences. University of Chicago students are asked which of the following two lies is more unfair. In one case, the seller sells a car to a buyer without revealing that it will cost \$250 to fix a pump in the car. In another, the seller does not reveal that the brakes are faulty, which will also cost \$250 for the gullible buyer to fix. A significantly higher percentage of the respondents judged the second scenario (faulty brakes) to be the more unfair of the two lies, despite identical monetary costs of repair. Subjects are not merely consequentialists but they apparently care about the higher potential risk to the buyer in the second case. If the responses of these subjects are more generally representative of the population, then we can trust a car salesman to more accurately report the condition of the brakes as compared to the air-conditioner.

Hurkens and Kartik (2009) reconsider the results of Gneezy (2005) and show that they are unable to reject the hypothesis that the data comes from a population of players of two types. Type-I never lies whatever the cost (ethical types) and Type-II always lies (unethical type).¹ However, if the types are fixed, this evidence begs the question of why the extent of lying responds to the incentives to lie in so many diverse contexts.

Gibson et al. (2013) conduct a decision theoretic lab experiment in which the confounds of strategic considerations and other-regarding preferences are eliminated. In the truth-telling task, subjects, in their role as the CEO of a company, had to decide on two possible earnings announcements that affected their own payoffs. A higher announcement is a lie but also increases their payoff. In several different treatments, the payoff differences between truth-telling and lying are different. For instance, in one treatment, the choice was between announcing 31 cents/share and 35 cents/share and the higher announcement led to 5 times higher earnings for the CEO. The subjects are told that the lie is within the bounds of what could be defended on accounting grounds.

If the subjects were of only two types, ethical and unethical, then we should not get any variation in the levels of truthfulness as the incentive to lie varies. The ethical types should never lie and the unethical types should always lie. In contrast, the extent of lying was sensitive to the incentive to lie. The authors find that the most significant factor in the decision to lie is the intrinsic cost of lying, which ties in with the results of Gneezy (2005). Hence, people seem to have underlying preferences over how much they are willing to lie as incentives for lying vary.

Fischbacher and Föllmi-Heusi (2013) introduce a novel method of measuring the extent to which people lie, even when the individual responses cannot be checked for lying. Subjects in the experiment are asked to roll a dice and privately observe the outcome of the dice. To ensure anonymity, the experimenter is not present in the room when each participant privately rolls the die. For the numbers 1, 2, 3, 4, and 5, the payoff equals the self-reported number that subjects claim has appeared on the dice. If the claimed number is 6, then payoff is zero. The experimental design ensures that subjects can never be found out if they cheat. The distribution based on the claimed numbers is then compared to the predicted statistical distribution of outcomes in which each number comes up with a probability 1/6 or 16.7%. Several control treatments were used that varied the level of stakes, created a negative externality on another player from lying, and created even higher levels of anonymity.²

¹ For an early theoretical attempt at modeling two-type models of this sort, see Koford and Penno (1992).

² In this treatment, participants shredded any pieces of paper they had written on and directly took prize money from an envelope without the intervention of the experimenter.




Figure 7.1 Results for the baseline treatment.

Source: "Lies in disguise: an experimental study on cheating," Fischbacher, U. and Föllmi-Heusi, F. *Journal of the European Economic Association*, 11(3): 525–47. Copyright © 2013 by the European Economic Association.

The results of the baseline treatment for 389 participants are shown in Figure 7.1. On the horizontal axis are the payoffs of the player that simply correspond to the claimed number on the dice (recall that for a report of 6 on the dice, the payoff is zero). The vertical axis measures the percentage who report being in each category; the theoretical prediction of 16.7% is also shown. The stars denote the significance levels for a two-sided binomial test of the differences in observed percentage of players who self-report in each category from the theoretical prediction of 16.7%. The results are as follows.

1. A Kolmogorov–Smirnov test rejects the hypothesis that the observed distribution is identical to the predicted uniform distribution where each number occurs with a chance of 16.7%.
2. Numbers below 4 and the number 6 are underreported as they lie below the 16.7% line, while the numbers 4 and 5 are overreported. Since only 16.7% of the subjects are predicted to get a 6 and anybody who does not get a 6 is unlikely to report a 6, it is significant that 6.4% nevertheless report 6. This is contrary to the assumption of selfish black liars in neoclassical economics. Assuming that the proportion of intrinsically honest players is independent of the number arising on the throw of the dice, then, since $\frac{6.4}{16.7} = 0.383\ 23$, the proportion of intrinsically honest subjects is 38.3%.
3. The number 5 is payoff maximizing, yet 16.7% of the subjects would have got a 5 anyway, hence, the percentage of people who are unethical income maximizers is $\frac{6}{5} (35 - 16.7) = 21.96\%$.
4. More than 16.7% of the subjects report the number 4. Hence, there is a set of subjects who lie but do not lie maximally. These may be termed as partial liars. Such behavior could arise, for instance, if subjects wish to maintain a positive self-image.

A similar elicitation method is used by Prucker and Sausgruber (2013) using field data. They conducted their field experiments in two towns in Austria. Several tabloids operate under an *honor system* in which newspapers are sold on the street using a booth filled with newspapers. Customers pay into a padlocked cashbox. The payments are not monitored, so customers can pay an amount below the price of the newspaper or not pay at all. There is very little gain from

not paying and a very small probability of non-payments being discovered. The authors run two different treatments by posting a note on the booth.

In the treatment LEGAL, a note is posted on the booth, reminding customers of the legal norm of paying for the newspaper. In treatment MORAL, the posted note appeals to the honesty of the customers, so it primes them to follow a social norm. The newspaper costs 0.60 euros and only a third of the customers pay a strictly positive amount. Average payment in the LEGAL treatment is 0.061 euros and average payment in the MORAL treatment is 0.14 euros. Thus, the MORAL treatment is more effective. The authors conjecture that there is, in fact, a moral norm of payment and when subjects are primed to consider it, self-interest is reduced somewhat in favor of a social norm of honesty.

Abeler et al. (2014) conducted a field study in Germany. Participants were contacted by phone, asked to toss a coin that only they privately observe, and report the result on the phone. Subjects received a payoff of 15 euros if they report tails. Under self-regarding preferences and amoral preferences, all subjects should report tails. Using the method introduced by Fischbacher and Föllmi-Heusi, the reported distribution of tails is compared to the predicted theoretical distribution. Almost all subjects told the truth; indeed the percentage reporting tails is lower than 50%. The authors conclude that the results are consistent with a cost of lying and that lab experiments understate the costs of lying and overstate the extent of amoral preferences.

Mazar et al. (2008) used two different treatments. In the first treatment, the experimenter checked answers to 20 maths questions and students were paid according to the number of correct answers. This establishes a benchmark distribution of marks. In the second treatment, subjects could mark the questions themselves (based on an answer key provided by the experimenter), then shred their answers and self-report the number of correct answers to claim the payment. Students in treatment 2 reported solving 10% more questions as compared to treatment 1. If students did not feel any unease at telling selfish black lies, then all of them should have reported 100% correct answers, given that this would have maximized their payoffs, particularly when the probability of being caught was zero.

Gneezy et al. (2013) introduce a third method of measuring lying aversion. Two individuals A and B are paired to play a simple sequential sender–receiver game. Only player A observes a randomly determined state of the world $s \in S = \{1, 2, \dots, 6\}$; each of the states is equally likely. Player A then sends a message $r \in S$ to player B. The message is truthful if $r = s$, otherwise it is a lie. Player B observes the message and then chooses an action $a \in S$. The payoff of player A, π_A , is linearly increasing in r but it does not depend on the state, s , or the action, $a \in S$, of player B. The payoff of player B, π_B , depends on his action a and the truthfulness of A. The exact payoffs are as follows.

$$\pi_A = 10 + 2r \text{ and } \pi_B = \begin{cases} 10 & \text{if } r = a \text{ and } r = s \\ 0 & \text{if } r = a \text{ and } r \neq s \\ 3 & \text{if } r \neq a \end{cases}.$$

Thus, player B gets the highest payoff when A sends a truthful message that is believed by player B. If player B chooses not to believe the message of player A (truthful or not), then he gets a fixed payoff of 3. If untruthful messages are believed, then B's payoff is zero. A player who is disposed to tell selfish black lies, as in neoclassical economics, would choose the message $r = 6$ in every state, s . The incremental payoff to player A from lying increases as s decreases.

The results with 72 subjects, who were rematched randomly over 24 rounds are as follows. The messages sent by player A in each state are shown in Figure 7.2. Each column corresponds to the




Figure 7.2 Messages sent by player A in each state.

Source: Reprinted from *Journal of Economic Behavior & Organization*, 93: 293–300. Gneezy, U., Rockenbach, B., and Serra-Garcia, A., “Measuring lying aversion.” Copyright © 2013, with permission from Elsevier.

percentage distribution of messages, $r \in S$, sent by player A for each state, $s \in S$. The sizes of the circles reflect the fraction of choices in each category. Along the diagonal, we get the percentage of truthful messages sent by A. Since the incentive to lie increases as s falls, it is not surprising that as s increases, messages are increasingly truthful.

Another indicator of this is the top row in Figure 7.2, which shows a message of $r = 6$ for each $s \in S$. As s increases from 1 to 5, the percentage of player As who lie maximally ($r = 6$) falls; these differences are statistically significant. These results are consistent with those of Gneezy (2005) reviewed above. As in Fischbacher and Föllmi-Heusi (2013), many players are partial liars but the percentages here are smaller; for instance, for states $s = 1, 2$, and 3 , there are 10% partial liars. There is also an element of other-regarding preferences in the results that is revealed over successive rounds of the experiment. Player As are less likely to lie in period t if their lie in period $t - 1$ reduced B's payoff.

As in many of the other results reported in this section, a very large fraction of player Bs choose to follow player A's message (i.e., $a = r$). For $r = 1, 2, \dots, 5$, in 85–89% of the cases we have $a = r$ but when $r = 6$ we have $a = r$ for only 46% of player Bs. There is evidence of learning over successive rounds; player Bs reduced their dependence on messages by A if they were lied to in the previous period. On the whole, excluding $r = 6$, player Bs significantly improved their payoffs from following the message of player A.

A developing theoretical literature introduces a cost of lying in the players' utility functions.³ Serra-Garcia et al. (2013) consider the efficacy of cheap talk in asymmetric information, sender-receiver, games. They consider two kinds of messages that can be sent by the sender—a message

³ See, for instance, Ellingsen and Johannesson (2004), Kartik et al. (2007), and Kartik (2009).

about the private information of the sender or a message about the action of the sender. In these kinds of games, there are two equilibria. In the *babbling equilibrium*, the receiver ignores the message sent by the sender. However, in the *influential equilibrium*, there is a separating equilibrium in which different types of senders send different messages and the receiver conditions his actions on the message of the sender. Under the *cost of lying assumption*, the only equilibrium that survives is the influential equilibrium.

There is a growing literature that examines the differential propensities of men and women to lie. For instance, Dreber and Johannesson (2008) find that men are more likely than women to tell a selfish black lie. Childs (2012) finds no gender differences in lying. However, when the stakes are reduced, in a similar game to Dreber and Johannesson, no gender differences are found by Gylfason et al. (2013). Erat and Gneezy (2012) show that men are more likely to tell a selfish black lie and a Pareto white lie. However, women are more likely to tell an altruistic white lie. This contributes to a developing literature that requires further replication and testing; for a survey on gender differences in preferences, see Croson and Gneezy (2009).

Conrads et al. (2013) find that individuals are more prone to lying under team incentives relative to piece rates. The reason is that under team incentives, players may be able to spread the consequences of lying among other team members. Cappelen et al. (2013) consider the non-economic correlates of lying. They find that the non-economic factors are as important as the economic ones. Participants who gave larger amounts in dictator games are also less likely to lie in sender–receiver games. They do not find a gender effect of lying, although females are slightly more prone to lie than males, a finding that is in contrast to earlier work. Other variables that correlated with lying were age, a measure of other-regarding preferences, and the nature of the lie, e.g., the incidence of lying falls when the message of the sender involves personal content.

7.2.2 A theoretical model of promises and threats

In neoclassical economics, promises or threats are not credible unless instrumentally rational self-regarding players respond to reputational concerns, or to third party enforcement. In the absence of these conditions, verbal commitments, in the form of *promises* or *threats*, are predicted to be ineffective. This is in contrast with the view taken in psychology and sociology, where such verbal communication can be efficacious; Part 4 of the book considers experimental evidence that supports this view. Ellingsen and Johannesson (2004), henceforth EJ, examine the predictions from these conflicting approaches in a classical hold up problem.

Consider a multi-stage, two-player, game between a seller (*S*) and a buyer (*B*). In stage I, the seller can, at a fixed cost, $F > 0$, make an investment decision $I \in \{0, 1\}$. $I = 0$ denotes foregoing the investment opportunity, and $I = 1$ denotes making the investment. The gains from trade in stage II are $g \in \{0, G\}$. If $I = 1$, then $g = G > F$, otherwise if $I = 0$, then $g = 0$.

If $I = 1$, then, in stage II, let x_S, x_B be, respectively, the seller's and buyer's shares of the gains from trade, G . Following, $I = 1$, the buyer proposes a split of G by making a proposal $x = \{x_S, x_B\}$ such that $x_S + x_B = G$. The seller makes the binary accept/reject decision $a(x) \in \{0, 1\}$. If the decision is to accept ($a = 1$), the proposal x is implemented, and if it is to reject ($a = 0$), both get zero. The outside option of the seller is zero. In the experiments, the unit of currency is Swedish Kroner (SEK), $F = 60$ and $G = 100$.

Let us first consider the theoretical predictions under the benchmark case of self-regarding preferences.

Proposition 7.22 *Suppose that the seller and the buyer have self-regarding preferences. Then, in the unique subgame perfect equilibrium of the game: $I = 0$, $x_B = G$, $x_S = 0$, and $a = 1$. Furthermore, any promises made by the buyer or the seller at any stage do not influence the outcome.*

Proof: In stage II, the seller will accept any non-negative offer because the outside option is zero, so $a = 1$. Knowing this, the buyer offers $x_S = 0$, so $x_B = G$. In stage I, the ex-ante utility of the seller from $I = 1$ is $-F < 0$ and from $I = 0$ it is 0. Thus, $I = 0$ is optimal. The reason that promises have no effect in the game is that players have self-regarding preferences and there is no cost of breaking a promise. ■

Let us now examine the experimental results to see if the predictions of the self-regarding model in Proposition 7.22 hold. EJ consider three treatments. In the first treatment, no communication is allowed. In the second treatment, the buyer sends a message to the seller before the seller makes the investment decision. In the third treatment, the seller sends a message to the buyer along with his investment decision $I \in \{0, 1\}$. Clearly, messages in the second treatment are of the nature of promises. For instance, the buyer might *promise* in his message to the seller that “I shall share the net gains from trade, $G - F$, equally with you.” Messages in the third treatment are of the nature of *threats*. For instance, in his message to the buyer, the seller might issue the threat that “If you make me a low offer, I will reject your offer and each of us will get zero.”

The experimental results reject the predictions of Proposition 7.22, and are as follows. In treatment 1, a third of the sellers choose $I = 1$ and the modal offer by buyers is to split the net gains from trade $G - F = 40$ equally, thus, the modal proposal is $x_S = 80$, $x_B = 20$. However, the mean offer in treatment 1 is $x_S = 48.57$ and mean profit of the seller, $x_S - F = -12.86 < 0$.

The presence of communication in treatments 2 and 3 increases the mean offer. The mean offer in treatment 2 (buyer communication) is the highest and it is significantly higher than in treatment 1; the mean profit of the seller in treatment 2 is 8.75. The two most frequent proposals made by the buyer in treatment 2 were $x_S = 60$ (so seller’s net profit, $x_S - F$, is zero), and $x_S = 80$ in which case the seller’s net profit is 20. Out of the 30 pairs of subjects, in three cases, sellers accepted an offer that made them a loss.

The mean offer in treatment 3 is lower but not statistically different from that under treatment 2. Nine out of the 21 messages in treatment 3 contained threats; however, these threats were not very credible and often ignored by the buyers. The vast majority of the seller messages (16 out of 21) suggested splitting the net gains from trade equally but in the majority of cases (9 out of 16), the buyers made lower offers to the seller. This is in keeping with the nature of the holdup problem; once the seller sinks the investment, the buyer can holdup the seller, ex-post.

The failure of the predictions of the model with self-regarding preferences is problematic because the hold up problem is of fundamental importance in industrial organization and in contract theory. EJ suggest that the explanation lies in other-regarding preferences, modified to take account of a cost of breaking the promise.

Suppose that the seller and the buyer in the EJ model have Fehr–Schmidt preferences as in (7.1), so that the Fehr–Schmidt utility of the seller is denoted by $U_S(x, I, a; \alpha_S, \beta_S)$, where α_S, β_S are the preference parameters of the seller. The buyer’s Fehr–Schmidt utility is denoted by $U_B(x; \alpha_B, \beta_B)$, where α_B, β_B are the preference parameters of the buyer and $x = \{x_S, x_B\}$. Suppose that there is mutual knowledge of the utility functions, U_S and U_B .

Proposition 7.23

- (a) *If $\beta_B > \frac{1}{2}$, then in the unique subgame perfect equilibrium of the model: $I = 1$, $x_B = \frac{G-F}{2}$, $x_S = \frac{G+F}{2}$ and $a = 1$.*

(b) If $\beta_B < \frac{1}{2}$, then in the unique subgame perfect equilibrium of the model, $I = 0$.

Proof: Consider a seller who, in stage II, is making his acceptance decision, having observed the buyer's proposal $x = \{x_S, x_B\}$, and conditional on having chosen $I = 1$ in stage I. The stage II payoffs of the seller and the buyer are, respectively, $x_S - F$ and x_B , where $x_S + x_B = G$. Suppose that the buyer does not make any overfair offers, so the buyer's payoff is relatively greater (we appeal to the ultimatum game for this sort of evidence⁴). Denote the utility of the seller from accepting ($a = 1$) and rejecting ($a = 0$) the proposal, respectively, by $U_S(x, 1, 1; \alpha_S, \beta_S)$ and $U_S(x, 1, 0; \alpha_S, \beta_S)$. Thus,

$$U_S(x, 1, 1; \alpha_S, \beta_S) = x_S - F - \alpha_S [x_B - (x_S - F)] \text{ and } U_S(x, 1, 0; \alpha_S, \beta_S) = -F - \alpha_S F. \quad (7.1)$$

Substituting $x_B = G - x_S$ in (7.1), the smallest offer x_S that the seller accepts such that $U_S(x, 1, 1; \alpha_S, \beta_S) > U_S(x, 1, 0; \alpha_S, \beta_S)$ is

$$\underline{x}_S = \frac{\alpha_S G}{1 + 2\alpha_S} > 0. \quad (7.2)$$

A rational buyer with mutual knowledge of Fehr-Schmidt preferences knows that any offer to the seller below $\underline{x}_S$ will be rejected (recall that under self-regarding preferences, buyers offered $x_S = 0$). Thus, the maximum monetary payoff of the buyer is $G - \underline{x}_S$.

In making his proposal, $x = \{x_S, x_B\}$, the buyer takes the subsequent behavior of the seller, in (7.2), as given. Suppose that we are interested in the case where the buyer demands more than half the net surplus, i.e., $x_B \geq \frac{G-F}{2}$, so that the buyer's payoff is relatively greater. Since $x_S = G - x_B$, the condition $x_B \geq \frac{G-F}{2} \Leftrightarrow x_S \leq \frac{G+F}{2}$. Also notice that $\frac{G+F}{2} > \underline{x}_S$. Thus, the buyer chooses x_B in order to maximize $U_B(x; \alpha_B, \beta_B) = x_B - \beta_B [x_B - (x_S - F)]$ subject to the constraints $x_S + x_B = G$ and $0 < \underline{x}_S \leq x_S \leq \frac{G+F}{2}$. This is equivalent to the problem of choosing x_B in order to maximize

$$U_B(x; \alpha_B, \beta_B) = x_B - \beta_B [x_B - (G - x_B - F)], \quad (7.3)$$

subject to $\frac{G-F}{2} \leq x_B \leq G - \underline{x}_S$. From (7.3) we have

$$\frac{dU_B}{dx_B} = 1 - 2\beta_B.$$

Suppose that $1 - 2\beta_B < 0 \Leftrightarrow \beta_B > \frac{1}{2}$, then the buyer chooses the lower limit $x_B = \frac{G-F}{2}$, so $x_S = G - x_B = \frac{G+F}{2}$. In this case, it is optimal in the first stage for the seller to choose $I = 1$.

When, $\beta_B < \frac{1}{2}$, $\frac{dU_B}{dx_B} > 0$, so the buyer chooses the upper limit $x_B = G - \underline{x}_S = G \left(\frac{1 + \alpha_S}{1 + 2\alpha_S} \right)$. The ex-ante utility of the seller from investing in this case is $U_S(\underline{x}_S, 1, 1; \alpha_S, \beta_S) = \underline{x}_S - F - \alpha_S [G - \underline{x}_S - (\underline{x}_S - F)] = -F - \alpha_S F < 0$, which is lower than the outside option of zero. Hence, $I = 0$. ■

The inequity aversion parameters of both players, α_S and β_B , determine the size of the offer and the investment decision. If the advantageous inequity parameter of the buyer, β_B , is high enough ($\beta_B > 1/2$), then the buyer offers to share equally in the gains from trade, inducing the seller to

⁴ This assumption comes with the caveat that in some cultures overfair offers are not unheard of; see Henrich et al. (2001).

make the investment in stage I. If $\beta_B < 1/2$, then the ex-post hold-up problem is serious enough for the seller to prevent any investment in stage I. From (7.2), we have that the seller's minimum acceptable share is increasing in the gains from trade, G , as well as the disadvantageous inequity parameter of the seller, α_S . The higher is α_S , the more likely the seller is to reject an offer by the buyer, eliciting higher offers for the seller in equilibrium.

Social preferences of the Fehr–Schmidt form are consistent with the experimental findings in EJ of strictly positive offers made by buyers and the investment choice, $I = 1$, of the sellers. However, these preferences still do not explain the role of communication in enhancing investment by sellers. In order to incorporate the role of communication, EJ introduce a cost of breaking a promise. This cost, l , is incurred only if a promise is broken.

Suppose that $\beta_B < \frac{1}{2}$, so that $1 - 2\beta_B > 0$. We know that, in this case, from Proposition 7.23, the seller is offered $\underline{x}_S$ but the seller does not find it worthwhile to invest ($I = 0$). Can the buyer *promise* to offer more to the seller, say, x_S^* , in order to induce first stage investment ($I = 1$)? Let us assume that the buyer does not make overfair offers, so he does not offer to share more than half the net gains from trade. Thus, the buyer promises to give a share, x_S^* , to the seller that satisfies

$$\underline{x}_S < x_S^* \leq \frac{G + F}{2}.$$

If the buyer keeps his promise, then the utility of the buyer is

$$U_B(x_S^*, \cdot) = G - x_S^* - \beta_B [G - x_S^* - (x_S^* - F)].$$

If the buyer does not keep his promise and reverts back to his ex-post optimal level of offer $\underline{x}_S$ then his utility is

$$U_B(\underline{x}_S, \cdot) = G - \underline{x}_S - l - \beta_B [G - \underline{x}_S - (\underline{x}_S - F)],$$

where l is the cost of breaking the promise. It is optimal for the buyer to keep his promise if $U_B(x_S^*, \cdot) \geq U_B(\underline{x}_S, \cdot)$, so the lowest promise of the seller's share is

$$x_S^* = \min \left\{ \frac{G + F}{2}, \underline{x}_S + \frac{l}{1 - 2\beta_B} \right\}.$$

Since we are considering the case $0 < 1 - 2\beta_B < 1$, we have $\frac{l}{1 - 2\beta_B} > l$. Thus, the presence of inequity aversion on the part of the players magnifies the effect of the costs of breaking the promise. If players had entirely self-regarding preferences, so $\beta_B = 0$, then the effect of promises does not get magnified.

Consider now the case of seller communication when there is a cost of not carrying out a threat. Suppose that the seller threatens to reject any unfair offers $x_S < F + \frac{G-F}{2} = \frac{G+F}{2}$. We are interested in the maximum offer $\hat{x}_S < \frac{G+F}{2}$ that the seller can credibly threaten to reject.

Using (7.1), we can find $\hat{x}_S$ from the following condition

$$\hat{x}_S - \alpha_S [G - \hat{x}_S - (\hat{x}_S - F)] - l = -\alpha_S F, \quad (7.4)$$

where l is the cost of not carrying out the threat by accepting any infinitesimal amount less than $\hat{x}_S$. Solving out for $\hat{x}_S$ from (7.4), and using the upper bound $\frac{G+F}{2}$ on $\hat{x}_S$ we get

$$\hat{x}_S = \min \left\{ \frac{l + \alpha_S G}{1 + 2\alpha_S}, \frac{G + F}{2} \right\}.$$

The seller will always get an ex-post offer of $x_S = \frac{\alpha_S G}{1 + 2\alpha_S}$ in (7.2). However, the cost of breaking threats, l , can give the seller an extra minimum payoff of $\frac{l}{1 + 2\alpha_S}$, which is decreasing in α_S . Thus, as the degree of disadvantageous inequity suffered by the seller, α_S , increases, the set of offers that he can credibly threaten to refuse, shrinks. This reduces the commitment value of threats. A simple application of L'Hopital's rule shows that $\lim_{\alpha_S \rightarrow \infty} \hat{x}_S = \frac{G}{2}$, thus, as $\alpha_S \rightarrow \infty$, the seller cannot threaten to reject even $\frac{G}{2}$, which splits the gross gains equally (i.e., not taking account of F). In this case, the ex-post monetary payoff of the seller is $\frac{G}{2} - F$, which for the EJ experiments ($F = 60$ and $G = 100$) is a negative number.

The experimental data in EJ does show that seller threats are not very potent, while buyer promises are far more effective; this qualitative feature is consistent with the proposed theoretical model. However, in order to explain the quantitative features of the data, one has to allow for a distribution of values for α and β (and possibly also of l) across the population.⁵

In contrast to EJ, Charness and Dufwenberg (2006) locate the explanation for the efficacy of promises in *guilt aversion*. Subjects who break a promise experience *guilt* when they let down the payoff expectations of others. An understanding of this mechanism requires the machinery of *psychological game theory* that is considered in Part 4 of the book. A priori, one possible objection to the explanation offered by EJ may be that they assume the very result they wish to prove, by including a cost of breaking a promise in their objective function. There are two responses. First, EJ make very precise predictions that differ between buyers and sellers, which can be tested. Second, one chooses among different explanations of the same phenomenon not by a-priori plausibility but by putting them to a stringent empirical test.

Vanberg (2008) pits the two competing explanations by EJ and Charness and Dufwenberg (2006) empirically against each other. He argues that the experiments by Charness and Dufwenberg are not able to distinguish between the two competing explanations. Thus, he designs experiments in which the predictions of the two explanations can be neatly distinguished. The main finding is that the results are consistent with a model in which people have a preference for keeping their promises, as in EJ.

7.3 Social identity

Individuals may “identify” themselves with certain *social categories*. *Social identity* refers to one's social category. For instance, an individual's identity could be manager or worker, blue collar or white collar worker, black or white, member of a club or a non-member, and soldier or civilian. Each social category is governed by a *norm* of behavior that specifies the *ideal behavior* of a member in each situation.⁶ Different social contexts may trigger different identities—a family identity—a regional identity, or a national identity (Turner et al., 1987).

Let us begin with an example. Akerlof and Kranton (2005, p. 9) describe the following initiation at the US West Point military academy. “On plebes' first day . . . they strip down to their

⁵ Indeed, Fehr and Schmidt (1997) use such a distribution from experimental data to explain their own findings. For instance, α takes the values 0, 0.5, 1, and 4 for, respectively, 30, 30, 30, and 10% of the players. The parameter β takes the values 0, 0.25, and 0.6 for, respectively, 30, 30, and 40% of the players.

⁶ For early developments in social identity theory, see Tajfel and Turner (1979).

underwear. Their hair is cut off. They are put in uniform. They then must address an older cadet, with the proper salute . . . must stand and salute and repeat, and stand and salute and repeat, until they get it exactly right, all the while being reprimanded for every tiny mistake.” How should an economist brought up on the theory of incentives and organizations make sense of this?

Akerlof and Kranton (2005, p. 9) offer an answer: “Economists’ current picture of organizations and work incentives has no place for the West Point program . . . The Army’s aim at West Point is to change cadets’ preferences. They wish to inculcate non-economic motives in the cadets, so that they have the same goals as the U.S. Army. Alternatively stated, the goal of West Point is to change the identity of the cadets, so they will think of themselves, above all else, as officers in the U.S. army. They will feel bad about themselves—they will lose utility—if they fall short of the ideals of such an officer.”

The effect of social identity has also been documented for some of the experimental games considered in Chapter 6. For instance, Ockenfels and Werner (2014) show that dictators transfer a higher amount to responders who are perceived to share a common group identity. Second order beliefs of the dictator are important too. When the dictator knows that the responder is unsure of the common group identity, he transfers a lower amount.

There are three main components of social identity theory. In the first, *categorization*, the individual players are classified into the relevant categories, say, gender, race, nationality, and profession. In the second, *identification*, individuals identify with the characteristics of the category in which they are categorized; perhaps individual identity imparts greater self-esteem. Other members of the same group are often called insiders or *ingroup members* and members of other groups are called outsiders or *outgroup members*. Identification might involve adopting the stereotypical behaviors of one’s group members. It also typically involves favoring the ingroup members over outgroups. In the third, *social comparison*, individuals compare their own group to other groups on some criteria, which is also possibly stereotypical.

Research in *social psychology* has typically employed the technique of inducing *artificial group identities* through a process of *priming* the individuals; see Tajfel et al. (1971) for the classic experiments. The typical experiment then observes the effect of such priming on the behavior of subjects towards ingroup and outgroup members. The experimental process involves finding the *minimal group* such that the weakest possible identification with the group produces discriminatory behavior towards ingroup and outgroup members.

An important finding of this research is that with most minimal groups, and even the most trivial categorization, say, into blue and red groups, there is relatively more favorable behavior towards ingroup members; for surveys, see Tajfel and Turner (1986) and McDermott (2009). Even when subjects know that the assignment to groups is random, they continue to discriminate in favor of insiders (Billig and Tajfel, 1973). These findings suggest that humans might be hardwired, perhaps for evolutionary reasons, to take on social identities that can create both socially harmful and socially helpful outcomes.⁷ Examples of socially harmful outcomes might be intolerance, discrimination, and prejudice, whereas socially helpful outcomes might arise from internalizing desirable norms of hard work and honesty; for instance, when Asians in the US are primed for their Asian identity, they achieve greater educational and professional success.

⁷ Henri Tajfel, one of the pioneers of social identity theory was a Polish Jew who suffered extreme hardships in Nazi concentration camps and lost his family. His interest in social identity theory was possibly fueled by the prejudice and discrimination that he personally experienced.

7.3.1 Evidence on social identity and other-regarding preferences

Benjamin et al. (2010) explain how priming of subjects to induce group identities can be used in experiments to explain the choices of subjects. Suppose that an individual belongs to some category C , and chooses some action, $x \in X$, where X is some finite set of actions; more generally, x could be a vector of actions. If the individual did not belong to any category, then his most preferred action would have been x_P . However, the target action (possibly a norm) for a member of category C is x_C . The individual must choose x in order to maximize his utility, given by

$$U = -w(s) (x - x_C)^2 - [1 - w(s)] (x - x_P)^2, \quad (7.5)$$

where $w(s)$ is the weight that the individual puts on conforming with his category norm, x_C . In (7.5), the individual derives disutility from a weighted combination of departures of his action, x , from the category norm, x_C , and his privately optimal choice, x_P . The variable s may be thought of as the strength of affiliation with one's category, and could reflect the extent of priming in the experiment. The restrictions on $w(s)$ are that $w(0) = 0$ and $w' > 0$, so greater priming induces a greater weight on conforming with group norm. The first order condition is sufficient and solving it we get the optimal choice

$$x^* = w(s)x_C + [1 - w(s)] x_P, \quad (7.6)$$

which is just a weighted average of the group norm and the privately optimal choice. Several conclusions can be drawn from (7.6). First, because $w' > 0$, as the strength of affiliation with one's category, s , increases, the optimal choice x^* becomes closer to the category norm, x_C . Suppose that starting from some level of strength, s , priming of subjects induces an extra affiliation, ε , to one's category, so that the final level of strength is $s + \varepsilon$. Then, expanding $x^*(s + \varepsilon)$ around $x^*(s)$ using a first order Taylor series approximation, and using (7.6),

$$x^*(s + \varepsilon) - x^*(s) \approx \varepsilon \frac{dx^*(s)}{ds} = \varepsilon w'(s) (x_C - x_P).$$

Since $w' > 0$, the effect of priming on optimal choices depends on whether $x_C \gtrless x_P$. Also, $\frac{d}{ds} \left(\frac{dx^*(s)}{ds} \right) = w''(s) (x_C - x_P)$, which depends on the sign of $w''(s)$ and whether $x_C \gtrless x_P$. Thus, a priori it is not necessarily the case that on the margin, priming is more effective for those individuals who identify more strongly with a category. These considerations could be important in understanding the behavior of individuals in social identity experiments.

Benjamin et al. (2010) consider the following stylized facts, using US data. First, Asian Americans accumulate more human capital and save more in tax-deferred savings accounts relative to white Americans. Controlling for observable demographic features, black Americans do relatively worse than white Americans on these counts. Second, women appear to take lower risks than men in financial investments, and also behave relatively more conservatively in the lab. Further, it has been advanced that the norms in each of these groups (by ethnicity and gender) may be responsible for these findings. For instance, it is believed that the Asian American identity is associated with the norm of patience, while the non-immigrant black identity with the opposite norm—impatience and risk taking. Similarly, the male identity may be associated with taking

more risks than the female identity.⁸ Clearly non-experimental data is fairly noisy in this respect because it includes many other confounds, such as peer pressure, income levels, and various socioeconomic aspects. Hence, the authors propose an experimental test.

After being primed for Asian, black, or gender identities, the risk preferences of the subjects and the time preferences are elicited. Priming took the form of background questionnaires in which the questions highlighted racial or gender identities. For instance, there were questions about languages spoken at home and how long the immigrant family had been in the US. In terms of the simple theoretical framework proposed above, x equals risk or time preferences and x_C are the relevant category norms. For instance, if $x_P = \theta_P$ is the discount rate of an unprimed Asian American, then $x_C = \theta_C < \theta_P$ is the discount rate of an Asian American who has been primed for his more patient Asian identity.

The results are as follows. Unprimed Asians are more risk-averse and more impatient than whites. However, when primed with their Asian identity, the impatient choices drop by 14%, which is significant at the 1% level. Primed Asians require much lower interest rates to defer payments from a current date to a future date.⁹ For a current amount of \$4 and payments deferred by a week, the elicited percentage interest rate falls by three-quarters when primed for the Asian identity.

The difference between the impatient choices of primed and unprimed whites and primed and unprimed immigrant blacks are not statistically significant. However, contrary to the received view, in terms of their time preference, primed native blacks respond in the same direction as primed Asians, i.e., become more patient. Individuals who are primed for gender do not make choices that are any different from the unprimed condition.¹⁰ Shih et al. (1999) find that when subjects are primed separately for Asian and gender identities, and given a maths test, they do best when primed for the Asian identity and worse when they are primed for the female identity.

Chen and Li (2009) consider the effect of social identity on other-regarding preferences. Each of their treatments has four stages. In the first stage, the subjects are categorized into groups. The groups are formed, based on the preferences of subjects for paintings by one artist over another artist.¹¹ In stages 2 and 3, the group identities are strengthened through two tasks. In the first task, subjects were quizzed about the names of artists for a set of paintings. Correct answers were rewarded with tokens that could be cashed for money from the experimenter. Subjects were allowed to consult their group members (but not the members of the other group) in searching for the correct answer. In the second task, subjects were given a set of tokens and had to decide, over five rounds, how to allocate them to other players. The identity of the other players was varied between ingroup and outgroup members. The players knew that one of the rounds would then be randomly picked and played for real money. Two further control groups were constructed in which the identity of the subjects towards their ingroup members was not primed in this manner.

⁸ We refer the reader to the large number of references on these claims in Benjamin et al. (2010).

⁹ Denote current income as Y_0 , next year's income as Y_1 , annual interest rate as r and the annual discount rate as θ . Then the elicited interest rate r^e makes the decision maker indifferent between the two choices: Y_0 and Y_1 . Thus, $Y_0 = \frac{Y_1(1+r^e)}{1+\theta}$, so a lower r^e reveals a lower θ , and a more patient individual.

¹⁰ This is consistent with the results of Cadsby and Maynes (1998) who also find that gender has little effect on contributions in public good experiments when subjects are primed for gender identities in a pregame questionnaire.

¹¹ The painters are two modern artists. Based on the names of the painters, the two groups are the "Klee" and "Kadinsky" groups. To control for possible salience effects of the names of the artists, two control groups—the Maize and Blue groups were also used.

Once group identity was induced and enhanced in this manner, in the final stage, stage 4 of the experiment, subjects played several two-person sequential move games from Charness and Rabin (2002). In these games, player 1 chooses between two actions, stay out or enter. If player 1 stays out, the game ends with exogenously given payoffs, and if player 1 enters, then player 2 chooses between two possible payoff allocations for the player. The strategy method is used to elicit the choices. Player 2 is asked to make a choice in each of the two cases, assuming that player 1 has stayed out and that player 1 has entered. The games are differentiated by the actions of both players. For instance, in making his choice, player 2 could choose an option where no sacrifice is involved in ensuring that player 1 gets a higher payoff, or where such a sacrifice is involved, or where the sacrifice actually reduces the payoff of player 1. The choices reveal the extent to which players appear to exhibit elements of other-regarding preferences towards others.

The results are as follows. For stage 2 of the treatment, players allocated significantly more tokens to ingroup members relative to outgroup members. This replicates the often found result that minimal groups favor ingroup members. From the choices made in stage 4 of the experiment, players exhibit altruism towards opponents if their payoff is greater and envy if it is lower. However, altruism and envy are relatively stronger towards ingroup members.

Other interesting results arise by observing the behavior of player 2 when the decision of player 1 (to enter or stay out) conveys information about intentions. In this case, results can be further sharpened by observing cases where player 2 sacrifices own-payoff to get player 1 a better payoff, and where the sacrifice in own-payoff leads to lower payoff for player 1. Here, the results, based on a logit specification, are as follows. Players are 18.6% likely to reward ingroup members more than outgroup members. Social preferences matter in the sense that player 2 is less likely to reward player 1 if player 2's payoff is relatively lower—so envy limits positive reciprocity.

For every 100 token disadvantageous gap, player 2 is less likely to reward player 1 by 7.8%. Since the strategy method is used to elicit the choices of player 2 for both possible choices of player 1 (stay or enter), we can determine if player 2 is likely to reward or punish unfavorable choices. It turns out that player 2 is more likely to reward favorable behavior by an ingroup player 1 and also more likely to forgive unfavorable behavior by an ingroup player 1. Finally both players are more likely to choose the joint payoff maximizing choice when they are ingroup members.

Cohn et al. (2013) conduct a field study of prison inmates from Switzerland's largest adult penitentiary. Their aim is to see if imprisonment fosters a criminal identity. In the baseline treatment, criminals were asked to report the number of heads in 10 tosses of a coin that they privately observed. The reported outcomes were then compared to the statistical average of 50% heads. Criminals misreported 20% of the times. When criminals were primed for their criminal identity (e.g., by asking what they were convicted for) they cheated even more—now they misreported 32% of the times. By contrast, in a placebo condition, subjects were recruited from the general population. In this case, when primed for a criminal identity (reading a short text profile describing a criminal or a non-criminal person), the percentage of misreported heads fell from 14% to 10%.

In Charness et al. (2007), subjects are assigned to one of two groups, row players (R) and column players (C), in the prisoner's dilemma game (PD) and the battle of sexes game (BOS). These games are shown in Figure 7.3.¹² The PD has a unique Nash equilibrium (B,B). The BOS game has two pure strategy Nash equilibria (A,A) and (B,B) and a mixed strategy Nash equilibrium in which each player plays his preferred strategy (A for row player and B for column player) with probability 0.75.

¹² Part 4 of the book considers the experimental results on these two class of games in detail.

Battle of sexes game			Prisoner's dilemma game		
	A	B		A	B
A	3, 1	0, 0	A	5, 5	1, 7
B	0, 0	1, 3	B	7, 1	2, 2

Figure 7.3 The battle of sexes (BOS) game and the prisoner's dilemma game (PD).

Source: Charness et al. (2007) with permission from the American Economic Association.

Players in the PD game are said to be *aggressive* if they play B and *accommodating* if they play A. In the BOS game, aggressive players play their preferred strategy, and accommodating players play the opponent's preferred strategy. Players are randomly assigned as row players (category label R) and column players (category label C). Once the category labels (R and C) are assigned to the players, they go to their respective rooms labeled R and C. If a game is played in room R, then the row player is termed as the host and the column player is termed as the guest. Similarly, in room C, the column player is the host and the row player is the guest.

The treatments exploit two dimensions. In the *audience treatment*, the other group members sit in a semi circle behind the player when he makes the decision, but do not otherwise interfere in actual choices made. This is intended to induce greater salience of group membership. In the *no-audience treatment*, other group members do not watch. In the second dimension, the *feedback treatment*, all people in the room immediately learn the outcome, while in the *no-feedback treatment*, the payoffs are revealed only at the end of the session. The payoffs of any player are made up of two components—their own-payoffs in games when they make active decisions and a one third stake in the payoffs of one's ingroup members when others in the ingroup play the games.

When group membership is more salient (audience treatment) subjects are more likely to choose the action that maximizes their own payoff. For the row player, these actions are B in the PD game, and A in the BOS game. However, when this salience is reduced (no-audience treatment) these aggressive choices are reduced. In the audience treatment, in the BOS, hosts are aggressive in 84% of the cases and guests accommodate in 65% of the cases. By contrast, in the no-audience treatment, these numbers are significantly lower at 60% and 38% respectively. Thus, hosts are more aggressive when they play in front of an audience of their ingroup members. Guests seem to anticipate this behavior and are more accommodating. Similarly, the feedback treatment makes group membership more salient and creates effects similar to those in the audience treatment. When group membership is most salient, in the treatment with audience effects and feedback effects, almost all hosts are aggressive (95%) and this is again well anticipated by the guests who accommodate even more.

In the PD game, hosts accommodate in 62% of the cases when there is no audience and feedback. However, when group membership is most salient (audience plus feedback) hosts only accommodate in 28% of the cases. The effect of group membership is smaller but still significant for guests. In the case of no audience and feedback, guests accommodate in 45% of the cases but this reduces to 34% in the audience and feedback treatment.

Goette et al. (2012) consider the external validity of minimal groups (MG) created in the lab in which ingroup members may have limited interaction with each other. They argue that real-world groups are, instead, characterized by strong interactions among ingroup members. These members have social ties and friendships between each other, which may generate empathy and prosocial behavior within the group. Denote the treatment with real social groups by SG treatment.

The main research objective is to examine the behavior of subjects in minimal and real social groups. In each of the MG and SG treatments, members of groups are randomly assigned, rather than endogenously assigned. Individuals only knew whether the player they interacted with is an ingroup or outgroup member but not the identity of the player; this rules out strategic, reputation building, motives.

The subjects in the experiment were Swiss army trainees who stayed in the same platoon once they joined the Joint Officer Training Program (JOTP). The subjects had performed a range of joint tasks together, possibly enhancing their group identity. The random assignment of the trainees to a platoon (each platoon has 21 members) comprises one *real* social group (SG). Assignment to a platoon enhances social identity among platoon members as they do a variety of joint tasks during a three week period. In contrast, the minimal group, MG, is formed simply on the basis of random choice, using the social security numbers of the army trainees. This is a minimal group because all members are part of the JOTP, yet, unlike the SG, may belong to different platoons, so have not necessarily interacted over the three weeks of platoon training. In the experiments, subjects are simply told of the ID number of the platoon of the opponent as a method of identifying ingroup and outgroup members.

The subjects played a modified trust game, akin to a simultaneous move prisoner's dilemma. Two players, A1 and A2, were given 20 points. They simultaneously decide to send up to 20 points to the other player. The experimenter doubled the points that were passed to the other player. Sending 0 points may be considered to be the *defection decision* and sending all 20 points may be considered to be the *cooperative decision*. In the ingroup condition, the opponent was a member of the same platoon, while in the outgroup condition, he was a member of another platoon.

In experiment 1, players simply choose the number of points to send to the other player. In experiment 2, the possibility of third party punishment was allowed with the introduction of players B1 and B2, each of whom was allocated 70 points. Having observed the actions of player A_j, $j = 1, 2$, player B_j could punish player A_j. Punishment took the form of each player B_j being allowed to spend up to 10 deduction points aimed at player A_j. For each point surrendered by player B_j, player A_j lost 3 points. This allows one to test for altruistic norm enforcement. Players B1 and A1 form an ingroup and players B2 and A2 form another ingroup. The payoffs (one point equalled 0.25 CHF) were mailed to the home addresses of the trainees, 10 days after the end of the course, on completion of the JOTP.




Figure 7.4 Cooperation rates and beliefs in the modified trust game. Bars represent standard errors of the means.

Source: Goette et al. (2012) with permission from the American Economic Association.

Panel A in Figure 7.4 shows the cooperation rates in the MG and the SG treatments. In Figures 7.4, 7.5 we use the convention that other A players are all referred to as A2 players and all players that a player B can punish are called A1 players. In both treatments, MG and SG, there is significant, positive, ingroup bias in cooperation rates. Averaged across both treatments, the cooperation rate with ingroup members is 65%, while that with outgroup members is 50%. The difference in cooperation rates towards ingroup and outgroup members in the MG treatment is 10%, but it is 20% in the SG treatment. Thus, there is a significant quantitative difference in the two treatments. The actual differences in cooperation rates are also evident in the beliefs of the players (Panel B); ingroup members are believed to be significantly more cooperative.

Figure 7.5 shows that the behavior in treatments MG and SG is qualitatively different. In the MG treatment (Panel B), when player A1 defects, the punishment is smaller for ingroup members relative to outgroup members. But when A1 cooperates, punishment of ingroup members and outgroup members is similar in both MG and the SG treatments; some punishment still takes place because players do not fully cooperate. In the SG treatment (Panel A), punishment of outgroup members for defection is nearly identical to ingroup members. Another interesting finding is that defection against ingroup members is punished relatively more highly in the SG treatment (see Figure 4 in Goette et al., 2012). By contrast, there is no such difference in the MG treatment. Thus, third party norm enforcement is stronger when the group identity is stronger (SG treatment).

These qualitative and quantitative differences between the SG and MG treatments can potentially reconcile the differences in existing results between minimal and real groups. For instance, for real groups (akin to the SG treatment) Bernhard et al. (2006) find that tribes in Papua New




Figure 7.5 Group membership and punishment. Bars represent standard errors of the means.

Source: Goette et al. (2012) with permission from the American Economic Association.

Guinea punish ingroup and outgroup members equally for defection. However, Chen and Li (2009) using minimal groups (akin to the MG treatment) find that ingroup members are punished less than outgroup members.

Yamagishi and Kiyonari (2000) find that the effects on cooperation rates are different in a simultaneous and sequential move prisoner's dilemma game. The argument is that in the sequential game, non-cooperative behavior is punished anyway by the second mover, removing the need for group identity to sustain cooperative behavior. Guala et al. (2013) show that common knowledge of group affiliation enhances cooperation rates in a prisoner's dilemma game.

Chen and Chen (2011) consider *minimum effort games* that are considered, in detail, in Part 4 of the book. In such games, the payoffs of players are decreasing in their own action, but increasing in the minimum action of the group. There are several possible symmetric Nash equilibria in which all players choose the same action. In the Pareto efficient Nash equilibrium, all group members choose the highest possible action. The authors show that when subjects are not primed for any identity, they choose an inefficient action. However, once individuals are primed for group identity (e.g., by undertaking group tasks), and learning is allowed, then individuals choose the efficient equilibrium.

Hargreaves Heap and Zizzo (2009) consider a trust game played by players from two groups of players. They verify the usual finding of more cooperative behavior towards ingroup members relative to outgroup members. They argue that while this clearly improves the welfare of ingroup members, it reduces the welfare of outgroup members. On net, the outcome may reduce aggregate welfare, which is a potentially negative outcome of group affiliation. What accounts for the psychological benefits from joining a group? In order to elicit the value that people place on belonging to a group, they run artificial markets in which participants can trade in their group membership for a price. The elicited prices suggest that people derive psychological benefits that exceed the material benefit of belonging to the group. When one adds these psychological benefits to the net welfare computed above, the overall social welfare of group formation is non-negative.

7.3.2 Identity and economic incentives

Identity economics is the study of identity, norms, and ideal behavior. It offers a radically different view of incentives, one that is not new to sociologists but would be to many economists. Several decades of experiments in social psychology show that it might be optimal for organizations to alter the identity of its members, so that they internalize the mission and objectives of the organization and think of themselves as *insiders* in the organization. Clearly, if this can be done successfully, then identity may substitute, perhaps partly, for the monetary costs of incentives to ensure that members work to fulfill the objectives of the organization that they work for. We draw on Akerlof and Kranton (2000, 2005, 2008) in our discussion of the relevant issues in this section.

Consider a prototype principal-agent model in which the agent can choose among two possible actions, high and low, H, L , which give rise, respectively, to the effort levels, $e_H > e_L$. The profits of the principal can either be high, π_H , or low, π_L , $\pi_H > \pi_L$. The actions of the agent induce a probability distribution over profits. Under action H , the profits π_H, π_L are equally likely, while under action L , the profit level π_L occurs with probability 1. The principal cannot observe the effort level but can condition wage payments to the agent, w , on the profit levels, π_L, π_H . For any wage-effort pair, (w, e) , define

$$u(w, e) = \ln w - e. \quad (7.7)$$

We shall enrich this prototype model below.

Denote the identity of the agent by $I_c \in \mathbb{R}$, where c indexes the identity of the agent. We distinguish between two identities. In the first, the agent identifies himself as part of the group that he works with; we call this the *insider identity* or *group identity* ($c = G$). In the second, the agent considers himself to be an outsider; we call this the *outsider identity* ($c = O$). Identity $c \in \{O, G\}$ is associated with the norm of an ideal effort level, e_c . The norm of effort for insiders is likely to be higher than for outsiders, so $e_G > e_O$. Thus, it is cheaper to elicit higher effort from insiders. The utility of the agent when his identity is c , is given by

$$U(w, e; c) = u(w, e) + I_c - \mu_c |e_c - e|, \quad (7.8)$$

where u is defined in (7.7), and $\mu_c > 0$ is a scaling parameter. In standard textbook economics models, the utility of the agent is given by $u(w, e)$. By contrast, in models of identity, the agent additionally derives direct utility from his identity, I_c , and disutility from deviations of effort from the norm for the identity group (last term in (7.8)).

Since $e_G > e_O$, the principal needs to offer relatively lower wage differentials (across the two effort levels, high and low) to insiders, for eliciting high effort. In other words, when there are two effort levels, then, under an insider identity one would expect to observe flatter wage profiles. This result need not hold under more than two effort levels, because an insider identity may also induce the principal to elicit an even higher effort level. Thus, incentives and identity are complements when there are two effort levels but may become substitutes when there are more than two effort levels.

The firm may consider altering the worker's identity from an outsider to an insider if the cost of doing so is low enough. The benefits are two fold. First, the incentive cost of eliciting the same effort level is lower. Second, a lower variation in wage is preferred by the risk averse insider, reducing the incentive costs of the firm. We now consider some of these issues in a model of optimal hierarchical organizational design (Akerlof and Kranton, 2008).

Suppose that a risk-neutral principal hires an agent with the preferences given in (7.8). The principal can either hire a supervisor or not. The supervisor obtains hard evidence of action L with probability $p > 0$, but cannot obtain verifiable evidence of action H . In the event that the supervisor reports hard evidence of a low level of effort, e_L , the principal levies a fine $f > 0$ on the agent.

The presence or the absence of the supervisor alters the identity of the agent and this is borne out by evidence from social psychology experiments (Jex and Britt, 2008). In the presence of the supervisor, the agent feels as an outsider (identity $c = O$) and in the absence of the supervisor, the agent feels trusted, so he has an insider or group identity (identity $c = G$). The principal understands, correctly, the agent's identity under various regimes.

The assumptions on the ideal effort levels are as follows: (1) The ideal effort level of an outsider is $e_O = e_L$. (2) The ideal effort level of an insider is the effort level $e_G = e_M$, corresponding to some action M on the part of the agent such that $e_L < e_M < e_H$. Under action M , the probability of getting high profits, π_H , is $\frac{\theta}{2}$ and the probability of getting low profits π_L , is $1 - \frac{\theta}{2}$, where $0 \leq \theta \leq 1$. The outside option of the agent is given by $\bar{u}$.

The principal offers a profit-contingent wage to the agent, (w_H^c, w_L^c) , where w_j^c is the contracted wage when the observed profit is π_j , $j = L, H$, and the identity of the agent is $c \in \{O, G\}$. Effort

is non-contractible. The expected profit of the principal from inducing the high effort level, e_H , when the agent's identity is c is given by

$$\Pi(c) = \frac{1}{2} (\pi_H - w_H^c) + \frac{1}{2} (\pi_L - w_L^c).$$

The expected profit, if other effort levels are to be induced can be found analogously.

PRINCIPAL APPOINTS A SUPERVISOR (REGIME-I)

When the principal appoints a supervisor (Regime-I), the agent takes the identity of an outsider, so $e_O = e_L$. Suppose that the principal wishes to induce the effort e_H . The expected profit of the principal is

$$\Pi(O) = \frac{1}{2} (\pi_H - w_H^O) + \frac{1}{2} (\pi_L - w_L^O).$$

The individual rationality constraint (IR), ensures that the agent receives at least as much as the outside option.

$$\frac{1}{2} \ln w_H^O + \frac{1}{2} \ln w_L^O - e_H + I_O - \mu_O |e_L - e_H| \geq \bar{u}. \text{ (IR)}$$

The incentive compatibility constraint (IC) of the agent ensures that action H is preferred to action L . Note that the identity of the agent stays the same under both actions and depends only on the organizational structure (supervisor or no supervisor). If the agent chooses action L , the outcome π_L obtains with probability 1. An audit by the supervisor reveals, with probability p , that the agent chose L in which case the agent is fined an amount f . Hence, the IC is given by

$$\frac{1}{2} \ln w_H^O + \frac{1}{2} \ln w_L^O - e_H - \mu_O |e_L - e_H| \geq (1 - p) \ln w_L^O + p \ln(w_L^O - f) - e_L. \text{ (IC)}$$

Finally, there is a limited liability constraint (LLC), which ensures that the agent's payoff after paying a fine cannot be driven below $l > 0$.

$$0 < l \leq w_L^O - f. \text{ (LLC)}$$

The IR binds in equilibrium. If not, then the principal can increase profits by reducing the incentive payments w_H^O and w_L^O . The IC also binds in equilibrium. If not then the principal can increase his profits by reducing w_H^O . The RHS of the IC is decreasing with respect to p , hence, the payoff of the agent from taking the action L is lower under better supervision. However, since the IC must bind, there are two ways of restoring the binding IC after an increase in p : Either w_H^O goes down or w_L^O goes up or a bit of both. We collect these results in the next lemma.

Lemma 7.8 *When the principal appoints a supervisor, then more effective supervision associated with a higher probability of detection, p , induces a flatter optimal wage profile (the gap between w_H^O and w_L^O is lower). This also reduces the expected cost of the incentive payment to the principal.*

THE PRINCIPAL DOES NOT APPOINT A SUPERVISOR (REGIME-II)

When the principal does not appoint the supervisor (Regime-II), he chooses not to monitor the agent, so the agent has an insider or group identity, $c = G$. The norm-driven ideal effort level of the agent is $e_G = e_M$. Suppose, initially, that the principal wishes to induce the high action, H . The expected profit of the principal is

$$\Pi(G) = \frac{1}{2} (\pi_H - w_H^G) + \frac{1}{2} (\pi_L - w_L^G).$$

The individual rationality constraint of the agent is

$$\frac{1}{2} \ln w_H^G + \frac{1}{2} \ln w_L^G - e_H + I_G - \mu_G |e_M - e_H| \geq \bar{u}. \text{ (IR)}$$

There are two incentive compatibility constraints, IC1 and IC2. The principal needs to ensure that the agent prefers the effort level e_H to e_M (IC1) and the effort level e_H to e_L (IC2). Recall that the probability of outcome π_H under the action e_M is $\frac{\theta}{2}$.

$$\frac{1}{2} \ln w_H^G + \frac{1}{2} \ln w_L^G - e_H - \mu_G |e_M - e_H| \geq \frac{\theta}{2} \ln w_H^G + \left(1 - \frac{\theta}{2}\right) \ln w_L^G - e_M \text{ (IC1)}$$

$$\frac{1}{2} \ln w_H^G + \frac{1}{2} \ln w_L^G - e_H - \mu_G |e_M - e_H| \geq \ln w_L^G - e_L - \mu_G |e_M - e_L|. \text{ (IC2)}$$

As in Regime-I, there is also a limited liability constraint in Regime-II. Note in IC1 and IC2 that the identity of the agent stays the same under different effort levels in this regime. The lack of supervision has costs and benefits. The cost is that, ex-post, the agent is not punished for exerting a low effort level. The benefit is that with a group identity, the ideal effort level of the agent is higher, $e_M > e_L$. Since $e_L < e_M < e_H$, it follows that $|e_M - e_H| < |e_L - e_H|$, thus, the individual rationality constraint is easier to satisfy in Regime-II relative to Regime-I. So, the wage costs are also relatively lower in Regime-II.

When might the principal not wish to induce the high effort level, e_H , in Regime-II and settle for the intermediate effort, e_M ? To answer this question, rewrite IC1 as follows

$$\frac{(1-\theta)}{2} (\ln w_H^G - \ln w_L^G) - (1 + \mu_G) (e_H - e_M) \geq 0. \text{ (IC1)}$$

As $\theta \rightarrow 1$, it becomes prohibitively expensive to elicit the effort level, e_H ; in this case one would need an extremely steep wage profile for IC1 to hold. The reason is that, as $\theta \rightarrow 1$, an agent with identity G finds that he gets only marginally higher wage payments by exerting the higher effort level, e_H , but substantially greater disutility from having to compromise on the ideal effort level, e_M . Under such identity concerns, the principal may find it optimal to induce the ideal effort level of the agent, e_M . In this case, the relevant incentive compatibility constraint, which ensures that the agent chooses the effort level e_M rather than e_L is given by IC3:

$$\frac{\theta}{2} \ln w_H^G + \left(1 - \frac{\theta}{2}\right) \ln w_L^G - e_M \geq \ln w_L^G - e_L - \mu_G |e_M - e_L|. \text{ (IC3)}$$

Or,

$$\frac{\theta}{2} (\ln w_H^G - \ln w_L^G) \geq (1 - \mu_G) (e_M - e_L). \text{ (IC3)}$$

The factors that make it easier for IC3 to hold and, hence, more likely that the principal will successfully induce e_M , are (1) $\theta \rightarrow 1$, $\mu_G \rightarrow 1$, and (2) the gap between e_M and e_L is small.¹³

¹³ For a richer and more encompassing model of identity, see Bénabou and Tirole (2011). Benjamin et al. (2010) consider the effect of ethnic, racial, and gender identity norms on risk and time preference.

CHAPTER 8

Incentives and Other-Regarding Preferences

8.1 Introduction

Traditional economic theory extolls the virtues of incentives, typically high-powered, in models where individuals have self-regarding preferences. By contrast, real-world incentives are typically low-powered. Consider the quote from the empirical study by Lemieux et al. (2009): “in the United States only 37 percent of individuals have some kind of pay-for-performance component to their compensation and that among those who do the median magnitude of these incentives is only 3.5 percent of their base wage.”

Models of *multitask agency* have already shown that high-powered incentives may be counter-productive. For instance, in the presence of high rewards on easily observed tasks, agents may put disproportionately less effort into tasks that are valuable for the principal, yet are hard to observe and reward formally. However, the experimental results show that even when agents have a single task, the optimal incentives may be low powered. This does not mean, of course, that incentives are unproductive. For instance, recall that the evidence in Chapter 5 shows that punishments (negative incentives) are efficacious in public goods games. However, the evidence from behavioral economics shows that the effect of incentives on human behavior is far more nuanced than can be accounted for by models of self-regarding preferences.

In Section 8.2, we consider the effects of other-regarding preferences on incentives in the presence of *moral hazard*. The agent’s effort is either not observable or is observable but not verifiable, hence, wage contracts can be conditioned only on the observable and verifiable outcomes, such as output levels.

Section 8.2.1 considers the effect of inequity aversion on optimal wage schemes when a principal observes a noisy signal of the agent’s effort level. Following Itoh (2004), consider two effort levels, two output levels, and a self-regarding principal. The agent is other-regarding (in the Fehr–Schmidt sense) and suffers inequity aversion with respect to the principal. The presence of moral hazard gives rise to rents for the worker, as in models of self-regarding preferences. However, the principal needs to offer even higher rents to an inequity averse agent to compensate for the agent’s disadvantageous inequity, relative to the principal.

Suppose that there are several agents who suffer from inequity aversion with respect to each other but not with respect to the principal. Each agent works on an independent project, but inequity aversion creates interlinkages between the outcomes on the projects. In this case, the choice among alternative forms of compensation schemes such as *team incentives*, *relative*

performance evaluation, and *independent contracts*, depends on the degree of advantageous and disadvantageous inequity suffered by the agents. For instance, if the advantageous inequity of the winning agent is too high, then relative performance evaluation creates higher disutility to the winning agent. Hence, to ensure satisfaction of the incentive compatibility constraint, the principal needs to offer a higher level of rents. By contrast, team incentives do not create inequality aversion, so they may be preferred to relative performance evaluation.

Inequity aversion reduces a player's utility when, relative to another player, his payoff is greater (altruism) or lower (envy). By contrast, *reciprocity* induces players to reciprocate kind actions of others by kindness and unkind actions by unkindness. Section 8.2.2 considers the implications of reciprocity for optimal contracts in a principal–agent problem with moral hazard. Ideally, one would wish to use the machinery of psychological game theory (see Part 4 of the book) to model issues of reciprocity. This allows for the endogenous determination of beliefs about others. However, here we follow Englmaier and Leider (2012) in modeling reciprocity in a simpler, reduced-form, manner.

Suppose that the outside option of the agent is zero. If the principal leaves positive rents to the agent, the agent perceives this to be a kind action and may reciprocate by aiding the principal in increasing profits. The implications for optimal contracts under moral hazard are profound, even in this reduced-form model. For instance, it has been a puzzle why real-world contracts are so low-powered relative to the predictions of agency models of moral hazard with self-regarding preferences. In the presence of reciprocity, however, it can be shown that even a *flat wage schedule* can implement high levels of effort, provided that the principal leaves sufficiently high rents with the agent. The reason is that the reciprocity induced by these rents elicits high effort level from the agent, as in gift exchange games.

Section 8.3 considers a principal–agent problem with non-verifiable effort and other-regarding preferences of the Fehr–Schmidt form (Fehr et al., 2007). The principal is uncertain whether the preferences of the agent are self-regarding or other-regarding. In addition to the standard *incentive contract*, we consider here a *trust contract* and a *bonus contract*. In the incentive contract, the principal can invest in a costly monitoring technology that gathers hard evidence of the shirking of the agent with some probability. In the event that shirking is discovered, the principal can levy a fine. Clearly, if the detection technology is effective and the level of fines is high enough, the agent will never shirk under an incentive contract. In a trust contract and a bonus contract, by contrast, the principal foregoes investing in the monitoring technology, even when such a choice is available. This is interpreted by the agent as *kind intentions* on the part of the principal. By contrast, the offering of an incentive contract is interpreted as unkind intentions.¹ Thus, other-regarding agents behave as if they were self-regarding in the presence of incentive contracts but other-regarding in the presence of trust and bonus contracts.

In a trust contract, the principal offers a fixed, contractible, wage despite not having the ability to monitor the agent. Clearly, the success of such a contract relies on the reciprocity of agents, as in gift exchange experiments. In a bonus contract, by contrast, the principal promises a bonus at the end of the game, after having observed the outcome. If the agents believe in the principal's reciprocity, then they know that if they work hard, they will be rewarded ex-post; this is true both for self-regarding and other-regarding agents. The incentive contract is least incomplete (effort can always be enforced with high enough detection probability and fines), the trust contract is more incomplete (effort can never be enforced), and the bonus contract is the most incomplete (effort and the bonus level are unenforceable).

¹ This is consistent with the empirical results that are reviewed below.

The benefit of the incentive contract is that it allows for the monitoring and punishment of shirking agents, while the cost is that such contracts cannot tap into the reciprocity of the other-regarding agents. Hence, they suffer on account of lower effort levels and output levels. The advantage of the trust and the bonus contracts is that they can tap into agents' reciprocity, but the disadvantage is that they cannot dissuade purely self-regarding agents from shirking. The eventual outcome depends on these competing trade-offs. The greater the probability that the agent is self-regarding, the more successful will an incentive contract be. The other-regarding principal suffers from disadvantageous inequity when he encounters a self-regarding agent. But if the disadvantageous inequity parameter of the principal is low, the costs of offering a trust and a bonus contract are low relative to the positive reciprocity obtained from the other-regarding agent.

The empirical evidence on contractual choice in the presence of such contractual incompleteness issues is considered in Section 8.3.3 (Fehr et al., 2007; Fehr and Schmidt, 2007). It turns out that the bonus contract is empirically the most widely chosen contractual form. The effort levels and the profits of the principal are also the highest under a bonus contract. By contrast, if the principal and the agent were purely self-regarding, the incentive contract always dominates the bonus contract.

We also examine the empirical choice of organizational forms in the presence of the classical hold-up problem between two parties (Fehr et al., 2008). The prediction of the model with self-regarding preferences is that residual control rights should be given to the party that makes the investment. Joint-ownership, by contrast, is predicted to stifle investment because it creates the problem of ex-post holdup. However, the empirical evidence shows that when given a choice, most players choose joint ownership because it is more conducive to reciprocal behavior.

In Section 8.4, we introduce the possibility of *two-sided asymmetric information* and *long-term contracts* in a principal–agent relationship (Brown et al., 2004, 2012). Our aim is to characterize the perfect Bayesian equilibrium in the model. The firm (principal) and the worker (agent) both have asymmetric information about each other's type; types can be either self-regarding or other-regarding. In a model with a finite number of periods, we allow firms to either hire workers in a spot market condition, or choose to endogenously maintain long-term relations with them.

When workers have other-regarding preferences, in the last period, they will not choose their effort opportunistically. Their inequity aversion induces an effort level higher than the minimum effort level that would be chosen by a purely self-regarding agent. Analogously, other-regarding principals will wish to offer rents in the last period. In the penultimate period, workers with self-regarding preferences may not wish to shirk because if they are revealed to be shirkers they will be fired from their job and forgo rents in the final period. One can work this logic backwards and ensure that workers get rents in equilibrium and, irrespective of their type, put in an effort level greater than the purely self-regarding one. The evidence supports the predictions of this model. One observes the endogenous formation of long-term relations. Effort levels, profits, and rents are higher in long-term relations, relative to spot-market or short-term relations. The formation of social ties in the relation induces even greater cooperation in long-term relations (Van Dijk et al. 2002).

Section 8.4.4 addresses the ambitious question of why firms exist. Or, under what conditions will an employer choose an *employment contract* over a *sales contract*? An employment contract gives greater ex-post flexibility to the employer, but also opens up the possibility that the employer may exercise his residual control rights in a manner that is disadvantageous to the employee. By contrast, a sales contract prevents the possibility of ex-post abuse of the contract by the employer but suffers from the absence of the ex-post flexibility of an employment contract. It turns out that the presence of other-regarding preferences and long-term endogenous contracts play a key

role in identifying a new set of transaction costs that explain the existence of vertical integration (Bartling et al., 2013). We shall consider the relevant theoretical framework, but particularly the empirical evidence, in some detail.

In contract theories with self-regarding preferences, cleverly designed incentives typically have a positive impact on effort, output, and profits. However, it is well known that even with self-regarding preferences, incentives may not always have the desired effects. First, incentives can be conditioned only on a limited number of variables that may be observable and verifiable, such as revenues or output. Second, under *multi-task agency*, where the agent simultaneously chooses effort levels on several distinct tasks, which are observed with error, incentives for more easily observable and verifiable tasks may shift attention away from other, valuable, but less-well observed tasks (Holmstrom and Milgrom, 1991). Third, when tournaments among workers are organized, they might sabotage each other or not cooperate with each other. A large amount of experimental and field evidence now demonstrates the presence of other-regarding preferences in worker–firm relations.² This creates further issues with the use of incentives in principal–agent problems.

An emerging literature explores the effectiveness of incentives, not only by using other-regarding preferences, but also by invoking additional features such as identity, reputational concerns, asymmetric information, self-signaling, and social signaling motives for one's actions.³ For instance, anonymous donations are rare, yet are considered to be very commendable, and boasting of one's charitable donations is looked down upon. Clearly, models that incorporate social signaling of one's intentions are a promising candidate to explain such phenomena. Models that explain monetary charitable giving as a signal of one's wealth cannot sometimes explain in-kind donations such as volunteering for social causes and donating blood. In taking such actions, the individual could be self-signaling respectable human qualities such as altruism or kindness.

In Section 8.5, we consider theoretical models of *intrinsic* and *extrinsic* motivation. Extrinsic motivation refers to human endeavor that arises purely as a response to external material incentives. Any remaining motivation that is independent of incentives offered by another party, is intrinsic. Section 8.5.1 considers first a principal–agent framework in which the agent has asymmetric information about the cost of the project (Bénabou and Tirole, 2003). The agent observes a private signal of the cost and the level of bonus offered by an informed principal. We show that in a perfect Bayesian equilibrium of the model, higher levels of bonus offered by the principal convey bad news to the agent about the cost. Hence, incentives may deter the effort of the agent. These sorts of models can potentially explain situations where one party to a formal or informal contract has private information (child rearing, education, monitoring of employees).

One needs a different theoretical framework to explain *prosocial decisions* such as blood donations, charitable giving, and voting (Bénabou and Tirole, 2006). In this case, incentives may signal to others that the individual's motivation is not intrinsic. Such incentives could, for instance, take the form of direct incentives for blood donations and tax-exempt charitable donations. Hence, there is a trade-off between appealing to an individual's extrinsic motivation and a cost in terms of the reputational payoffs that the individual perceives are being lost by responding to such incentives. These theoretical models do not, however, address the question of whether they can explain the experimental evidence from the dictator, ultimatum, trust, and public goods games.⁴

² See for instance, Fehr et al. (1997), Fehr and Schmidt (1999), Fehr and Gächter (2000), Bandiera et al. (2005), Falk et al. (2005), Fehr et al. (2007), DellaVigna (2009), Sloof and Sonnemans (2011).

³ See Bénabou and Tirole (2003, 2006, 2011) for an extensive discussion of these issues.

⁴ But also see Benabou and Tirole (2011), whose motivation is to explain experimental evidence that may be anomalous for some of the models of social preferences, such as reluctant altruism in dictator games.

Section 8.6 considers the empirical evidence on extrinsic and intrinsic motivation. Section 8.6.1 shows that incentives may signal a *hostile intent*. In gift exchange games, the effort level is higher when no incentives are used, relative to negative incentives (fines) and positive incentives (bonuses). Within the two incentive treatments, effort levels are higher under positive, as compared to negative incentives (Fehr and Gächter, 2001). In trust games, when a fine is available but not used, it produces the highest level of trustworthiness (Fehr and Rockenbach, 2003).

Section 8.6.2 shows that incentives may conflict with the desire of agents for *autonomy* and *independence* (Deci and Ryan, 1985; Deci et al., 1999). In trust games, when the trustor announces a target back-transfer from the trustee, the degree of trustworthiness falls (Falk and Kosfeld, 2006). Relatively high monetary incentives can arouse the agent excessively and the agent may be distracted by the high reward. It is possible that under these conditions, the agent *chokes* and produces a lower level of performance, as, for instance, in a penalty kick taken in a high pressure football game (Ariely et al., 2009a).

When employers are given an opportunity to screen the past history of the workers, relative to a base treatment where such screening is not allowed, one observes important implications for contract choice (Bartling et al., 2012). Employees who develop good reputations by putting in high effort levels in the past are rewarded with good jobs that have high wages and are subject to low levels of control. On the other hand, employees who do not develop such reputations are offered bad jobs (low wages, high control, and low surplus). The possibility of screening spurs workers to work harder in order to build a reputation. Even self-regarding workers mimic other-regarding workers to partake in this reputation building exercise in order to increase their payoffs.

Section 8.6.3 considers the possibility that economic agents may *switch off their ethicality* in the presence of incentives; a phenomenon that psychologists refer to as *moral disengagement*. The moral frame of the dictator game can be changed if it is framed as an exchange game, leading to lower transfers from the dictator (Hoffman et al., 2008). Similarly, when a *competitive environment* is introduced in ultimatum games, proposers make lower offers (Schotter et al., 1996). In field studies of common property resources, regulation produces less cooperative outcomes as compared to communal cooperative efforts (Cardenas et al., 2000).

Fines on parents for arriving late in day care centres in Haifa induced parents to arrive even later. In the absence of fines, a reciprocity relation between carers and parents ensured that they did not arrive as late (Gneezy and Rustichini, 2000a). When students were primed to collect money for a good cause, but not given incentives, they produced greater collections, relative to low incentives. Higher incentives increased collections relative to low incentives, but not relative to the case of no incentives and priming for a good cause (Gneezy and Rustichini, 2000b).

Section 8.6.4 shows that incentives may reduce the ability of people to signal their human virtues, hence reducing prosocial behavior. Titmuss (1971) argued that blood donations are likely to be higher on a voluntary basis as compared to incentivized blood donations. Swedish data on blood donations showed that while there is no crowding-out effect of incentives for men, there is a statistically significant crowding-out effect for women (Mellström and Johannesson, 2008). However, the emerging evidence points to an even more nuanced set of effects. Non-monetary incentives may actually improve blood donations (Lacetera et al., 2012; Goette and Stutzer, 2008). In contrast, several studies find no effect of incentives on blood donations (Reich et al., 2006; Goette et al., 2010).

Monetary incentives given to high school students, preparing for their Bagrut certification in Israel, are ineffective for boys but effective for girls whose predicted grades are higher than

the average (Angrist and Lavy, 2009). When charitable donations are publicly observed, then incentives reduce the signaling value of such donations, reducing the level of donations (Ariely et al., 2009b).

Section 8.6.5 considers the effect of the economic environment. For instance, the legal environment and the strength of law enforcement has an effect on the trustworthiness of agents in a principal–agent relation (Bohnet et al., 2001). *Moral suasion* may be effective in inducing contributions in public goods games, while regulation to contribute a certain minimum level may have a long-run detrimental effect (Reeson and Tisdell, 2008). Bicycle couriers who receive performance related pay are likely to cooperate less in sequential prisoner’s dilemma games. By contrast, those on a fixed hourly wage or members of cooperatives are more likely to cooperate (Burks et al., 2009). This also contributes towards establishing the external validity of lab experiments.

Section 8.6.6 considers situations in which there is synergy between incentives and prosocial behavior. Punishments in public goods games are less likely to be *antisocial punishments* (non-contributors punishing contributors) when punishments for non-cooperation elicit shame and guilt (Herrmann et al. 2008; Carpenter et al., 2009). When sanctions are imposed in public goods games to punish non-contributors, based on a referendum among players (rather than sanctions imposed by the experimenter), they are likely to be more effective. Such sanctions are also more likely to elicit feelings of shame and guilt (Tyran and Feld, 2006).

No significant link is found between intrinsic motivation and incentives in a field study conducted in US schools. However, while incentives for output (final grades) are ineffective, incentives for inputs, e.g., attendance and good behavior, are more effective (Fryer, 2011). The effects of short-term educational incentives are likely to be gender specific and age specific. Further, when incentives are given with delay they are less effective (Levitt et al., 2012). Incentives to exercise over a long term are more effective than short-term incentives (Charness and Gneezy, 2009).

8.2 Moral hazard and other-regarding preferences

In this section, we consider a traditional principal–agent framework in which there is moral hazard and only the agent has other-regarding preferences. We will, in turn, consider the effects of *inequity aversion* (Section 8.2.1) and *reciprocity* (Section 8.2.2) on optimal contracts. Here we restrict attention only to state-contingent wage contracts. The subsequent sections will consider a richer range of contracts such as trust and bonus contracts, as well as issues of asymmetric information.

8.2.1 *Inequity aversion and moral hazard*

Consider the following simple model, due to Itoh (2004), using the standard moral-hazard framework, that should be familiar to undergraduate students in economics. A risk-neutral principal hires a risk-neutral agent to engage in a project that has two possible observable and verifiable outcomes (outputs), $0 \leq y_L < y_H$. The agent can choose one of two unobservable (or observable but non-verifiable) actions $a_0 < a_1$. The cost of the actions is given by $c(a_i)$, $i = 0, 1$. Action a_i results in the outcome y_H with probability p_i , and $p_1 > p_0$. Thus, the higher action induces a higher probability of the higher outcome. The reservation utility of the agent is normalized to zero.

A contract (w_L, w_H) specifies a wage $w_j \geq 0, j = L, H$, contingent on the only observable and verifiable outcome, i.e., y_j . The constraint $w_j \geq 0$ reflects *limited liability*. We rule out renegotiation of contracts.

The principal has self-regarding preferences and his profit is given by

$$\pi_j = y_j - w_j, j = L, H. \quad (8.1)$$

The agent has other-regarding preferences with respect to the payoffs of the principal. His utility function is of the Fehr–Schmidt form and is given by

$$u(w_j, \pi_j, a_i) = \begin{cases} w_j - c(a_i) - \alpha\gamma v(w_j - \pi_j) & \text{if } w_j \geq \pi_j \\ w_j - c(a_i) - \alpha v(\pi_j - w_j) & \text{if } w_j \leq \pi_j \end{cases} ; j = L, H, i = 0, 1, \quad (8.2)$$

where $\alpha > 0, |\gamma| < 1$, and v is an increasing function. Typical, under Fehr–Schmidt preferences, $\beta = \alpha\gamma$ (see (6.1)). If $\gamma > 0$, then the agent suffers from advantageous inequity. But if $\gamma < 0$, then the agent derives satisfaction or pride in coming out ahead; see also the discussion on these issues in Section 6.4. When $w_j = \pi_j$, we have $u(w_j, \pi_j, a_i) = w_j - c(a_i)$, so there are no inequity concerns. Notice that in (8.2), the disutility of effort does not enter into the comparison of the agent's payoff with the principal's payoff; this assumption can be relaxed.

For pedagogical reasons, we make the following two simplifying assumptions.

- A1. The principal always wishes to induce the higher action, a_1 , from the agent (possibly because y_H is high enough).
- A2. The second assumption involves the following relatively innocuous normalization

$$c(a_0) = 0, c(a_1) = c > 0, y_L = 0. \quad (8.3)$$

Let us first consider the textbook case of self-regarding preferences ($\alpha = 0$). The results are summarized in the next proposition.

Proposition 8.24 *Suppose that the agent has self-regarding preferences ($\alpha = 0$).*

- (i) *Let the action of the agent be observable and verifiable. Any combination (w_L, w_H) such that $w_L + p_1(w_H - w_L) = c$, is able to implement the first best contract.*
- (ii) *Under moral hazard, the first best contract does not satisfy the incentive compatibility conditions of the agent.*
- (iii) *Under moral hazard, the unique optimal second best contract specifies $w_L = 0$ and $w_H = \frac{c}{p_1 - p_0}$.*
- (iv) *The outside option of the agent gives a payoff of zero, yet the agent earns positive rents in the second best contract.*

Proof: Define the following first differences.

$$\Delta w = w_H - w_L, \Delta p = p_1 - p_0. \quad (8.4)$$

The expected utility of the agent from the contract (w_L, w_H) , when the agent puts in effort a_1 is $(1 - p_1)w_L + p_1w_H - c$; this can be written as $w_L + p_1\Delta w - c$. The incentive compatibility (IC) constraint and the individual rationality (IR) constraint are given by

$$\begin{cases} IC: & w_L + p_1 \Delta w - c \geq w_L + p_0 \Delta w \\ IR: & w_L + p_1 \Delta w - c \geq 0 \end{cases}$$

Using the notation in (8.4) the IC and IR constraints can be written compactly as

$$\begin{cases} IC: & \Delta p \Delta w \geq c \\ IR: & w_L + p_1 \Delta w \geq c \end{cases}$$

(i) If the action of the agent is observable and verifiable, then the only relevant constraint for the principal is the IR constraint. The principal can directly specify an action and implement it. Leaving any rents to the agent, beyond the agent's outside option of zero, reduces the principal's profits, hence, the IR constraint binds. Thus, any combination (w_L, w_H) such that $w_L + p_1 (w_H - w_L) = c$, is able to implement the first best contract.

(ii) From (8.3), $y_L = 0$, hence, we pick the first best contract $(w_L, w_H) = \left(0, \frac{c}{p_1}\right)$ such that the IR constraint binds. Substitute into the LHS of the IC constraint, we get $\Delta p \Delta w = \Delta p w_H = c \left(1 - \frac{p_0}{p_1}\right) < c$, which violates the IC constraint.

(iii) The problem of the principal under moral hazard is to minimize expected cost, $EC = w_L + p_1 \Delta w$, subject to the IC, IR, and the limited liability constraints. Both IC and IR constraints cannot bind at the same time. For instance, using $(w_L, w_H) = \left(0, \frac{c}{p_1}\right)$, where the IR constraint binds, we get the contradiction in part (ii). In particular, the agent must get some rents under moral hazard from taking the higher action a_1 . Thus the IR constraint does not bind. The IC constraint must bind in equilibrium. If not, then the principal can reduce w_L, w_H by a small amount $\varepsilon > 0$. This does not alter the IC and for a sufficiently small ε , the IR constraint continues to hold. However, the principal's expected costs are reduced. Thus, IC must bind, so, $\Delta p (w_H - w_L) = c$. Thus, the principal's expected cost is $EC = w_L + c \frac{p_1}{\Delta p}$, which is minimized by choosing $w_L = 0$ and choosing $w_H = \frac{c}{\Delta p}$. Thus, the optimal contracts under moral hazard specify $(w_L, w_H) = \left(0, \frac{c}{\Delta p}\right)$.

(iv) Substituting the wage in (iii) contracts on the LHS of the IR constraint, we get $c \frac{p_1}{\Delta p} > c$. Thus, the agent gets strictly positive rents under the optimal contract. ■

Let us now consider optimal contracts when the agent has other-regarding preferences ($\alpha > 0$). We restrict attention to optimal contracts where $w_L = 0$. This restricts the class of optimal contracts, but under certain conditions it is without loss of generality.⁵

Using (8.1), (8.2) the assumptions in (8.3) imply that when $j = L$ we have $w_L = \pi_L = 0$. So, in this case, there are no inequity aversion issues to consider. Thus, inequity aversion has bite only when $j = H$. Using (8.1),

$$w_H \geq \pi_H \Leftrightarrow w_H \geq \frac{y_H}{2} \text{ and } w_H \leq \pi_H \Leftrightarrow w_H \leq \frac{y_H}{2}. \quad (8.5)$$

The incentive compatibility (IC) condition of the agent requires that the agent prefers the high action, a_1 , over the low action, a_0 . Using (8.2) we know that there are two cases to consider:

⁵ See Itoh (2004) for the conditions under which such a contract is optimal. For instance, if there is a limited liability condition on the principal, then there is no loss in generality in assuming $w_L = 0$ because no output is produced in this state of the world (from (8.3), $y_L = 0$).

$w_H \geq \pi_H (\Leftrightarrow w_H \geq \frac{y_H}{2})$ and $w_H \leq \pi_H (\Leftrightarrow w_H \leq \frac{y_H}{2})$. Using (8.1), (8.2), (8.3), (8.5), when $w_H \geq \frac{y_H}{2}$, the IC condition of the agent requires that

$$p_1 (w_H - c - \alpha\gamma v(2w_H - y_H)) + (1 - p_1) (0 - c - 0) \geq p_0 (w_H - \alpha\gamma v(2w_H - y_H)) + (1 - p_0) \times 0. \quad (8.6)$$

When $w_H \leq \frac{y_H}{2}$, the IC condition requires that

$$p_1 (w_H - c - \alpha v(y_H - 2w_H)) + (1 - p_1) (0 - c - 0) \geq p_0 (w_H - \alpha v(y_H - 2w_H)) + (1 - p_0) \times 0. \quad (8.7)$$

Simplifying (8.6), (8.7) we can write the incentive compatibility condition as

$$\begin{cases} w_H - \alpha\gamma v(2w_H - y_H) \geq \frac{c}{p_1 - p_0} & \text{if } w_H \geq \frac{y_H}{2} \quad (IC^+) \\ w_H - \alpha v(y_H - 2w_H) \geq \frac{c}{p_1 - p_0} & \text{if } w_H \leq \frac{y_H}{2} \quad (IC^-) \end{cases}. \quad (8.8)$$

We have denoted the incentive compatibility conditions in the advantageous inequity and the disadvantageous inequity cases, respectively, by IC^+ and IC^- . Given the evidence on ultimatum games, where proposers typically do not offer more than half the share of the pie, we focus here on the case represented by IC^- . The optimal contract in this case, when the agent has other-regarding preferences, is described in the next proposition; readers interested in exploring the implications of the case, IC^+ , may do so as an exercise.

Proposition 8.25 Suppose that $\frac{y_H}{2} \geq \frac{c}{\Delta p}$, where Δp is defined in (8.4).

(i) In the optimal second best contract, $w_L = 0$ and w_H satisfies

$$w_H - \alpha v(y_H - 2w_H) = \frac{c}{\Delta p}.$$

(ii) The expected payment made by the principal in the optimal contract is increasing in α , and strictly increasing in α if $\frac{y_H}{2} > \frac{c}{\Delta p}$.

(iii) In the optimal second best contract, the agent earns positive rents.

Proof: (i) Since we are interested only in the case described by IC^- , let us find the set of values of w_H for which IC^- holds. For $w_H = \frac{y_H}{2}$, IC^- implies that $w_H = \frac{y_H}{2} \geq \frac{c}{\Delta p}$, which is also the assumed condition in the Proposition. The LHS of IC^- is increasing in w_H , so let us find the minimum value of $w_H = \hat{w}_H$ for which IC^- holds; this value is given by the solution to

$$\hat{w}_H - \alpha v(y_H - 2\hat{w}_H) = \frac{c}{\Delta p}. \quad (8.9)$$

Thus, the set of contracts for which IC^- holds is given by the interval $w_H \in [\hat{w}_H, \frac{y_H}{2}]$, where $\frac{y_H}{2} \geq \frac{c}{\Delta p}$; the higher the output in the good state of the world, the wider is the interval of incentive compatible contracts. It is now sufficient to show that $\hat{w}_H$ satisfies the individual rationality constraint of the agent, IR , when the principal wishes to induce the high action, a_1 . The relevant IR is given by

$$\widehat{w}_H - \alpha v(y_H - 2\widehat{w}_H) \geq \frac{c}{p_1}. \quad (8.10)$$

Using (8.9), we can rewrite (8.10) as $\frac{c}{\Delta p} > \frac{c}{p_1}$, which is true, hence, *IR* is satisfied (by assumption, the optimal $w_L = 0$).

(ii) The minimum expected payment of the principal is $p_1 \widehat{w}_H$. It suffices to show that $\widehat{w}_H$ is increasing in α . Differentiating the LHS of (8.9) with respect to α we get $\frac{\partial LHS}{\partial \alpha} = -v(y_H - 2\widehat{w}_H)$, which is strictly negative if $y_H > 2\widehat{w}_H$, thus, the optimal level of $\widehat{w}_H$ must increase to compensate for the fall in the LHS, in order to restore the equality in (8.9). If $\frac{y_H}{2} = \frac{c}{\Delta p}$, then we get the case $\widehat{w}_H = \frac{y_H}{2} = \frac{c}{\Delta p}$, but then $\widehat{w}_H$ is independent of α and $\frac{\partial LHS}{\partial \alpha} = 0$.

(iii) From the last sentence of the proof to (i), $\frac{c}{\Delta p} > \frac{c}{p_1}$, thus, the *IR* holds with a strict inequality and the agent earns positive rents. ■

Comparing Propositions 8.24, 8.25, the agent earns rents in the optimal contract when effort is unobservable, irrespective of whether he has self-regarding or other-regarding preferences. From Proposition 8.25(ii), the other-regarding preferences of the agent hurt the principal. The reason is that when realized output is y_L , there are no equity comparisons. However, when the realized output is y_H , the agent suffers from disadvantageous inequity if the wage in that state, w_H , is low. In order to ensure that the *IC* constraint of the agent binds in that state, the principal ends up offering a higher wage, and reducing his own profits.

In the presence of multiple agents, it may be that the reference group of the agent is not the principal, but rather the other agents. Consider two agents and to each of the agents, the principal offers an independent project (we rule out joint production). Each project has the same technology as in the case of the single agent. Thus, each agent exerts two possible levels of effort, $a_0 < a_1$, and there are two possible probabilistic outcomes for each agent, $0 \leq y_L < y_H$. The higher outcome is more likely for each agent under higher effort, i.e., $p_1 > p_0$. It is possible in this case that $\gamma < 0$, i.e., agents derive greater utility by coming out ahead of the other agent. Let w_{jk}^n be the wage offered to agent $n = 1, 2$ if the outcome of that agent's project is y_j and the outcome of the other agent is y_k such that $j, k = L, H$. For pedagogical simplicity we make the following assumption.

- A3. The agents are symmetric in the sense that they have identical α, γ . The function v is also identical for both agents and is given by $v(x) = x$. The reference group for each agent is the other agent. As in the case of a single agent, we restrict attention to the case in which if an agent gets a low outcome y_L , then his wage is zero, i.e., $w_{LL}^n = w_{LH}^n = 0$ for $n = 1, 2$.⁶ The optimal contract is symmetric, so $w_{jk}^1 = w_{jk}^2 = w_{jk}$, $j, k = L, H$.⁷ The objective of the principal is to minimize the expected payment across both agents.

Remark 8.2 *An immediate implication of assumption A3 is that it suffices to find the optimal values of only w_{HH} and w_{HL} .*

Assume now that assumptions A1, A2, and A3 hold. From assumption A1, the principal wishes to implement the high action from both agents.

The action a_1 , by each of the agents can result in four possible pairs of outcomes (y_H, y_H) , (y_H, y_L) , (y_L, y_H) , and (y_L, y_L) . These outcomes correspond to the respective pairs of wages

⁶ Itoh (2004) shows that this restriction is without loss in generality.

⁷ To be sure, w_{jk}^1 is the wage paid to agent 1 when his output is y_j and the output of agent 2 is y_k . Analogously, w_{jk}^2 is the wage paid to agent 2 when his output is y_j and the output of agent 1 is y_k .

$(w_{HH}^1, w_{HH}^2), (w_{HL}^1, w_{LH}^2), (w_{LH}^1, w_{HL}^2), (w_{LL}^1, w_{LL}^2)$ that occur with the respective probabilities $p_1^2, p_1(1-p_1), (1-p_1)p_1$, and $(1-p_1)^2$.

Suppose that we get the outcome pair (y_j, y_k) , so that the corresponding tuple of wages is (w_{jk}^1, w_{kj}^2) . In this state, the utility of agent 1 from action a_1 is

$$u(w_{jk}^1, w_{kj}^2, a_1) = \begin{cases} w_{jk}^1 - c - \alpha\gamma(w_{jk}^1 - w_{kj}^2) & \text{if } w_{jk}^1 \geq w_{kj}^2 \\ w_{jk}^1 - c - \alpha(w_{kj}^2 - w_{jk}^1) & \text{if } w_{jk}^1 \leq w_{kj}^2 \end{cases}. \quad (8.11)$$

The utility of agent 2 can be written analogously. From assumption A3, the contracts are symmetric, hence, we can dispense with superscripts for players. Using (8.11), we can write the expected utility of any agent from the action, $a_i, i = 0, 1$,

$$\begin{aligned} Eu(W, a_i) &= p_i^2(w_{HH} - c(a_i)) + p_i(1-p_i)[w_{HL} - c(a_i) - \alpha\gamma w_{HL}] \\ &\quad + (1-p_i)p_i[-c(a_i) - \alpha w_{HL}] - c(a_i)(1-p_i)^2. \end{aligned} \quad (8.12)$$

From any agent's perspective, the first and the last terms correspond to the outcomes (w_{HH}, w_{HH}) and (w_{LL}, w_{LL}) that occur with respective probabilities p_i^2 and $(1-p_i)^2$. In each of these states, there are no equity concerns, and by assumption, $w_{LL} = 0$. The middle two terms correspond to the outcomes (w_{HL}, w_{LH}) and (w_{LH}, w_{HL}) that occur with respective probabilities $p_i(1-p_i)$ and $(1-p_i)p_i$. In each case, inequity concerns are important (see (8.11)) and we have assumed that $w_{LH} = 0$. Simplifying (8.12) we get

$$Eu(W, a_i) = p_i[p_i w_{HH} + (1-p_i)w_{HL} - (1-p_i)\alpha w_{HL}(1+\gamma)] - c(a_i). \quad (8.13)$$

Since we would like to implement the high action, a_1 , thus, it follows that the individual rationality constraint of any agent is $Eu(W, a_1) \geq 0$, or

$$p_i[p_i w_{HH} + (1-p_i)w_{HL} - (1-p_i)\alpha w_{HL}(1+\gamma)] - c(a_i) \geq 0. \quad (\text{IR}) \quad (8.14)$$

Using (8.3) and (8.13), the incentive compatibility constraint is $Eu(W, a_1) \geq Eu(W, a_0)$, or

$$p_1 w_{HH} + (1-p_1)w_{HL} + \alpha w_{HL}[p_1 - 1 - (1-p_1)\gamma] \geq \frac{c}{p_1 - p_0}. \quad (\text{IC}) \quad (8.15)$$

From Remark 8.2, the objective of the principal is to choose w_{HH} and w_{HL} in order to minimize his expected cost

$$2p_1(w_{HL} + p_1(w_{HH} - w_{HL})),$$

subject to the IR and the IC of the agents in (8.14) and (8.15).

We introduce some useful terminology about the kinds of contracts that can be offered.

Definition 8.22 In a “team contract,” both agents receive a higher wage if their projects simultaneously succeed relative to the case where at least one project fails, i.e., $w_{HH} > w_{HL}$. In a “relative performance contract,” if one project fails and the other succeeds, the successful agent gets a higher wage relative to the case where both projects succeed, i.e., $w_{HH} < w_{HL}$.

In an “independent contract,” $w_{HH} = w_{HL}$, thus, the success or failure of the project of one agent is irrelevant for the wage of the other agent.

As one may expect, under self-regarding preferences, the independent contract is optimal because the problems of both agents are completely de-coupled; the proof is left for the reader.

Proposition 8.26 *Under self-regarding preferences ($\alpha = 0$), the optimal contract offered by the principal is non-unique. Any pair (w_{HL}, w_{HH}) that satisfies $p_1 w_{HH} + (1 - p_1) w_{HL} = \frac{c}{(p_1 - p_0)}$ is optimal. In particular, the independent contract in which $w_{HH} = w_{HL} = \frac{c}{p_1 - p_0}$ is also optimal.*

The optimal contract under other-regarding preferences is described next; the proof is left for the reader. The intuition behind the result is provided below.

Proposition 8.27 *Under other-regarding preferences ($\alpha > 0$), the optimal contract offered by the principal depends on the parameters of the model in the following manner.*

(i) *If $\gamma > \frac{p_1}{1-p_1}$, then the optimal contract is an extreme “team contract”*

$$w_{HL} = 0 \text{ and } w_{HH} = \frac{c}{p_1(p_1 - p_0)}.$$

The expected payment of the principal is independent of α and γ .

(ii) *Let $\gamma < \frac{p_1}{1-p_1}$ and $\gamma \leq \frac{1}{\alpha} - \frac{L_H}{L_L}$, where $L_H = \frac{p_1}{p_0}$ and $L_L = \frac{1-p_1}{1-p_0}$ are the likelihood ratios of getting the high and the low outcomes under, respectively, the high and the low actions. Then, the optimal contract is an extreme “relative performance evaluation contract,”*

$$w_{HH} = 0 \text{ and } w_{HL} = \frac{c}{(p_1 - p_0)} \frac{1}{(1 - p_1) + \alpha [p_1 - \gamma (1 - p_1)]}.$$

The expected payment of the principal is decreasing in α and increasing in γ .

(iii) *If $\gamma < \frac{p_1}{1-p_1}$ and $\gamma > \frac{1}{\alpha} - \frac{L_H}{L_L}$, then*

$$w_{HL} = \frac{c}{(p_1 - p_0)} \frac{1}{\alpha L_H} \text{ and } w_{HH} = \frac{c}{(p_1 - p_0)} \frac{1 - p_1}{p_1 L_H} \left(\frac{L_H}{L_L} - \frac{1 - \alpha \gamma}{\alpha} \right),$$

and the expected payment from the principal to the agents is increasing in γ , but is independent of α .

We now discuss the results in Proposition 27. Inequity averse agents are motivated by two considerations that do not influence agents with self-regarding preferences. Such agents dislike being behind, as reflected in the size of α , and dislike being ahead if $\gamma > 0$, as reflected in the size of $\alpha \gamma$ (which equals β in the Fehr–Schmidt model). Agents are behind when their action leads to the output y_L , while the other agent’s action leads to the output y_H . Agents are ahead in the converse case. The outcome depends on which of these two influences dominate.

For instance, in Proposition 8.27(i), when γ is large (or $\frac{p_1}{1-p_1}$ is small), then agents derive an even larger disutility from being ahead. Thus, when an agent produces y_H and the other agent produces y_L , his wage is $w_{HL} = 0$. An extreme form of team incentives is preferred to relative performance incentives in this case by paying a positive wage only if both agents are successful in producing y_H . Since there is no inequity in equilibrium, the payments of the principal are independent of α and γ . Thus, extreme team incentives ensure that agents never get left behind each other (provided they both put in the high effort level, a_1).

In Proposition 8.27(ii), we have the opposite case when γ is small. An extreme example of this case arises when $\gamma < 0$, i.e., an agent derives pride from coming out ahead. In this case, the inequity cost to the agent from being ahead is small (or negative if $\gamma < 0$), so, the opposite of team incentives, relative performance evaluation, is efficient.⁸ However, in this case, since positive wages are paid to one of the agents only when that agent produces a high output and the other a low output, issues of inequity aversion will be central to the satisfaction of the incentive compatibility conditions. Hence, the expected payment of the principal does depend on α and γ (and it is increasing in γ , as suggested by the intuition, above).

In general, the inequity aversion of the agents makes the individual rationality constraint of the agents most costly for the principal to satisfy if $\alpha(1 + \gamma) > 1$ (see (8.14)), resulting in larger wage payments relative to the self-regarding case. When $p_1 \rightarrow 1$, the signal (i.e., the realized output) is increasingly informative of the agent's effort. Both conditions, $\gamma < \frac{p_1}{1-p_1}$ and $\gamma \geq \frac{1}{\alpha} - \frac{L_H}{L_L}$, are likely to hold, so we have the case in Proposition 8.27(iii). In this case, a contract that is intermediate between an extreme team incentive contract and an extreme relative performance evaluation contract is offered. Each agent is offered a positive wage whenever his project does not fail. Comparing Propositions 8.26, 8.27, the optimal contract under self-regarding preference is non-unique but the introduction of even a small amount of other-regarding preferences ($\alpha > 0$) gives rise to a unique contract.⁹

Rey-Biel (2008) also considers inequity aversion in a situation where a principal hires two agents. The agents are inequity averse towards each other but not towards the principal. The main difference from the model in Itoh (2004) is that there is *joint production* of a single good between the two agents and there is no uncertainty. Hence, the outcome provides a noiseless signal of the effort levels of the agents. The principal who has self-regarding preferences exploits the inequity aversion of the two agents in designing the optimal wage contracts in the following manner.

The optimal wage levels are high and equal when both agents work hard (so a high output level is produced), thus, there is minimal inequity aversion. However, when the outcome signals with certainty that one of the agents has not worked hard but the other has, then the wage levels are used to deliberately create high inequality among the agents. Since the agents dislike inequality, it is easier for their incentive compatibility conditions to be satisfied in the presence of inequity aversion. Relative to purely self-regarding agents, the principal is able to induce a higher effort level at a lower wage cost when agents are inequity averse.

8.2.2 Reciprocity and moral hazard

Having considered the effects of *inequity aversion* in the presence of moral hazard, we now consider the effect of *reciprocity*. We follow the model of Englmaier and Leider (2012) but restrict attention to two actions, two outcomes, a risk-neutral principal and a risk averse agent. The rigorous modeling of these issues requires the use of psychological game theory; see Part 4 of the book. By contrast, Englmaier and Leider (2012) take a simpler, direct, and reduced-form approach.

The model is similar to that in Section 8.2.1. The agent chooses among two unobservable actions $a_0 < a_1$. The cost of action $i = 0, 1$ is c_i . Each action probabilistically determines one of two observable output levels, $0 \leq y_L < y_H$. Action a_i results in the outcome y_H with probability

⁸ Further issues of the optimality of relative performance evaluation under inequity averse agents are explored in Bartling (2011).

⁹ For a consideration of issues similar to Itoh (2004) see Bartling and von Siemens (2010) and, in a continuous outcome framework, Englmaier and Wambach (2010).

p_i such that $p_1 > p_0$. The utility function of the agent, defined over material rewards, u , is increasing and concave (we did not invoke risk aversion in Section 8.2.1). The outside option of the agent is zero. The *monotone likelihood ratio property* (MLRP) holds, i.e., $\frac{p_1}{p_0} > \frac{1-p_1}{1-p_0}$, so the relative probability of getting a higher signal under a higher action is greater.

A contract specifies, for $j = L, H$, a triple (w_L, w_H, a) , where $w_j \geq 0$ is the wage if the realization of output is y_j and $a \in \{a_0, a_1\}$ is the *non-contractible, desired level of effort*. The agent has limited liability ($w_j \geq 0$) and there is no renegotiation of contracts. We assume that the principal always wishes to induce the higher action, a_1 . Thus, we wish to focus attention on contracts of the form (w_L, w_H, a_1) in which the principal optimally chooses w_L, w_H . Let $u_H = u(w_H)$ and $u_L = u(w_L)$ and define the following

$$\bar{u}(a_i) = p_i u_H + (1 - p_i) u_L - c_i, \bar{\pi}(a_i) = p_i y_H + (1 - p_i) y_L, \quad (8.16)$$

where $\bar{u}(a_i)$ and $\bar{\pi}(a_i)$ are, respectively, the expected utility of the agent (net of effort costs) and the expected output of the principal when the agent chooses the action a_i , $i = 0, 1$. Conditional on any level of expected wage, the expected output, $\bar{\pi}$, is directly proportional to the principal's profit. Thus, when considering the effort choice of the agent, it is pedagogically convenient to refer to $\bar{\pi}$ as the expected profit of the principal. Define the following first differences

$$\Delta c = c_1 - c_0, \Delta \bar{\pi} = \bar{\pi}(a_1) - \bar{\pi}(a_0).$$

Consider the contract (w_L, w_H, a_i) , in which the desired, but non-contractible, effort is a_i . The utility of the agent from this contract when he chooses the action $a_j \in \{a_0, a_1\}$ is denoted by $U(a_j, a_i)$ and is given by

$$U(a_j, a_i) = \bar{u}(a_j) + r \bar{u}(a_i) \bar{\pi}(a_j), \quad r > 0. \quad (8.17)$$

Conditional on the contract (w_L, w_H, a_i) , the first term on the RHS is the expected utility of the agent when he chooses the action a_j . The second term captures considerations of reciprocity. In judging the kindness of the principal, the agent takes the outside option, normalized to zero, as the fair level of utility. But if $\bar{u}(a_i) > 0$, the agent derives positive rents from the contract, so he reciprocates by engaging in gift exchange. This is reflected in the interaction of the agent's rents from the intended contract with the expected profit of the principal from the action chosen by the agent, a_j . The parameter $r > 0$ reflects the relative weight placed on reciprocity considerations. The greater is r and the higher the level of rents, the greater is the level of internalization of the principal's expected profit by the agent.

We assume that the principal always finds it optimal to implement the action a_1 . For the principal, maximizing expected profits $p_1(y_H - w_H) + (1 - p_1)(y_L - w_L)$ is equivalent to minimizing expected wage costs. Thus, the principal's problem, P0, is

$$w_L^*, w_H^* \in \arg \min_{w_j \in [0, y_j], j=H,L} \bar{w}(a_1) = p_1 w_H + (1 - p_1) w_L \quad (8.18)$$

subject to

$$U(a_1, a_1) \geq U(a_0, a_1) \quad (\text{IC})$$

$$U(a_1, a_1) \geq 0. \quad (\text{IR})$$

The optimal wage profile w_L^*, w_H^* must induce the choice of a high effort level, a_1 (incentive compatibility constraint, IC) and provide the agent with at least his reservation utility of zero (individual rationality constraint, IR).

The main difference from the standard moral hazard problem is that the principal can induce a high effort level through two channels. (1) By optimally offering a mix of w_L, w_H , as in models of self-regarding preferences. (2) By offering rents to the worker, which tap into the worker's reciprocity and internalize the objective of producing high expected output for the principal.

Consider the transformation introduced by Grossman and Hart (1983). Let

$$w_j = h(u(w_j)), h = u^{-1}, \text{ so } h' > 0 \text{ and } h'' > 0, j = H, L. \quad (8.19)$$

The implications in (8.19) are drawn from the standard properties of inverse functions and the fact that $u' > 0$ and $u'' < 0$.¹⁰ Using (8.17), (8.19), and the notation $u_j = u(w_j), j = L, H$, we can rewrite the principal's problem, P0. The new, but equivalent, problem, P1, eliminates the wage levels and is expressed in terms of utilities of the agent in the two states $j = L, H$:

$$u_L^*, u_H^* \in \arg \min V = p_1 h(u_H) + (1 - p_1) h(u_L)$$

subject to

$$\bar{u}(a_1) - \bar{u}(a_0) + r\bar{u}(a_1)(\bar{\pi}(a_1) - \bar{\pi}(a_0)) \geq 0 \quad (\text{IC})$$

$$\bar{u}(a_1) + r\bar{u}(a_1)\bar{\pi}(a_1) \geq 0. \quad (\text{IR})$$

The optimal choices $u_j^*, j = H, L$ are restricted to the compact interval $[u(0), u(y_j)]$. In problem P1, in order to induce the high effort, a_1 , the principal minimizes a convex function (recall $h'' > 0$) subject to linear constraints. Thus, the satisfaction of the first order conditions with the relevant complementary slackness conditions is necessary and sufficient for an optimal solution.

We may now speak of the optimal contract as the triple (u_L^*, u_H^*, a_1) , which allows us to recover the optimal wage contract (w_L^*, w_H^*, a_1) through the relation $w_j^* = h(u_j^*)$.

Let the Lagrangian of the constrained minimization problem, P1, be denoted by L and let λ_{IC} and λ_{IR} be the Lagrangian multipliers on the IC and IR constraints. The first order conditions with respect to u_L and u_H are

$$u_L : h'(u_L) = \lambda_{IC} \left[1 - \frac{1-p_0}{1-p_1} + r\Delta\bar{\pi} \right] + \lambda_{IR} (1 + r\bar{\pi}(a_1)). \quad (8.20)$$

$$u_H : h'(u_H) = \lambda_{IC} \left[1 - \frac{p_0}{p_1} + r\Delta\bar{\pi} \right] + \lambda_{IR} (1 + r\bar{\pi}(a_1)). \quad (8.21)$$

In addition, we have the Kuhn–Tucker complementary slackness conditions, i.e.,

$$\lambda_{IC} \frac{\partial L}{\partial \lambda_{IC}} = 0, \lambda_{IR} \frac{\partial L}{\partial \lambda_{IR}} = 0 \quad (8.22)$$

¹⁰ For any feasible point, x , in the domain of h , $h'(x) = \frac{1}{u'(h(x))}$ and $h''(x) = -\frac{1}{(u'(h(x)))^2} u''(h(x))h'(x)$. Thus, if u is increasing and concave, then the inverse, h , is increasing and convex.

Denote the minimum value of the objective function of the principal by

$$V(r, \Delta\pi) = p_1 h(u_H^*) + (1 - p_1) h(u_L^*) \quad (8.23)$$

We first state the benchmark textbook result in the absence of reciprocity ($r = 0$).

Proposition 8.28 (*Optimal contracts in the absence of reciprocity*): Suppose that reciprocity is absent ($r = 0$). Then, in the optimal contract, (u_L^*, u_H^*, a_1) :

(i) The first order conditions are

$$h'(u_L) = \lambda_{IC} \left[1 - \frac{1 - p_0}{1 - p_1} \right] + \lambda_{IR}, \quad (8.24)$$

$$h'(u_H) = \lambda_{IC} \left[1 - \frac{p_0}{p_1} \right] + \lambda_{IR}, \quad (8.25)$$

which, in conjunction with the complementary slackness conditions, (8.22), give the optimal solution.

(ii) The IR constraint binds.

(iii) The principal rewards the agent more for outcomes that are relatively more likely under the desired action, a_1 .

Proof: (i) The first order conditions follow by setting $r = 0$ in (8.20), (8.21). This gives the two equations (8.24), (8.25). Since we have a convex minimization problem subject to linear inequality constraints, (8.24), (8.25), in conjunction with (8.22), give the solution to the optimal contract.

(ii) Suppose that the IR constraint does not bind at the optimal contract (u_L^*, u_H^*, a_1) . Then consider a new contract $(u_L^* - \epsilon, u_H^* - \epsilon, a_1)$, where $\epsilon > 0$ is a small number. The IC and IR constraints continue to hold (because $r = 0$) and the expected payment to the agent falls, increasing the expected profits of the principal. Hence, the original contract (u_L^*, u_H^*, a_1) at which the IR constraint does not bind is not optimal.

(iii) Suppose $\lambda_{IC} = 0$, so from (8.24), (8.25), $h'(u_L) = h'(u_H) = \lambda_{IR}$. In this case, a flat wage $w_L = w_H$ is optimal. When there are no problems of incentives, then it is optimal to fully insure the risk averse agent. But this creates a problem of incentives in implementing a high effort level. To see this, a flat wage gives $u_L = u_H = \hat{u}$. Then, using (8.16) we have $\bar{u}(a_1) = \hat{u} - c_1$ and $\bar{u}(a_0) = \hat{u} - c_0$. From the binding IR constraint we have $\hat{u} = c_1$. The IC constraint then implies the contradiction $0 \geq \Delta c$. Thus, a flat wage is not optimal.

Since the Lagrangian multipliers are positive in the presence of inequality constraints, the other possibility is $\lambda_{IC} > 0$, so the IC constraint must bind (from (8.22)). Notice that in each of (8.24), (8.25), since $h'' > 0$, an increase in the RHS increases the corresponding wage levels, w_L, w_H . From the MLRP, stated above, we know that $\frac{p_1}{p_0} > \frac{1-p_1}{1-p_0}$, hence, $w_L < w_H$. Thus, the principal rewards the agent for observing an outcome that is more likely under the desired action, a_1 . ■

Let us now see how the results are modified in the presence of reciprocity ($r > 0$).

Proposition 8.29 Let $r > 0$. Then the following hold.

(i) If an agent with the reciprocity parameter $r = r_1$ is induced to choose action a_1 by the contract (u_L, u_H, a_1) , then any agent with a reciprocity parameter $r = r_2 > r_1$ also chooses the action a_1 .

- (ii) The minimum expected cost of the principal, $V(r, \Delta\bar{\pi})$, defined in (8.23), is non-increasing in r and $\Delta\bar{\pi}$.
- (iii) When the rent of the agent is high enough, a flat wage schedule can implement the action a_1 . As the reciprocity parameter $r \rightarrow \infty$, the action a_1 can be implemented by a flat wage schedule with infinitesimal rents.

Proof: (i) Differentiating the LHS of the IC constraint with respect to r , we get $\bar{u}(a_1)(\bar{\pi}(a_1) - \bar{\pi}(a_0)) > 0$. Similarly, differentiating the LHS of the IR constraint with respect to r we get $\bar{u}(a_1)\bar{\pi}(a_1) > 0$. Thus, an increase in r relaxes both constraints. Hence, if the optimal contract can be implemented at r_1 it can be implemented at $r_2 > r_1$.

(ii) From part (i), an increase in r relaxes the IC constraint, and the IR constraint continues to hold. Hence, the expected profits of the principal can be no lower as r increases. In particular, if IR does not bind, then the expected profits of the principal will strictly increase as r increases, otherwise they are unchanged (but no lower). An increase in $\Delta\bar{\pi}$ relaxes the IC and IR constraints, hence, the expected profits of the principal are also increasing in $\Delta\bar{\pi}$.

(iii) Suppose that there is a flat wage profile, so $w_L = w_H$ but the contract leaves positive rents for the agent. Constant state-contingent wages imply that $u_L = u_H = \tilde{u} > 0$. The outside option of the agent is zero, thus, the IR constraint is

$$(\tilde{u} - c_1)(1 + r\bar{\pi}(a_1)) > 0.$$

Evaluating the IC constraint at the stated contract, we get

$$r(\tilde{u} - c_1)\Delta\bar{\pi} > \Delta c.$$

Thus, for high enough $\tilde{u}$, the IC constraint is satisfied. Clearly, the larger the reciprocity parameter r , the smaller is the level of rents, $\tilde{u}$, required to satisfy the IC constraint. ■

Comparing Propositions 8.28, 8.29 we see the following. In the absence of reciprocity, there is a conflict between insurance and incentives, so a flat wage is good for insurance but bad for incentives. However, in the presence of reciprocity, both insurance and incentives point in the same direction, hence, this may give rise to an optimal contract with flat wages. Indeed, the prediction of high-powered incentives in the absence of reciprocity (Proposition 8.28) is often violated in practice. See, for instance, the quote, due to Lemieux et al. (2009) in the introduction.

Lemma 8.9 *In the presence of reciprocity, the IC constraint binds for all $r > 0$.*

Proof: First suppose that both constraints, IC and IR, are slack when evaluated at the contract (u_L, u_H, a_1) . Then the principal's profits can be increased by the contract $(u_L - \epsilon, u_H - \epsilon, a_1)$, hence, this case cannot arise. The other possibility is that IC is slack but IR binds. Since IR binds, the agent gets no rents, hence, reciprocity of the agent become irrelevant. In effect, the behavior of the agent in this case is like a non-reciprocal agent ($r = 0$). If the IC is slack, so $\lambda_{IC} = 0$, then from (8.20), (8.21) we have $h'(u_L) = h'(u_H)$, so $w_L = w_H$, i.e., a flat wage schedule. But we know in this case that the self-regarding agent does not put in any effort, thus IC should be violated, contradicting the initial supposition that it does hold. Thus, the IC should bind when $r > 0$. ■

Lemma 8.9 shows that IC will always bind, so we have two additional possibilities. (i) IR binds. (ii) IR is slack. The next proposition indicates when each of these cases is likely to arise.

Proposition 8.30 *There exists some finite reciprocity parameter r^* such that (i) for all $r \leq r^*$, IR binds, and (ii) for all $r > r^*$, the IR constraint is slack and the optimal contract provides strictly positive rents to the agent.*

We leave the proof of Proposition 8.30 as an exercise for the reader. In case (ii) of Proposition 8.30, since IR is slack we have $\lambda_{IR} = 0$. Substituting $\lambda_{IR} = 0$ in (8.20), (8.21) and using the fact that $h'(u_L) = \frac{1}{u'(w_L)}$ (see footnote 10 above on the properties of an inverse function) we have that the optimal contract specifies

$$\frac{u'(w_H)}{u'(w_L)} = \frac{1 - \frac{1-p_0}{1-p_1} + r\Delta\bar{\pi}}{1 - \frac{p_0}{p_1} + r\Delta\bar{\pi}}. \quad (8.26)$$

Thus, the ratio of marginal utilities depends not just on the likelihood ratios but also on the extent of reciprocal feelings of the agent. We leave it to the reader to explore the effects of the degree of informativeness of signals on the optimal contract. Englmaier and Leider (2012) show that when signals are more informative about the effort choice of the agent, then the principal relies more on incentives (as compared to reciprocity) in order to induce the desired effort level. One testable prediction of this model, therefore, is to compare the incidence of incentive contracts in firms with different levels of noise in the production technology. These sorts of models are also potentially useful to study the optimality of different organizational forms.¹¹

8.3 Incomplete contracts under other-regarding preferences

Most real-life contracts are incomplete. In particular, in principal–agent relations, the effort level of agents is difficult to specify in full detail. High-powered incentive contracts predict high levels of effort when agents have self-regarding preferences. However, real life incentives are typically not high-powered, and apart from top management, fixed wage payments or mildly increasing wage schemes are quite common.

In this section, we consider a simple model of moral hazard in a principal–agent relation but with (1) a rich menu of contracts, (2) the presences of other-regarding preferences, and (3) uncertainty about preferences (self-regarding or other-regarding). We show that in the presence of other-regarding preferences, a range of contracts, such as *bonus contracts* and *trust contracts* become potentially profitable for a principal; indeed, even more profitable than purely *incentive contracts*. Under purely self-regarding preferences, the bonus and the trust contracts are shown to be dominated by incentive contracts. The converse may be true under other-regarding preferences.

The treatment in this section is based on Fehr et al. (2007). We give a detailed treatment for two reasons. First, understanding the basic theoretical framework is important in the economics of incentives. Second, the predictions of this model have been empirically tested; we consider the empirical results in Section 8.3.3, below.

¹¹ Englmaier and Leider (2012) make progress in this direction by identifying three different organizational forms in which there are two managers of a firm. They consider various heirarchical relations between them.

Consider a single, risk-neutral, principal who contracts a single, risk-neutral, agent to work for him. The reservation utilities of both parties are normalized at zero. The agent can expend effort $e \in [\underline{e}, \bar{e}]$ at some private cost $c(e)$ such that $c(\underline{e}) = 0$, $c' > 0$ and $c'' > 0$. The utility of the agent is given by

$$u = w - c(e) - C_A, \quad (8.27)$$

where C_A are other possible costs associated with a contract that are incurred by the agent, such as fines for non-fulfillment of contractual obligations; these are specified in more detail below. w is the wage received by the agent from the principal.

The output of the principal is given by $\theta v(e)$, where $\theta > 0$ is a constant productivity parameter, and $v(\underline{e}) = 0$, $v'(e) > 0$, $v''(e) < 0$. We also assume that

$$c'(\underline{e}) < \theta v'(\underline{e}), \quad (8.28)$$

i.e., the marginal benefit of the lowest effort level to the principal, exceeds the marginal cost.

Let e_c be some effort level that the principal desires to induce. The principal can write e_c as part of the contract, in which case we call it the contracted effort level, but when the effort level is non-verifiable, then e_c is non-binding. Effort is observed by both parties. However, the principal cannot provide hard evidence of the effort level to a third party unless he invests in a verification technology that costs $k > 0$. The verification technology unearths hard evidence of the actual effort level with some probability $p \in (0, 1)$. Based on this evidence, the principal can impose a fine $f > 0$ on the agent. The maximum possible fine is $\bar{f}$. We also make a technical assumption

$$c(e_c) - c(\underline{e}) \leq p\bar{f}, \forall e_c. \quad (8.29)$$

The profits of the principal are given by

$$\pi = \theta v(e) - w - C_P, \quad (8.30)$$

where C_P captures other possible costs (positive or negative) associated with the contract that are incurred by the principal, such as investment in a verification technology, or receipt of fines from the agent.

The sequence of moves in the multi-stage game is as follows. In Stage 1, the principal makes a *take it or leave it* offer to the agent after choosing (a) the type of contract to offer, and (b) whether to invest in the verification technology. In Stage 2, the agent decides (a) to accept or reject the contract, and (b) the effort level to expend. In Stage 3, if the initial investment in the verification technology has been made, then hard evidence of the agent's effort level is obtained with probability $p > 0$ and the terms of the initial contract are enforced.

Consider the following three types of contracts.

1. *Incentive contracts*: This is the typical form of contracts in classical contract theory. Assume that the principal invests in the monitoring technology, so hard evidence of the agent's effort level can be produced with probability $p > 0$ and the agent expends an effort level e . The contract specifies a triple (w, e_c, f) , where e_c is the contracted effort level and f is the

fine to be paid if verifiable evidence of shirking (i.e., $e < e_c$) is discovered. The ex-post utilities of the agent and the principal, u, π , are given respectively by:

$$u = \begin{cases} w - c(e_c) & \text{if } e = e_c \\ w - c(e) & \text{if } e < e_c \text{ (with probability } 1 - p) \\ w - c(e) - f & \text{if } e < e_c \text{ (with probability } p) \end{cases} \quad (8.31)$$

$$\pi = \begin{cases} \theta v(e_c) - w - k & \text{if } e = e_c \\ \theta v(e) - w - k & \text{if } e < e_c \text{ (with probability } 1 - p) \\ \theta v(e) - w - k + f & \text{if } e < e_c \text{ (with probability } p) \end{cases} \quad (8.32)$$

2. *Trust contract*: The trust contract specifies a pair (w, e_c) and there is no investment in the monitoring technology. Thus, shirking by the agent is never discovered and no fines are levied either. The contractual effort, e_c , is, therefore, non-binding. The wage payment $w \geq c(e)$ to the agent is not conditioned on the actual effort level of the agent. The contractual wage payment can be enforced by a court if necessary, so the agent is guaranteed to receive it, irrespective of the actual effort level, e . The respective utilities of the agent and the principal in a trust contract are:

$$u = w - c(e). \quad (8.33)$$

$$\pi = \theta v(e) - w. \quad (8.34)$$

Since e_c is unenforceable, the trust contract is more incomplete than the incentive contract.

3. *Bonus contract*: This contract is similar to a trust contract except for the presence of a non-enforceable bonus payment, $b \geq 0$, from the principal to the agent at Stage 3, conditional on having observed the Stage 2 outcome. Hence, the contract can be represented by a triple (w, e_c, b) . As in the case of the trust contract, e_c is not enforceable. The respective, ex-post, utilities of the agent and the principal are:

$$u = \begin{cases} w - c(e) & \text{when no bonus is paid} \\ w - c(e) + b & \text{when bonus is paid} \end{cases} \quad (8.35)$$

$$\pi = \begin{cases} \theta v(e) - w & \text{when no bonus is paid} \\ \theta v(e) - w - b & \text{when bonus is paid} \end{cases} \quad (8.36)$$

The bonus contract is the most incomplete of the three contracts because b and e_c are both unenforceable.

We now make an important assumption that concerns the nature of reciprocity. We assume that this assumption holds in the subsequent analysis in this section. For the empirical evidence behind this assumption, see Section 8.6.

- A1. When faced with an incentive contract, all types of principals and agents (other-regarding and self-regarding) behave selfishly. *This is the sense in which the model captures concerns for reciprocity.* Incentive contracts are perceived as unkind behavior by the principal and induce agents to behave in a self-regarding manner. Trust and bonus contracts (see below) are perceived by the agent as kind behavior by the agent and induce other-regarding preferences (unless the agent is inherently self-regarding). Furthermore, in a bonus contract, if the principal exhibits reciprocity, then the agent may put in high effort in the hope of receiving an, ex-post, bonus payment.

8.3.1 Optimal contracts under self-regarding preferences

If agents have self-regarding preferences, then reciprocity-based contracts, such as trust and bonus contracts, invite one-sided opportunism (trust contracts) or two-sided opportunism (bonus contracts). Self-regarding agents exert zero effort level because they are neither influenced by other-regarding preferences, nor by reciprocity. This contrasts with the observed evidence, say, in the gift exchange game.

Incentive contracts can, however, be very effective in eliciting the contracted effort, provided that the monitoring technology is effective (high p) and the maximum available fine, $\bar{f}$, is high enough. Proposition 8.31, below, formalizes these ideas and also characterizes the first best contract when both contractual parties have self-regarding preferences. The proofs in this section should be standard for most undergraduate students in economics, so they are all relegated as exercises.

Proposition 8.31 *Suppose that the principal and the agent have self-regarding preferences.*

- (i) *Consider the first best contract that maximizes the total surplus $u + \pi = \theta v(e) - c(e)$. The socially optimal effort level, e^* , is given by*

$$\theta v'(e^*) = c'(e^*), \quad (8.37)$$

and $u + \pi > 0$.

- (ii) *In the trust and the bonus contracts, the solution is summarized by*

$$e = \underline{e}, w = b = 0, u = 0, \pi = 0, u + \pi = 0. \quad (8.38)$$

- (iii) *Suppose that the principal invests in the monitoring technology. If $p\bar{f} \geq c(e^*) - c(\underline{e})$, then the solution under the incentive contract is summarized by*

$$e = e^* > \underline{e}; u = 0; \pi = \theta v(e^*) - c(e^*) - k \geq 0; u + \pi \geq 0. \quad (8.39)$$

If the deterrence technology is potent enough ($p\bar{f} \geq c(e^*) - c(\underline{e})$), the equilibrium effort under the incentive contract is identical to the first best contract. However, the aggregate payoff is lower, on account of the cost of the verification technology, $k > 0$. Hence, this is the second best outcome. The trust and the bonus contracts elicit the lowest effort levels $\underline{e}$; since $v(\underline{e}) = 0$, the firm's profits are zero in both cases.

The testable prediction of the model with self-regarding preferences, therefore, is that the principal will choose the incentive contract over the trust and the bonus contracts, when all three contractual forms are available; the evidence, described in Section 8.3.3, rejects this prediction.

8.3.2 Optimal contracts in the presence of other-regarding preferences

Now suppose that the principal and the agent have other-regarding preferences that take the Fehr–Schmidt (1997) form (abbreviated as FS preferences). Given assumption A1, individuals with FS preferences exhibit other-regarding preferences in the case of bonus and trust contracts but behave like self-regarding agents in the case of incentive contracts.

If all this is commonly known, then, in the trust contract, the principal could rationally hope that higher wages will be reciprocated by higher effort (as in the gift exchange game). Analogously, in a bonus contract, the agent rationally hopes the principal will reciprocate a high effort with a bonus payment. These factors are conducive for the agent to put in a high effort. Reciprocity

arises here on account of inequity aversion and not beliefs, which is the case when we apply the machinery of psychological game theory (see Part 4).

For pedagogical simplicity, we make the following assumptions.

$$v(e) = e; \quad c(e) = \frac{1}{2}e^2; \quad e \in [0, 1]. \quad (8.40)$$

$$\theta = 1. \quad (8.41)$$

Equation (8.40) restricts the marginal benefit of effort to be linear, the effort cost function to be quadratic, and the effort choice to be made from the unit interval. Given these assumptions, the marginal product of effort, $\theta v'(e) = 1$, is strictly greater than the marginal cost of effort, $c'(e) = e$, for all $e \in [0, 1)$ and equal at $e = 1$. Thus, from (8.37), a social planner would always want to elicit $e^* = 1$.

Under these specific functional forms, the utility of a self-regarding agent and a self-regarding principal is given, respectively, by

$$u = w - \frac{1}{2}e^2 - C_A \text{ and } \pi = e - w - C_P. \quad (8.42)$$

If the monetary payoffs of the agent and the principal are given by u and π as in (8.42) (which is a special case of (8.27), (8.30)), then their respective FS preferences, U and Π are given by:

$$U(u, \pi) = u - \alpha_A \max\{\pi - u, 0\} - \beta_A \max\{u - \pi, 0\} \quad (8.43)$$

$$\Pi(u, \pi) = \pi - \alpha_P \max\{u - \pi, 0\} - \beta_P \max\{\pi - u, 0\}, \quad (8.44)$$

where $\beta_i \in [0, 1]$, $\beta_i \leq \alpha_i$, $i = A, P$.

Let $0 < \lambda \leq 1$ denote the probability that the agent has self-regarding preferences, i.e., $\alpha_i = \beta_i = 0$. With probability $1 - \lambda$, the agent has other-regarding preferences. At the time of designing the incentive scheme of the agent in Stage 1, there is uncertainty: The principal does not know the agent's type, but knows the distribution of types. The principal cannot condition wage payments on the type of the agent; perhaps because the type information is non-verifiable.

In line with the experimental evidence, we make the following assumption that applies to the subsequent analysis in this section.

A2. The preference parameters obey the restrictions, $\alpha_i \geq \beta_i > 0.5$.¹²

Due to uncertainty about the type of the agent, the profit of the principal, π , is a random variable. Theory does not provide sufficient guidance on how to extend (8.44) in the presence of risk. We make the simplest pedagogical assumption that the principal maximizes his expected payoffs $E\P$,

$$E\P = \lambda \Pi(u, \pi \mid \text{self-regarding agent}) + (1 - \lambda) \Pi(u, \pi \mid \text{other-regarding agent}), \quad (8.45)$$

where the first and second terms on the RHS in (8.45) condition the profits of the principal on the type of the agent, self-regarding or other-regarding.

Suppose that the agent is other-regarding. It can then be shown that other-regarding agents are: (i) willing to transfer money to the other party to reduce disadvantageous inequality, and

¹² The assumption $\alpha_i > 0.5$ has better support than the assumption $\beta_i > 0.5$; see Table 6.1, above. However, the assumption $\beta_i > 0.5$ affords pedagogical simplicity.

(ii) engage in costly punishments, provided that punishments are not too costly and the envy parameter α is high enough. This is formalized in Lemma 8.10.

Lemma 8.10 *Suppose that $\alpha_A \geq \beta_A > 0.5$ (assumption A2) and the agent is other-regarding. Consider the Stage 2 problem of the agent to choose an effort level.*

- (1) *If $u > \pi$, the agent is willing to reduce inequality by transferring resources to the principal.*
- (2) *Suppose that the agent can, at a cost $\gamma < 1$, reduce the payoff of the principal by a unit. If $u < \pi$, the agent is willing to engage in costly punishment in order to reduce inequality if $\gamma \leq \frac{\alpha_A}{1+\alpha_A} < 1$.*

Proof: Let x_A and x_P be, respectively, the monetary payoffs of the agent and the principal. There are two cases to consider.

(a) $x_A > x_P$. In this case, the utility of the agent is given by $U(x_A, x_P) = x_A - \beta_A(x_A - x_P)$. Suppose now that the agent were to transfer a small amount $\varepsilon > 0$ of his own income to the principal, then his utility is

$$U(x_A, x_P) = x_A - \beta_A(x_A - x_P) + \varepsilon(2\beta_A - 1).$$

Since $\beta_i > 0.5$ it follows that the agent is better off by such a transfer of income. Hence, it follows that when inequality is to their advantage, other-regarding agents are willing to spend resources to benefit the disadvantaged partner.

(b) $x_A < x_P$. In this case, the utility of the agent is $U(x_A, x_P) = x_A - \alpha_A(x_P - x_A)$. Suppose that the agent can, at a cost $\gamma < 1$, reduce the payoff of the principal by a unit (it never makes sense to do so if $\gamma \geq 1$). The payoff of the agent in this case is

$$\begin{aligned} U(x_A, x_P) &= (x_A - \gamma) - \alpha_A[(x_P - 1) - (x_A - \gamma)] \\ \Leftrightarrow U(x_A, x_P) &= x_A - \alpha_A(x_P - x_A) + \alpha_A(1 - \gamma) - \gamma. \end{aligned} \quad (8.46)$$

Thus, the agent will prefer to punish the principal to reduce inequality if $\alpha_A(1 - \gamma) - \gamma \geq 0$ or equivalently if

$$\gamma \leq \frac{\alpha_A}{1 + \alpha_A} < 1. \quad (8.47)$$

This completes the proof. ■

TRUST CONTRACTS

From Proposition 8.31, under a trust contract, the self-regarding agent puts in the effort level $e = 0$ (this result holds in the special case of (8.40)). We now consider the effort choice of an other-regarding agent in Stage 2.

Lemma 8.11 *(a) Consider an other-regarding agent who accepts a trust contract. Then, in any “incentive compatible” trust contract, the agent will choose an effort level such that the payoffs of the principal and the agent are equated.*

The locus of (w, e) points that satisfy the incentive compatible contract are given by the condition

$$g(w, e) = e^2 + 2e - 4w = 0. \quad (8.48)$$

The effort reaction function of the other-regarding agent under an incentive compatible trust contract is given by¹³

$$e^T(w) = \sqrt{4w + 1} - 1 \geq 0. \quad (8.49)$$

(b) A self-regarding agent chooses $e = 0$, irrespective of the wage paid.

From (8.49), the effort level of the other-regarding agent is increasing in the wage paid, as in gift exchange experiments and models of efficiency wage. This arises on account of the other-regarding preferences of the agent (see Lemma 8.10).

Lemma 8.12 *In equilibrium, in a trust contract, the individual rationality constraints of the self-regarding and the other-regarding agents are respectively given by $w \geq 0$ and $w - \frac{1}{2}e^2 \geq 0$, where e is some incentive compatible effort level.*

Proof of Lemma 8.12: We know from Lemma 8.11 that in a trust contract, an other-regarding agent chooses effort $e^T(w) \geq 0$ and the effort choice of the agent equates the payoff with the principal. Using (8.43), since the outside option of the agent is zero, the principal must offer a wage such that $w - \frac{1}{2}e^2 \geq 0$. From Lemma 8.11 the self-regarding agent puts in a zero effort level, hence, he accepts the contract if $w \geq 0$. ■

We now state the next result that is simple to verify.

Lemma 8.13 *For a fixed $0 < e \leq 1$, the wage, w , that satisfies (8.48), the incentive compatibility condition of an other-regarding agent, is higher than the wage that satisfies the individual rationality constraint at equality for any strictly positive effort level, i.e., $w = \frac{1}{2}e^2$. When $e = 0$, the wage level that satisfies these two constraints is identical at $w = 0$. Hence, the satisfaction of the incentive compatibility constraint, for any strictly positive effort level by an other-regarding agent, will leave some rents to the agent.*

The result in Lemma 8.13 is analogous to the result in Proposition 8.30. Inequity averse agents repay the principal in terms of higher effort if the principal leaves them with rents.

We have considered, above, the Stage 2 problem for the case of an other-regarding and a self-regarding agent. We now consider the Stage 1 problem of the principal, separately, for the cases of other-regarding and self-regarding principals.

I. Trust contract offered by an other-regarding principal

Suppose that the principal has other-regarding preferences given in (8.44). Consider the Stage 1 problem of the principal.

With probability λ , the principal encounters a self-regarding agent who puts in zero effort level (see Lemma 8.11) and no output is produced. By definition, in a trust contract, the principal must honor the promise to pay a wage of w even if no output is produced. Thus, the principal faces disadvantageous inequity. Using (8.44), the principal's payoff in this case is $\Pi = -w(2\alpha_p + 1) < 0$.

With probability $1 - \lambda$, the principal encounters an other-regarding agent. Lemma 8.11 describes the outcome and (8.49) gives the effort level. The effort choice of the agent equalizes the payoffs

¹³ We are being a bit informal here because we also require that $e \in [0, 1]$, thus, $e^T(w) = \min\{\sqrt{4w + 1} - 1, 1\}$. We shall assume that $\min\{\sqrt{4w + 1} - 1, 1\} = \sqrt{4w + 1}$ to abstract from technical issues of boundary points. We trust that the reader can supply the necessary, but simple, formalism if required.

of both parties, so there are no equity considerations in equilibrium. Given (8.40), (8.41), the principal's payoff is given by $\Pi = e^T(w) - w$, which is identical to the material payoff π .

Combining the two possibilities, and using (8.45), the expected payoff of the other-regarding principal is given by

$$E\Pi = -\lambda w(2\alpha_P + 1) + (1 - \lambda)(e^T(w) - w), \quad (8.50)$$

where $e^T(w)$ is defined in (8.49). The problem of the other-regarding principal is to maximize (8.50) subject to the individual rationality constraint of the other-regarding agent, $w \geq \frac{1}{2}(e^T)^2$; the effort reaction function e^T is already incentive compatible. Since e^T is increasing in w , it is not necessarily optimal to choose w , so that the individual rationality constraint of the other-regarding agent binds. We show this in the next proposition; w_F^T is the wage offered by the other-regarding (or fair) principal in a trust contract.

Proposition 8.32

- (i) If $\lambda \geq \frac{1}{2(1+\alpha_P)}$, then in the only possible equilibrium, $(w_F^T, e_c) = (0, 0)$, and the individual rationality constraint of both types of agents binds. The payoffs of both types of agents and the payoff of the principal are zero.
- (ii) If $\lambda < \frac{1}{2(1+\alpha_P)} < \frac{1}{2}$, then there is an interior solution to the wage rate given by $w_F^T = \frac{(1-\lambda)^2}{(1+2\lambda\alpha_P)^2} - \frac{1}{4}$. The effort level of the other-regarding agent is given by $e_F^T = \frac{2(\lambda-1)}{(2\lambda\alpha_P+1)} - 1$. w_F^T is decreasing in α_P and λ , and both types of agents derive strictly positive rents. The expected profit of the other-regarding principal in a trust contract is

$$\Pi_F^T = -\lambda w_F^T(2\alpha_P + 1) + (1 - \lambda)(e_F^T - w_F^T). \quad (8.51)$$

- (iii) It is never the case that the other-regarding principal elicits the maximum feasible effort level, $e = 1$, from the agent.

From Proposition 8.32(ii), if the probability of the self-regarding agent, λ , is low enough, then it is worthwhile for the principal to offer high enough rents. In response to these rents, the other-regarding agent responds with a high level of effort, on account of his inequity aversion. The self-regarding agent also benefits from the rents being offered but he puts in a zero effort level. This does create disadvantageous inequity for the other-regarding principal. But if the disadvantageous inequity parameter, α_P , of the principal is low, the costs are low relative to higher effort obtained from the other-regarding agent. Otherwise, the case in Proposition 8.32(i) obtains.

Finally, the surplus maximizing, or the most efficient effort choice of the agent is $e = 1$. The model predicts that we should not observe this choice. In other models, where agents care for efficiency, say, in Charness and Rabin (2002), agents may well choose $e = 1$. However, the empirical evidence shows that in gift exchange games, agents typically do not make the surplus maximizing choice (see Section 5.3, above).

II. Trust contract offered by a self-regarding principal

We can recover the case of the self-regarding principal from that of the other-regarding principal by setting $\alpha_P = 0$ in Proposition 8.32.

The comparison between the regimes of an other-regarding and a self-regarding principal under trust contracts is summarized below. The simple proof is omitted.

Remark 8.3 *The wage rate offered by an other-regarding principal is always lower relative to that of a self-regarding principal. The other-regarding agent reciprocates a higher wage by a higher effort, hence: (i) The effort level of the other-regarding agent is higher when the principal is self-regarding. (ii) Expected output is larger in the regime of a self-regarding principal as compared to the regime of an other-regarding principal.*

When $\lambda > \frac{1}{2}$, no principal (other-regarding or self-regarding) offers a positive wage, hence, the trust contract breaks down. The reason is that the self-regarding principal does not suffer from inequity aversion (which is costly for an other-regarding principal) when he meets a self-regarding agent. Thus, a self-regarding principal faces a lower marginal cost from a marginal increase in λ . If $\lambda \in \left(\frac{1}{2(1+\alpha p)}, \frac{1}{2}\right)$, then the self-regarding principal offers a trust contract but the trust contract breaks down if the principal is other-regarding. If $\lambda < \frac{1}{2(1+\alpha p)}$, then both types of principals, self-regarding and other-regarding, offer a trust contract. The effort level of the self-regarding agent is zero, regardless of the type of the principal.

Why does a self-regarding principal offer a higher wage than an other-regarding principal under a trust contract? The intuition is that the other-regarding principal, unlike a self-regarding principal, dislikes disadvantageous inequity. When the other-regarding principal encounters a self-regarding agent, he suffers disadvantageous inequity, which is increasing in the wage offered. Since the wage contract is offered prior to the type of the agent being revealed, the other-regarding principal offers a relatively lower wage.

INCENTIVE CONTRACTS

Recall our assumption that in the presence of an incentive contract, all parties to the contract behave as if they had self-regarding preferences (assumption A1). Consider now the incentive contract described above in which a self-regarding principal meets a self-regarding agent (with probability 1).

If the self-regarding agent were to shirk, relative to the contractual effort level, e_c , then it would be best to put in the minimum effort level, $e = 0$. This is because the probability of detection and fines is independent of the level of effort. Thus, the incentive compatibility condition of the self-regarding agent is

$$(1-p)w + p(w-f) \leq w - \frac{1}{2}(e_c)^2. \quad (8.52)$$

Rewriting (8.52),

$$e_c \leq \sqrt{2pf}, \quad (8.53)$$

gives the set of effort levels of the agent that are incentive compatible.

The self-regarding principal chooses a contract (w, e_c) to maximize expected profits

$$E\pi = (1-p)(e-w-k) + p(e-w-k+df) = e-w-k+pdf, \quad (8.54)$$

where d is a binary variable such that $d = 0$ if $e = e_c$ and $d = 1$ if $e \neq e_c$. There are two constraints, the incentive compatibility constraint in (8.53) and the individual rationality constraint $w - \frac{1}{2}e^2 \geq 0$. The satisfaction of these constraints at the contracted effort, e_c , ensures that $d = 0$. The objective function $E\pi$ is decreasing in w , and the incentive compatibility

constraint is independent of w , hence, the individual rationality constraint always binds, i.e., $w = \frac{1}{2}e^2$. Denote the maximum value of the expected profit of the principal in an incentive contract by $E\P^I$. In light of the discussion, above, the results in the next proposition should be unsurprising.

Proposition 8.33 *In an incentive contract, the optimal contracted effort level is given by*

$$e_I = \min \left\{ 1, \sqrt{2pf} \right\}. \quad (8.55)$$

There are two main cases.

- (i) *If $\sqrt{2pf} \geq 1 \Leftrightarrow pf \geq 1/2$, then $e_I = 1$. The individual rationality constraint of the agent binds, so $w = 1/2$ and the expected profit of the firm is given by $E\P^I = \frac{1}{2} - k$.*
- (ii) *If $\sqrt{2pf} < 1 \Leftrightarrow pf < 1/2$, then $e_I = \sqrt{2pf}$. The individual rationality constraint of the agent binds, so $w = pf$ and the expected profit of the firm is given by $E\P^I = \sqrt{2pf} - pf - k$.*

BONUS CONTRACT OFFERED BY AN OTHER-REGARDING PRINCIPAL

The bonus contract is described above. Analogous to the results in Lemma 8.10, 8.11, an other-regarding principal chooses a bonus, b , in Stage 3, in order to equate the monetary payoffs of the two parties, so,

$$e - w - b = w + b - \frac{1}{2}e^2.$$

Solving out for b we get

$$b = b(e, w) = \frac{1}{2}e - w + \frac{1}{4}e^2. \quad (8.56)$$

The bonus payment is increasing in the effort level of the agent, which captures the reciprocity of the principal. The bonus also depends negatively, and one to one, on the wage level because these are substitute instruments in ensuring equality of payoffs between the principal and the agent. The remaining analysis is similar to the case of trust contracts except that both parties know that in Stage 3, the bonus payment is given by (8.56).

The preferences of a self-regarding agent are given by $u = w + b(e, w) - \frac{1}{2}e^2$, which, using (8.56), is strictly concave in e . From the first order condition, we find that in a bonus contract, the optimal effort choice of a self-regarding agent is $e_S^b = 1$. Hence, even a self-regarding agent puts in the maximum possible effort level if he knows that the principal is other-regarding, who will, optimally, choose to honor his commitment to pay a bonus.

Now consider an other-regarding agent. From Lemma 8.10, 8.11 the optimal effort choice of the other-regarding agent will equate his monetary payoff with the principal, i.e.,

$$w + b(e, w) - \frac{1}{2}e^2 = e - w - b(e, w). \quad (8.57)$$

It turns out that (8.57) is satisfied for any value of e (because the choice of $b(e, w)$ ensures that payoffs are equated ex-post). Hence, the payoff of the other-regarding agent is identical to that of a self-regarding agent and is given by $U = w + b(e, w) - \frac{1}{2}e^2$. Substituting $b(e, w)$ from (8.56), we get $U = \frac{1}{2}e - \frac{1}{4}e^2$. Clearly this payoff is also maximized at $e_F^b = 1$.

Given that the Stage 3 bonus plays the same role as a wage, and is offset one for one with the wage, we can focus on the case where the Stage 1 wage offer of the principal is zero. Since the other-regarding and the self-regarding agents both put in the maximal effort level, the other-regarding principal's expected payoff in a bonus contract is $E\Pi = e - b(e, w)$, which is independent of λ . On substitution of $e = 1$, $w = 0$ and (8.56), the expected profit in a bonus contract is given by

$$E\Pi_F^B = \frac{1}{4}. \quad (8.58)$$

We summarize these results below.

Proposition 8.34 *In a bonus contract, when the principal is other-regarding, the self-regarding and the other-regarding agents exert identical effort levels equal to 1, i.e., $e_S^b = 1 = e_F^b = 1$. The expected profit of the principal is $E\Pi_F^B = \frac{1}{4}$.*

COMPARISON OF PROFITS UNDER THE THREE CONTRACTS

Does the other-regarding principal prefer an incentive contract or a bonus contract? The answer depends on the parameters and there are many possibilities. Classical contract theory under self-regarding preferences predicts that the incentive contracts should always dominate. We consider only one example here, in which bonus contracts dominate all others. Using Proposition 8.33, and (8.58), when $\sqrt{2pf} \geq 1$,

$$E\Pi^I \geq E\Pi_F^B \text{ as } \frac{1}{4} \geq k,$$

and when $\sqrt{2pf} < 1$

$$E\Pi^I \geq E\Pi_F^B \text{ as } \sqrt{2pf} - pf \geq \frac{1}{4} + k. \quad (8.59)$$

Thus, the incentive contract dominates the bonus contract when (1) the deterrence parameters p, f are high enough, so that hard evidence of shirking is discovered with high enough probability and then punished with severe enough consequences, and (2) the cost of the monitoring technology, k is low enough. When $\lambda > \frac{1}{2}$ the trust contract is not offered by either an other-regarding or a self-regarding principal (see Remark 8.3). In this case if $k > \max\{\frac{1}{4}, \sqrt{2pf} - pf - \frac{1}{4}\}$, then the bonus contract dominates all other contracts.

These results are only illustrative of the departures one gets from the standard model when individuals have other-regarding preferences. A classical game theorist might argue that the existence of non-competitive rents embodied in wage levels above the market level can equally well be explained using infinitely repeated games between purely self-regarding firms and workers. However, such arguments are not persuasive on empirical grounds. (1) The evidence strongly confirms the existence of economic agents who exhibit other-regarding preferences; much of Part 2 of the book attempts to establish the relevant evidence base. (2) The model with other-regarding preferences generates these results without the machinery of repeated games and reputation formation. One can then design appropriate experiments to test the relative merit of these two alternative explanations, a task that we turn to in the next section.

8.3.3 Evidence on contract choice under contractual incompleteness

Fehr et al. (2007) examine, experimentally, the choice between incentive, trust, and bonus contracts. They ran several sessions with 10 periods in each session and 20–24 individuals in each session. Agents are randomly matched with a different principal in each period. When principals are given a choice between trust and incentive contracts (the TI treatment), there is a strong trend towards the incentive contracts over successive rounds of the experiment. In the final three periods, more than 80% of all contracts offered were incentive contracts. Trust contracts were tried in one of the rounds by at least 70% of the principals but were not preferred. Effort was significantly higher in the incentive contracts relative to the trust contracts.

Profits were negative, on average, in the incentive contracts because the majority of the incentive contracts, 58.5%, violated the no-shirking constraint in (8.53). Most principals used the maximum fine and asked for an effort level that was too ambitious. However, the demanded effort fell over time and the share of incentive compatible contracts increased with time. The incentive-compatible incentive contracts are more profitable than the trust contracts but the non-incentive compatible ones are less profitable than trust contracts. The trust contracts in Fehr et al. (2007) do not perform as well as those in Fehr et al. (1993). A potential reason is that effort–wage relation is far steeper in the latter and the set of effort choices is discrete.

When principals are given a choice between bonus and incentive contracts (the BI treatment), then the overwhelming majority of principals chooses the bonus contract. The results, averaged across all sessions, are shown in Figure 8.1. From period 6 onward, the bonus contract is chosen in approximately 90% of the cases and approaches 96% in the final period. In the BI treatment, in the bonus contract, the average effort level is relatively higher and the gap between the demanded and actual effort is the lowest.

On average, the principal's profit in the bonus contracts is about 4 times higher. Roughly 40% of the compensation of agents is paid out as a bonus by the principal. The principals are not fully




Figure 8.1 Share of bonus contracts and incentive contracts over successive rounds.

Source: Fehr et al. (2007) with permission of The Econometric Society.




Figure 8.2 Relation between the actual bonus paid by the principal and effort level of the agent.

Source: Fehr et al. (2007) with permission of The Econometric Society.

committed to their announced bonuses. Rather, the actual bonus paid out depends heavily on the effort level of the agent, as shown in Figure 8.2. In conclusion, across all three contractual forms, the principals strongly prefer the bonus contract to the other contracts.

Fehr and Schmidt (2007) address one shortcoming of the Fehr et al. (2007) paper. They allow principals to combine a voluntary bonus (the carrot) with an incentive contract that allows for fines (the stick). It was hoped that this more general contract would probably perform the best. However, they found that two-thirds of the principals still preferred the pure bonus contract to the more general contract and the latter did not elicit greater effort.

A lower bonus is offered in the combined contract as compared to a bonus contract alone. Since agents might not know the type of the principal, they may take the choice of a combined contract by the principal as a signal that the principal is self-regarding and, so, less likely to offer a bonus. In anticipation, agents put in a low level of effort, which leads to lower profits in the combined contract relative to the bonus contract. The belief that principals who are less fair are also the ones more likely to choose the more general contract and less likely to pay a bonus is confirmed in the experimental findings.¹⁴ This evidence further underpins assumption A2 in the theoretical model made above.

The evidence in Fehr et al. (2007) indicates that implicit contracts were chosen in 88% of the cases even when an explicit contract was available. Implicit contracts also elicited higher effort and were more profitable. The key to the success of the implicit contracts was the mutual

¹⁴ In an insightful paper, Bénabou and Tirole (2003) consider the conflict between *intrinsic* and *extrinsic* motivation as a result of asymmetric information in a principal-agent framework. They find conditions under which incentives will or will not be an effective instrument. For instance, roughly paraphrased, one of their results is that incentives may crowd out intrinsic motivation when the principal is better informed and the agent has low ability in performing a task. Under these conditions, giving the agent a high-powered incentive may lead the agent to infer that his ability is actually low, which may be demotivating. These results also have a bearing on the difference in experimental practice in economics and psychology. Experimental economists emphasize the role of incentives, or extrinsic motivation, while experimental psychologists emphasize the role of intrinsic motivation and downplay the use of incentives in experiments. Most of the experimental results that we report are conducted under full information, so they are not suitable to test or apply the results of Bénabou and Tirole (2003).

reciprocity of the principal and the agent. For instance, reciprocal principals were expected to honor their promise to pay a bonus, hence, eliciting greater effort. Explicit contracts might on the other hand be perceived to be hostile, so they elicit lower effort.

Suppose that there are two parties to a contract and at least one party must make relation-specific investment in the relationship. An influential view in the property rights approach predicts that the institution of *joint ownership* will hamper short-term *relation-specific investment* because each party will fear that, ex-post, its investment will be held up by the other party (Hart, 1995). One proposed solution is to vest *residual control rights* in the hands of the party that makes the investment in the physical asset (Grossman and Hart, 1986).

A simplified version of the Grossman and Hart (1986) framework is considered in the two-player experiments conducted by Fehr et al. (2008). In the first stage, one or both players are asked to choose the form of their relationship, e.g., *single* or *joint* ownership (see below). In the second stage, the two players, A and B, make respective investments, $a, b \in \{1, 2, \dots, 10\}$, that lead to a joint surplus $v(a, b)$. The cost of the investment equals the level of the investment. Investments are made sequentially: B invests first, followed by A who observes B's investment decision before choosing his own investment. Two possible regimes are considered.

1. Under *A-ownership*, A hires B on a fixed wage contract that pays a wage w . Net of the respective investments, the payoffs of A and B are respectively given by $v(a, b) - a - w$ and $w - b$.
2. Under *joint ownership*, each player gets one half of the total surplus, net of their respective investment costs. So the respective payoffs of A and B are: $\frac{1}{2}v(a, b) - a$ and $\frac{1}{2}v(a, b) - b$.

The game has been set up in such a way (by an appropriate choice of $v(a, b)$) that the first best levels of investments are $a = b = 10$. The predictions of the standard incomplete contracts model for self-regarding preferences are as follows. Under A ownership, B has no incentive to invest, so b takes its lowest value, $b = 1$. Under joint ownership, each individual has to share half the marginal return with the other party. Given the setup of the problem, this induces the lowest possible investment for each player $a = b = 1$.¹⁵

The predictions of the model with self-regarding preferences are rejected by the experimental evidence. Player A chooses joint ownership in more than 80% of the cases. The respective investment levels chosen by the players, on average, under joint ownership, are $a = 6.5$ and $b = 8.9$. Furthermore, A earns more, on average, if he chooses joint ownership, relative to A ownership. The explanation is that joint ownership is more conducive to reciprocal behavior. Player B expects A to behave reciprocally. This would be true even if player B is self-regarding but knew that the other player has other-regarding preferences. Such expectations (which are fulfilled in equilibrium) induce player B to invest.

To summarize: the experimental evidence shows that incentive contracts are often outperformed by bonus contracts; incentive contracts can be viewed as hostile; and the institution of joint ownership can tap into the reciprocity of economic agents rather than holding up investment through opportunistic channels.

Suppose that n ex-ante identical agents work in a team for a principal and everyone has self-regarding preferences. Each of the n agents undertakes a sub-project, whose success is guaranteed by working hard; under shirking, the sub-project succeeds with probability $0 \leq p_i < 1$.

¹⁵ Essentially, A's choice of $a = a(b)$ is given from $0.5v_a(a(b), b) = 1$. Substituting in the first stage, the optimal choice of b is found from $0.5 \left[v_a(a(b), b) \frac{da}{db} + v_b(a(b), b) \right] = 1$. For the functional forms chosen in the experiments, it turns out that $a = b = 1$.

Effort on the sub-project or its outcome cannot be contracted on. Each sub-project is part of a grand project whose success can be verified by all parties, hence, the wage vector of the n agents, $(w_1, w_2, \dots, w_n)$, is conditioned solely on its outcome. Let $0 \leq \hat{n} \leq n$ agents work hard on their sub-projects. Then the probability that the grand project succeeds is $p(\hat{n})$ where p is an increasing and convex function. The aim of the principal is to minimize aggregate wage payments such that all the agents find it incentive compatible to work hard.

Under these conditions, Winter (2004) showed that it is optimal to pay unequal wages for equal work. Thus, wage inequality (or discrimination) among agents is efficient. The intuition is along the following lines. Since effort has the nature of a public good, agents would wish to free-ride on the effort levels of others. Each agent, if rational, can foresee this. Hence, he may withhold costly effort because if others do not work hard, the project is unlikely to succeed, which, in turn, does not justify his effort. If some agents are paid high enough discriminatory wages, so as to raise their stakes in the success of the grand project, the remaining agents feel sufficiently reassured to also work hard. Thus, discrimination in wages allows agents to coordinate their beliefs on the success of the grand project. Conditional on these beliefs, they indeed work hard, which fulfills the original beliefs.

Cabrales et al. (2010) and Goerg et al. (2010) report experimental support for this framework. The experiments show that efficiency is lower under equal pay for equal work, relative to pay discrimination. However, the experimental results of Abeler et al. (2010) that we have examined in Chapter 6 show that equal pay for unequal effort levels reduces effort levels in subsequent rounds and reduces output.

In actual practice, the implementation of a discriminatory wage scheme may run into difficulties with various employment legislations.¹⁶ A more fundamental concern with this approach is that other-regarding preferences have been widely documented, not only in experimental games (Part 2 of the book is a testament to this claim) but also in a range of principal-agent situations. Insofar as principals can rely on the other-regarding preferences of even a fraction of the agents, or on their reciprocity, the optimal contracts are likely to be markedly different. Indeed, the empirical contracts in Cabrales et al. (2010) and Goerg et al. (2010) concentrate on a narrow range of contracts that rule out reciprocity-based contracts such as bonus and trust contracts, which are found to be efficacious, as the evidence in this section suggests.

There might be substantial heterogeneity with respect to worker characteristics such as their risk aversion and other-regarding preferences. Will workers of different productivity levels and different personal characteristics self-select or sort themselves into fixed wage contracts and incentive contracts (such as piece rate wages, tournaments, or revenue sharing)? It is quite difficult in field studies to observe the sorting decision. Hence, it is not clear if the observed positive effects of incentives on effort are the result of pure-incentive effects or the result of sorting of more productive workers into incentive contracts. By contrast, experiments allow a clean separation of these different motives.

Dohmen and Falk (2011) address these questions within a novel experimental design. They find that more productive workers self-select themselves into incentive schemes relative to fixed wage schemes. Output and effort are higher under the incentive contracts but these effects are driven largely by the sorting effects of workers. Furthermore, more risk averse subjects, and women, are

¹⁶ However, it is possible that firms may find other means of engaging in such practices. For instance, non-pecuniary discrimination or seniority based pay, etc. However, non-pecuniary discrimination and seniority based pay have other well-known sources also.

more likely to self-select themselves into the fixed wage contract. Somewhat surprisingly, other-regarding preferences play only a marginal role in the self-selection decision.

Dohmen and Falk (2011) assume that it is not possible for firms to know the relevant traits of potential employees, hence the use of contractual form as a screening device. Englmaier et al. (2014) give the complementary argument that several proxies about the characteristics of workers are available to firms. These include the CV of the workers (firms might shortlist, say, 4 candidates from an application pool of 100), ability, personality tests, and interviews.¹⁷

In a novel set of experiments, Englmaier et al. (2014) first elicit the productivity and reciprocity parameters of individuals. Productivity is elicited by giving subjects a real effort task that requires matching words to a four digit code in a list. Reciprocity is elicited by the degree of *trustworthiness* in a trust game. Firms may be given no information (treatment N), only information about trust (treatment T), only information about productivity (P), or information about productivity and trust (treatment PT). Once the relevant information (which includes no information in treatment N) is given to the firm, the firm and the worker play a standard gift exchange game. One can then compare the wage, effort level, and profit level of the firm in each treatment.

The results are as follows. There is a substantial wage premium for productivity and trust. More productive and more trustworthy subjects receive significantly higher wages. However, the wage premium for productivity is twice as large as the premium for trust. In regression analysis, with the effort level of the worker as the dependent variable, once the wage is included as an explanatory variable, then productivity and trust add no further explanatory power. The reason is that wages are endogenous, and they correctly incorporate the effects of productivity and trustworthiness. Firms make the highest profits when the workers have both desirable features, productivity and trust. More trustworthy workers also induce a smaller variation in the profits and output levels of the firm.

8.4 Reciprocity and long-term contracts

The explanation of the structure of long-term relations, say, between firms and their suppliers, is of fundamental importance in industrial organization. Consider infinitely repeated games in which considerations of future costs and benefits of current actions give rise to *instrumental reciprocity*. In addition, suppose that individuals also exhibit *intrinsic reciprocity*. What is the interaction between these two different types of reciprocity? For instance, are these two types of reciprocity complementary or substitutes? These issues are addressed by Brown et al. (2004), which we now consider.

A repeated experiment was performed with 15 trading periods and 238 subjects who were assigned to play either the role of a firm or a worker. To model the importance of issues of reputation formation, each firm and worker were assigned an *ID* number. Each firm could only hire one worker in any period. Any period consisted of two subperiods. In the first subperiod, a firm could offer a contract (w_b, e_n, ID) , i.e., the contract specifies a *binding* wage level, w_b , a *non-binding* effort level, e_n , and the firm's identification number, *ID*. Firms can make either a *public offer* (open to all workers) or a *private offer* (open only to a worker with a particular *ID*

¹⁷ For instance, in the hiring process for academics in universities, the ability and personality characteristics are sought to be discovered through references, CV, one-to-one meetings with staff, job talk in the department, and interviews with committees.

number). In each period, and in each treatment, there were ten workers and seven firms, so there was always an excess supply of workers.

The experiments considered the following three treatments.

1. In the complete contracts condition, or *treatment C*, there is third party enforcement, i.e., the experimenter makes sure that the contracted effort level is always honored. Thus, the contracted effort level, e_n , equals actual effort level, e_a .
2. In the incomplete contracts condition, or *treatment ICF*, there is no third party enforcement of the contract, so it is possible that $e_n \neq e_a$.

In each of these two treatments, there is the possibility that the firm hires the same worker for more than one period, which constitutes a *long-term relation*. Hence, long-term relations, if they are observed, are purely endogenous.

3. To better understand the effectiveness of long-term relations, a control treatment, *treatment ICR*, is introduced. This is identical to treatment ICF, except that long-term relations are ruled out. In the actual experiments, this was achieved by jumbling the ID numbers of players at the end of each period, so that firms could not guarantee hiring the same worker through private offers even if they wanted to.

In the event of a successful contract between a firm and a worker, the payoff of the self-regarding firm, π , and the self-regarding worker, u , are respectively

$$\pi_s = 10e - w; u_s = w - c(e). \quad (8.60)$$

The cost function of effort, $c(e)$, is increasing and convex; it is shown in Table 8.1.

If a firm or a worker is unable to enter into a contract with the other party, then their respective payoffs are $\pi_s = 0$, $u_s = 5$, where the worker's outside option, e.g., unemployment benefits, is given by 5. Thus, the total social surplus in the worker–firm relationship (net of the worker's outside option) is given by

$$S = \pi_s + u_s - 5 = 10e - c(e) - 5. \quad (8.61)$$

8.4.1 The outcome with self-regarding preferences

We consider self-regarding preferences first, followed by other-regarding preferences in the next section.

Lemma 8.14 *The socially efficient level of effort in any period is given by $e = 10$.*

Proof: In any period, marginal benefit of a unit of effort is 10, while the highest marginal cost is 3 (see Table 8.1). Further, the payoffs are independent across periods. Hence, in any period, it is optimal to choose the highest effort level, $e^* = 10$. ■

Table 8.1 The cost of effort function.

e	1	2	3	4	5	6	7	8	9	10
$c(e)$	0	1	2	4	6	8	10	12	15	18

Source: Brown et al. (2004).

Lemma 8.15 *In the ICF and ICR treatments, in any pairwise interaction, and in every period, the worker chooses the lowest effort level $e_I = 1$ and the firm pays a wage $w_I = 5$. The total social surplus $S_I = 5$ is reaped entirely by firms.*

Proof: Effort is not enforceable by a third party in the ICF and ICR treatments, workers receive a binding wage, w_b , and there is full information in each of the finite number of periods. Consider the last period. Self-regarding workers choose the least possible effort, i.e., $e_I = 1$. Anticipating this, firms choose a wage rate to ensure that the workers get no more than their outside option, so $w_I = c(1) + 5 = 5$. Substituting $w_I = 5$, $e_I = 1$ in (8.61) and using Table 8.1 we get that $S_I = 5$. However, since workers simply get their outside option, the entire surplus is reaped by firms. Now consider the penultimate period. In this period, the matching of firms and workers may be different but the solution is identical to that in the last period. Continuing backwards in this manner, we obtain an identical solution in each period. ■

Lemma 8.16 *In the complete contracts treatment, treatment C, in every period, the optimal effort level is $e_c = 10$, wage is $w_c = 23$ and the surplus is $S_c = 77$.*

Proof: A third party, the experimenter, ensures that the contracted effort level in treatment C is binding. Consider the last period. As in the proof of Lemma 8.14, it is optimal to induce the highest effort level. The wage must now ensure that the worker's individual rationality constraint is satisfied, so $w_c = c(10) + 5 = 23$. Substituting $e_c = 10$ in (8.61) and using Table 8.1 we get that $S_c = 77$. By backward induction, the same outcome occurs every period.

Lemma 8.17 *Under self-regarding preferences, firms are indifferent between making private and public offers in the three treatments, C, ICF, ICR.*

Proof: Since there are no issues of reputation building, there is no additional gain from hiring a worker more than once. In any period, any worker is as good as another. Hence, private offers are just as good as public offers. ■

8.4.2 The outcome with other-regarding preferences

Suppose that at least some workers have other-regarding preferences of the Fehr–Schmidt form (FS). Recall that there are 15 trading periods. We make the following assumptions.

- A1. In any period, each individual compares his/her payoffs only with their trading partner. Thus, if i, j are trading partners, with respective material payoffs x_i and x_j , the FS utility of player i is:

$$U_i = x_i - \alpha_i \max \{x_j - x_i, 0\} - \beta_i \max \{x_i - x_j, 0\}. \quad (8.62)$$

- A2. $\alpha_i = \beta_i = 0.5 + \epsilon$ where ϵ is some small but strictly positive number.
A3. By giving up a dollar, each player can punish the trading partner by an amount $\delta > 1 + \frac{1}{0.5+\epsilon}$.
A4. 60% of all players are other-regarding, while 40% are purely self-regarding (this is roughly consistent with some of the empirical evidence reviewed in Chapter 12).
A5. Other than in treatment ICR, in any period, firms can choose to fire the existing worker or hire the worker for the next period.
A6. There is no discounting of future payoffs.

Lemma 8.18 *Under assumptions A1–A6, if players have other-regarding preferences, they prefer to equalize payoffs with their trading partners.*

Proof: Suppose first that $x_j > x_i$. Using (8.62), the payoff of player i is $U_i = x_i - \alpha_i (x_j - x_i)$. Suppose that player i now gives up n units to punish the trading partner in order to equate payoffs, i.e., $x_j - n\delta = x_i - n$. Thus, $n = \frac{x_j - x_i}{\delta - 1}$. The post-punishment payoff of player i is $U_i = x_i - n = x_i - \frac{x_j - x_i}{\delta - 1}$. Given the original supposition that $x_j > x_i$, an equalization of payoffs improves player i 's utility if $x_i - \frac{x_j - x_i}{\delta - 1} > x_i - \alpha_i (x_j - x_i)$ or $\alpha_i > \frac{1}{\delta - 1}$. Substituting $\alpha_i = 0.5 + \epsilon$ we get the condition $\delta > 1 + \frac{1}{0.5 + \epsilon}$, which is true by assumption A3 above.

Now suppose that $x_j < x_i$. Using (8.62), the payoff of player i is $U_i = x_i - \beta_i (x_i - x_j)$. In order to equalize payoffs, suppose that player i makes a transfer of $(x_i - x_j) / 2$ to player j . The post-transfer payoff of player i is $U_i = x_i - \frac{x_i - x_j}{2}$. It is worth making the equalizing transfer if $x_i - \frac{x_i - x_j}{2} \geq x_i - \beta_i (x_i - x_j)$, which holds because $\beta_i > \frac{1}{2}$. Thus, both players find it optimal to equalize payoffs. ■

THE SOLUTION UNDER ASYMMETRIC INFORMATION

The remaining theoretical discussion is entirely organized around Proposition 8.35 below, which collects all the relevant results in this section. Workers have asymmetric information about the type of the firm; the type of the firm can be self-regarding or other-regarding. Prior to hiring a worker, firms also do not know if the worker is self-regarding, although, from assumption A4, they know that 60% are other-regarding and 40% are self-regarding. Thus, there is two-sided asymmetric information. Let us first offer the insights that underpin Proposition 8.35.

The finite period setup is realistic (e.g., because workers eventually retire). It also allows the workers' types to separate in the last period because self-regarding workers will always shirk in the last period. Since other-regarding workers respond to higher wages with higher effort (as in gift exchange games), the firm may find it profitable to offer rents to the workers. While self-regarding workers do not *intrinsically reciprocate*, however, the attractiveness of rents also induces them to *instrumentally reciprocate*. Hence, self-regarding workers would like to build a reputation for reciprocity and enjoy rents—they milk this reputation for as long as it is possible in the finite horizon game. This behavior suits other-regarding firms because profits are increasing in the effort levels of workers. Thus, *reputation effects interact in an important manner with reciprocity effects*.

The possibility that a fraction of the workers are other-regarding (assumption A4) means that even in the last period of the game, the firm can offer rents to the workers and, rationally, hope that other-regarding workers will respond with greater effort. Consider a self-regarding worker in the penultimate period of the game. If the worker shirks, then he will be fired, hence, he will be unable to avail himself of the last period rents. The contract offered in the penultimate period must ensure that the current gain to shirking of the self-regarding worker is smaller as compared to the expected future loss of rents, as in models of efficiency wages (Shapiro and Stiglitz, 1984).

Such repeated contracts that satisfy incentive compatibility in each period, provide the necessary inducement for self-regarding workers to put in greater effort. This, in turn, facilitates the endogenous formation of long-term contracts and we can expect the ICF treatment to foster greater effort and rents relative to the ICR treatment, which does not allow long-term relations. One may, thus, also expect a higher share of private offers relative to public offers in the ICF treatment.

Why does the firm not offer very low rents in the penultimate period? The reason is that it will then be perceived as being self-regarding itself and workers with other-regarding preferences would react negatively to lower wages offered by reducing their effort. As is typically the case with repeated games, there are a large number of equilibria. Finally, in treatment C, because the contracts are always enforced with the help of a third party, there is no need to give any incentives or rents. And there is no need to form long-term relations either. So one would expect a preponderance of public offers in treatment C.

We now use these insights to formalize the relevant results in Proposition 8.35.

Proposition 8.35 *Consider two-sided asymmetric information in which the type of the firm and the worker (self-regarding or other-regarding) is private information. The game is played over 15 periods and the outside option of any worker is 5. Assumptions A1–A6 hold. In the perfect Bayesian equilibrium of the model:*

- I. *The strategy of the other-regarding and the self-regarding workers is as follows:*
 - (a) *Other-regarding workers choose their effort levels to equalize payoffs with their trading partners every period and are willing to work if their individual rationality constraint is satisfied.*
 - (b) *Faced with a contractual, but non-binding, effort level, e_n , self-regarding workers choose to shirk (i.e., choose $e < e_n$) in any period if the cost of effort, $c(e_n)$, is more than the expected loss from being fired.*
- II. *The strategy of the other-regarding and self-regarding firms is as follows:*
 - (a) *Firms renew the contract of a worker in period t if in period $t - 1$ the worker put in $e = e_n$. Otherwise the worker is fired and earns his outside option in all future periods. Public offers are made by the firm in the first period and in any period following the decision of workers to shirk ($e < e_n$).*
 - (b) *Both types of firms (self- and other-regarding) engage in a pooling contract for the first 14 periods but offer separating contracts in the 15th period. In periods 1–13, the contract offered is $(w_b, e_n) = (59, 10)$. In period 14, the contract offered is $(w_b, e_n) = (40, 7)$. If $e = e_n$, then the payoffs of the firm and the worker are equalized in the first 14 periods. In period 15, the self-regarding firm offers the contract $(w_b, e_n) = (34, 6)$ but it is not unique. The other-regarding firm offers the contract $(w_b, e_n) = (5, 1)$.*
- III. *The out of equilibrium beliefs are as follows:*
 - (a) *All workers believe that the firm will terminate employment in period $t + 1$ if it does not pay a fair wage (that equalizes payoffs) in period t .*
 - (b) *Higher average effort levels can be sustained in the ICF treatment relative to the ICR treatment. Long-term relations arise endogenously in the ICF treatment and there are greater private offers in the ICF treatment relative to the ICR treatment.*

Proof: Consider self-regarding firms first. From Lemma 8.18 we know that an other-regarding worker will always equalize payoffs, i.e.,

$$10e - w = w - c(e). \quad (8.63)$$

Denote the solution to this equation by $e(w)$; this is the effort reaction function of the agent. For pedagogical convenience we abuse notation by writing $c'(e)$ for the marginal cost of effort, although the effort levels are actually discrete. Using (8.63),

$$\frac{de(w)}{dw} = \frac{2}{10 + c'(e)} > 0, \quad (8.64)$$

hence, other-regarding workers will always reciprocate a higher wage offer with higher effort, as in gift exchange games.

We solve the game using backward induction.

(1) Consider first, the last period, period 15: Given the stated strategies in the proposition, a pooling equilibrium is played in the first 14 periods, hence, the type of the firm, other-regarding or self-regarding, is not revealed to the workers. Using (8.60) and assumption A4, the expected payoff of a self-regarding firm is given by

$$E\pi_s = 0.6 [10e(w) - w] + 0.4 [10 - w]. \quad (8.65)$$

The other-regarding worker reciprocates higher wage with higher effort (8.64). But the self-regarding worker chooses to shirk and put in the minimum effort level, $e = 1$, because it is the last period. Differentiating (8.65) with respect to w and using (8.64) we get

$$\frac{\partial E\pi_s}{\partial w} = \frac{12}{10 + c'(e)} - 1. \quad (8.66)$$

Since the cost of effort function is weakly convex throughout, for an interior solution, solving (8.66) gives $c'(e) = 2$. From Table 8.1, $c'(e) = 2$ for all effort levels in the set $\{3, 4, \dots, 8\}$, hence, we have multiple equilibria. Let us choose the particular solution $e = 6$. To induce the other-regarding worker to undertake an effort level $e = 6$, the wage must be such that the payoffs of the worker and the firm are equalized (see Lemma 8.18). Substitute $e = 6$, $c(6) = 8$ in (8.63) to get $w = 34$. Hence, the (non-unique) contract offered by the self-regarding firm in period 15 is $(w_b, e_n) = (34, 6)$.

Since a pooling equilibrium is played in the first 14 periods, the type is unknown at the start of period 15. The expected payoff of the self-regarding firm $E\pi_s$ is given in (8.65). The payoff of the worker is $u = w - c(e)$. Using (8.62), the utility of an other-regarding firm is given by

$$\pi_f(k) = \pi_s - \alpha_f \max\{u_k - \pi_s, 0\} - \beta_f \max\{\pi_s - u_k, 0\},$$

where $k = f, s$ is the type of the worker—other-regarding or self-regarding—and π_s denotes the self-regarding material payoff of the firm, given in (8.60). Using assumption A4, the expected profit of the other-regarding firm from the contract (w_b, e_n) is given by

$$E\pi_f(w_b, e_n) = 0.6\pi_f(f) + 0.4\pi_f(s). \quad (8.67)$$

What is the optimal contract offered by the other-regarding firm in the last period? One can show that the other-regarding firm will prefer the contract $(w_b, e_n) = (5, 1)$ to any other contract that induces an effort level from the other-regarding worker belonging to the set $\{3, 4, \dots, 8\}$. Recall that this is the set of effort levels that the self-regarding firm finds it optimal to induce from the other-regarding worker. The self-regarding worker, of course, only puts in the minimum effort level $e = 1$ in period 15 because there is no future to worry about. We show that it is non-optimal for the other-regarding firm to elicit the effort level $e = 3$ (a similar check can be made against the other effort levels, i.e., 4, 5, $\dots$, 8).

Suppose first that the other-regarding firm offers the contract $(w_b, e_n) = (5, 1)$. It is simple to check that the other-regarding worker also chooses the same effort as a self-regarding worker, i.e., $e = 1$, because it equalizes the payoffs with the firm (see Lemma 8.18) and ensures that the individual rationality constraint is met (recall that the outside option of the worker is 5). Since both types of worker choose $e = 1$, using (8.67), we get that, corresponding to the contract $(w_b, e_n) = (5, 1)$, the expected profit of the other-regarding principal is

$$E\pi_f(5, 1) = 5. \quad (8.68)$$

Let us now compute the payoff of the other-regarding firm from the contract $(w_b, e_n) = (w_b, 3)$, where the non-contractible effort level is $e_n = 3$. The actual effort level supplied by the self-regarding worker $e_a = 1$ because it is the last period. What level of wage, w_b , does the firm need to offer, so that the other-regarding worker chooses $e = 3$? Since the other-regarding agent equalizes material payoffs with the principal, substitute $e = 3$, $c(3) = 2$ in (8.63) to get $w_b = 16$. Substitute $w_b = 16$, $e = 3$ in (8.65) to get:

$$E\pi_s = 0.6 [10 \times 3 - 16] + 0.4 [10 - 16] = 6. \quad (8.69)$$

The respective payoffs of the self-regarding and other-regarding workers are $16 - c(1) = 16$ and $16 - c(3) = 14$. It is simple to calculate that $\pi_f(f) = 14$ and $\pi_f(s) = -6 - 22\alpha_f$. Thus, we get the payoff of the other-regarding firm as

$$E\pi_f(16, 3) = 0.6 \times 14 + 0.4 (-6 - 22\alpha_f). \quad (8.70)$$

Comparing (8.68), (8.70), $E\pi_f(5, 1) > E\pi_f(16, 3) \Leftrightarrow 5 > 6 - 8.8\alpha_f \Leftrightarrow 8.8\alpha_f > 1$, which holds because $\alpha_f > 0.5$ (assumption A2). Hence, the contract $(5, 1)$ dominates the contract $(16, 3)$ for the other-regarding firm. A similar exercise shows that it is not optimal for the other-regarding firm to induce any effort level other than $e = 1$.

We conclude that in period 15, the self-regarding firm offers $(w_b, e_n) = (34, 6)$, while the other-regarding firm offers $(w_b, e_n) = (5, 1)$. Thus, the two types of firms offer a separating contract in period 15.

(2) Now consider period 14: Workers do not know the type of the firm. But given the strategies, they are aware of the contracts that will be offered in period 15 by both types of firms. From assumption A4, they know the respective proportions of the types of the firms. Hence, the expected future payoff of other-regarding workers, $E\pi_f$, from not shirking in period 14, so ensuring that they get hired by the same firm in period 15 is

$$E\pi_f = 0.6(5 - c(1)) + 0.4(34 - c(6)) = 13.4. \quad (8.71)$$

A self-regarding worker, on the other hand, will put in $e = 1$ in period 15, hence, his expected future payoff, $E\pi_s$, from not-shirking in period 14 is

$$E\pi_s = 0.6(5 - c(1)) + 0.4(34 - c(1)) = 16.6.$$

If the worker shirks in period 14, then given the strategies in the proposition, he is fired and earns his outside option of 5. So the expected future net gain from not-shirking in period 14 for the

other-regarding and self-regarding workers is respectively $13.4 - 5 = 8.4$ and $16.6 - 5 = 11.6$. Since the worker receives a binding wage, the current gain to a self-regarding worker from shirking in period 14 is $c(e_n) - c(1) = c(e_n)$ (from Table 8.1, $c(1) = 0$) where e_n is the non-binding effort level in period 14. Hence, the self-regarding worker shirks if the current gain from shirking is greater than the expected future loss from shirking, i.e.,

$$c(e_n) > 11.6. \quad (8.72)$$

From Table 8.1, the set of effort levels which satisfy this inequality is $\{8, 9, 10\}$. Hence, in period 14, any effort level $e \leq 7$ will induce no-shirking from the self-regarding worker. A higher effort level leads to greater output and an other-regarding worker equalizes payoffs. Hence, inducing a higher effort level from the set of incentive-compatible effort levels for the self-regarding worker allows the principal to enjoy a higher payoff. Thus, it is optimal for the firm to induce $e_n = 7$ in period 14.

Using (8.63) and Table 8.1, the firm has to then pay a payoff equalizing wage $w_b = \frac{10 \times 7 + c(7)}{2} = 40$. Why is it not optimal to pay $w > 40$? If $w > 40$, then other-regarding workers will reciprocate with extra effort, however, self-regarding workers will not reciprocate because their no-shirking constraint will be violated. But does the extra effort from other-regarding workers make it worthwhile to offer $w > 40$? The answer is in the negative. The firm makes no additional profits from self-regarding workers. The extra profit made on account of the other-regarding workers is $\frac{d}{dw} (10e(w) - w)$, which when evaluated at $e = 7$ gives $\frac{20}{12} - 1 = 0.67$. However the extra wage cost on account of the self-regarding workers is 1. Hence, the net change in expected profit is $-1 + 0.67 = -0.33 < 0$. It is also not optimal to pay $w < 40$. Given out of equilibrium beliefs, $w < 40$ is construed as an unfair wage, hence, the expectation is that workers will not be rehired by the firm for period 15. But this induces full shirking from workers in period 14. Hence, $w_b = 40$ is the optimal wage for period 14.

The firm must also make a private offer to the worker at the end of period 14. If it did not do this, then there is a (large) possibility that the current worker will not be rehired. This dilutes the incentives of the worker to work in period 14.

(3) Let us now focus on period 13: Workers know (from the firm's stated strategy) that for period 14, the pooling contract offered is $(w_b, e_n) = (40, 7)$, while for period 15, with probability 0.6, the firm is other-regarding and offers the contract $(w_b, e_n) = (5, 1)$, while with probability 0.4 the firm is self-regarding and offers the contract $(w_b, e_n) = (34, 6)$. The expected future rent (in periods 14 and 15) to a self-regarding worker who does not engage in shirking in period 13 is

$$[40 - c(7) - 5] + [0.6(5 - c(1)) + 0.4(34 - c(6)) - 5] = 33.4. \quad (8.73)$$

The two square brackets capture expected rents in periods 14 and 15 respectively. Analogous to (8.72), the self-regarding worker will shirk in period 13 if

$$c(e_n) > 33.4.$$

Hence, from Table 8.1, the set of effort levels that will not induce shirking is the entire set, i.e., $\{1, 2, \dots, 10\}$. Since additional effort is valuable for the firm, in order to induce an other-regarding worker to supply $e = 10$, the firm must offer a payoff equalizing wage $w_b = \frac{10 \times 10 + c(10)}{2} = 59.0$. It is, therefore, optimal to offer the contract $(w_b, e_n) = (59, 10)$ in period 13. Using backward induction it is optimal to offer this contract also in the first 12 periods.

Since only one-shot interactions take place in the ICR treatment, the maximum effort level $e = 10$ cannot be sustained in this condition. Finally, accepting a wage equal to the outside option in period 15 is rational for an other-regarding worker because there is an oversupply of workers. If some of them are perceived to be self-regarding workers who will accept the outside option, then an other-regarding worker cannot punish the firm by rejecting a low offer. ■

8.4.3 Experimental results

The experimental results are in line with the predictions of the theoretical model (Proposition 8.35). These results are as follows:

1. The relative share of private offers is substantially greater in the ICF treatment relative to the ICR and the C treatments. Furthermore, most of the private offers (63–79% after period 5) are made to workers already working for the firm. Hence, voluntary long-term relations are maintained through private offers when there is an opportunity to do so.
2. As predicted, about 90% of the trades in treatment C took place in interactions that lasted for one or two periods only. This contrasts with the ICF condition where only a third of the trades took place in one-shot encounters. About 50% of the trades in the ICF condition took place in relationships that lasted four or more consecutive periods and about 40% in relationships that lasted eight or more consecutive periods. The separation rate between workers and firms declined significantly after period 3, suggesting the possibility of a search phase for a good match in the initial periods. A comparison of the length of the relationship under the two treatments is shown in Figure 8.3, which plots the cumulative frequency of the length of relations. The cumulative distribution function under treatment ICF first order stochastically dominates the distribution function under treatment C.
3. Under the ICF treatment, firms pay a rent to workers in order to induce them to put in the contracted effort level. Since rents enhance effort, due to the gift exchange mechanism, they can be interpreted as a form of *efficiency wages*. On the other hand, in treatment C, there is no need to pay efficiency wages because workers put in the contracted effort due



Figure 8.3 A comparison of the length of the worker–firm relationship under the two treatments.

Source: Brown et al. (2004) with permission of the Econometric Society.

to third party intervention. In the ICR treatment, on the other hand, there is no future to worry about, hence, efficiency wages are pointless. Thus, a clear prediction of the model is that wages and effort would be highest in the ICF treatment. These predictions are verified in Figures 8.4, 8.5. From Figure 8.4, it can be seen that the average wage is always highest in the ICF treatment except for the first few periods when the wage in the C treatment is higher. Further, the wage gap between the ICF and the ICR treatments increases over time and the gap in wages between the ICF and the other treatments is statistically significant. From Figure 8.5, there is an increase in effort over time in the ICF treatment. Effort drops off towards the end, as predicted in theory. Not surprisingly, effort is highest in treatment C, because it can be costlessly implemented through a third party and additional effort, net of the wage costs, is valuable. Furthermore, as predicted by the theory, the experimental results show that there is contingent contract renewal, i.e., a worker is more likely to have his contract renewed if he undertakes higher effort. For a worker providing the maximum effort level, the probability that the contract will be renewed in the following period is close to 1.

4. In choosing between alternative contractual forms, the incomplete contracts literature often takes the view that the contractual form which maximizes the total surplus in the relationship will be the institutional solution; for instance, the formation of firms in preference to the conduct of transactions at arms-length relationships (Hart, 1995). The experimental results indicate that the joint surplus in a relationship is higher, the longer it lasts. Since the ICF treatment is the one most conducive to long-term relations, it is not surprising that, on average, it has a higher joint surplus relative to the ICR treatment. The highest surplus is in treatment C because the contract can be costlessly enforced. Insofar as contracts cannot be costlessly enforced in the real world, and treatment C is an idealized benchmark case, the predicted institutional solution is the ICF treatment. Figure 8.6 shows the association between profitability and the length of the relationship in the three treatments.




Figure 8.4 Average temporal wage in each treatment.

Source: Brown et al. (2004) with permission of the Econometric Society.




Figure 8.5 Average temporal effort in each treatment.

Source: Brown et al. (2004) with permission of the Econometric Society.

In Figure 8.6, the ICR and the C treatments are represented at the two extremes with the ICF treatment in the middle. Two important features of these results are the following. First, the surplus is shared less equally in the C and the ICR treatments relative to the ICF treatment. Second, the total surplus in the ICF treatment is increasing in the length of the relationship.

The experimental results suggest that initial conditions are the crucial determinant of a relationship developing into a long-term one. Of the initial relationships between workers and firms, 75% do not survive. However, 80% of the relations that survive for at least two periods go on to survive for the full duration of the experiment. Of the relationships that survive, an important factor seems to be the magnitude of the wage. If the initial wage is high, then initial effort is also high and the relation is more likely to survive. There is a gradual increase in wage and effort in relationships that survive longer than 6 periods. One can conjecture that this gradual increase in wage and effort creates greater mutual trust and makes the relationship more enduring.

In Brown et al. (2004) considered above, there is an excess supply of labor. In contrast, Brown et al. (2012) implement an experimental setup with an excess demand for labor. Denote the contracts C and ICF in Brown et al. (2004) by C-S and ICF-S to reflect the excess supply condition. Likewise denote the corresponding contracts under excess demand by C-D and ICF-D; contracts ICR-D and ICR-S are defined analogously. The point of the exercise is to see if the extent of unemployment has a bearing on the formation of relational contracts. Figure 8.7 compares the mean effort levels under the following two comparisons, ICF-D and ICF-S; ICR-D and ICR-S.

The evidence shows that firms are equally likely to offer contract renewal for non-shirking workers under conditions ICF-S and ICF-D. However, in the excess demand condition, workers turn down the contract renewals more often, so they switch firms more often. Hence, there is a lower formation of relational contracts. But the proportion of trades taking place in condition ICF-D is still larger than those taking place under C-D. In this sense, bilateral long-term contracts continue to play an important role even under excess demand. The effort levels under excess supply and excess demand for labor are similar.




Figure 8.6 Average earnings per trade of the workers and firms in each treatment.

Source: Brown et al. (2004) with permission of The Econometric Society.

What other factors are conducive for improving payoffs in finite period relationships? Van Dijk et al. (2002) produce experimental evidence to show that the continuation of relations produces *social ties* between the parties to a relationship. This strengthens their willingness to cooperate in the relationship, hence, providing an additional mechanism, other than instrumental reciprocity, that can sustain cooperative behavior. Gächter and Falk (2002) find that subjects in experiments cooperate more if they know there is a possibility that they will be paired with the same partners in the future.¹⁸ Engelmann and Fischbacher (2009) confirm similar findings by giving a fraction of the subjects an opportunity to develop a reputation. They find that altruistic rewarding and the desire to build a reputation are important considerations for their subjects in experiments.

One argument that has been made in countering the results of cooperative behavior in one-shot games is that subjects in experiments are unable to distinguish between one-shot and repeated interactions. Hence, cooperation is some form of *maladaptation* that results from cognitive failure; for instance, Roth and Erev (1995) and Binmore et al. (1995). We have already considered these arguments, above, in the context of the ultimatum game. We have seen that individuals understand the strategic differences between one-shot and repeated interactions.

8.4.4 The nature of the firm and other-regarding preferences

When should transactions be carried out within a firm and when should they be carried out in arms-length relationships? This question has had a distinguished history in economics, particularly since Coase (1937). Simon (1951) posed this question as the choice between *employment contracts* within the firm and *sales contracts* offered by one entity to another. The advantage of an employment contract is that it offers greater flexibility in decisions to the employer. Within legal bounds, it allows the employer to control more closely the type of task that an employee will undertake. The cost of an employment contract is that employees may be treated more opportunistically by the employer (within legal bounds); this may reduce the intrinsic

¹⁸ Dal Bó (2002) finds similar results for the prisoner's dilemma game. Cooperation increases when the probability of future interaction with the same partner increases from 0.5 to 0.75.




Figure 8.7 Comparison of mean effort levels under various treatments.

Source: "Competition and relational contracts: the role of unemployment as a disciplinary device," Brown, M., Falk, A., and Fehr, E., *Journal of the European Economic Association*, 10(4): 887–907. Copyright © 2012 by the European Economic Association.

motivation of employees and make them reluctant to accept such a contract unless rents are high. The sales contract, by contrast, offers less flexibility to the employer but prevents opportunistic treatment of the employee, making the latter more likely to accept the contract at a lower rent.

Bartling et al. (2013) formalize similar trade-offs between two parties, a principal and an agent. The two parties sign one of two contracts, a sales contract or an employment contract, prior to the realization of one of two equiprobable states of the world—state 1 and state 2. There are three possible tasks, task 1–3, that the principal can assign to the agent. Table 8.2 shows the relevant, state-contingent and task-contingent, economic information. This includes the revenues of the

Table 8.2 The relevant economic information in each state of the world.

		Task 1	Task 2	Task 3
State 1	Revenue	160	80	190
	Cost	0	40	130
	Surplus	160	40	60
	Principal's profit	160–wage	80–wage	190–wage
	Agent's income	wage	wage–40	wage–130
State 2	Revenue	30	200	230
	Cost	0	40	130
	Surplus	30	160	100
	Principal's profit	30–wage	200–wage	230–wage
	Agent's income	wage	wage–40	wage–130

Source: Bartling et al. (2013).

principal, the cost of the agent, the total surplus to be shared between the two parties, the principal's profit, and the agent's income. Task 3 always gives the highest revenues in each state but agents incur very high costs in undertaking it, hence, it is never efficient in any state. In state 1, task 1 is most efficient (total surplus is 160) and in state 2, task 2 is most efficient (total surplus is 160). The outside option of the principal and the agent is respectively 0 and 30.

A *sales contract* takes the form (w_b, task) , where w_b is a binding wage and task specifies one of the three tasks. Thus, before any uncertainty about the state of the world is realized, the agent, in a sales contract, already knows which contractual task has been specified. An important cost of such a contract is that the ex-ante choice of the task may not be ex-post optimal. An *employment contract*, takes the form $(w_b, \text{principal chooses task ex-post})$. Thus, the principal has *residual control rights* on the choice of one of the three tasks after observing the state of the world. Such a contract gives greater flexibility to the principal on task choice, but leaves open the possibility that the principal might choose the more exploitative task 3 and *abuse his legal power*. Thus, this is a problem of *employer moral hazard*. Once the state of the world is realized, the contracts are executed.

The game is played between eight principals and eight agents, over 15 periods. All players are assigned ID numbers at the start of the game. In the first stage, the principal makes a *private offer* that is addressed to a particular agent by specifying the ID number. This is followed by a *public offer* to other agents. In the second stage, agents observe these offers. If an agent receives more than two private offers he can accept only one. Once agents have accepted the private offers, if any, the unmatched agents are randomly matched with the public offers of the principal. In the end of this process, at least one agent is matched with one principal.

There are two main treatments. In treatment RANDOM, the ID numbers are randomly reassigned at the beginning of each period. Thus, when a principal makes contract offers or agents examine the contract offers, they cannot identify the past behavior of the other party. By contrast, in the treatment FIXED, the ID numbers stay fixed over the 15 periods, hence, principals can specify the ID number of the agent they would like to be matched with. In this case, the agent can use his past experience with the principal to make a more informed acceptance decision. In the treatment RANDOM, there is no point for the principal to build a reputation for not abusing authority (say, by not choosing task 3). However, building a reputation is critical in treatment FIXED. A third treatment, UNFAIR, is identical to RANDOM except that with an 80% chance, task 3 is exogenously chosen.

Proposition 8.36 *Suppose that the principal and the agent are self-regarding. The optimal contract for the principal is to offer a sales contract of the form $(w_b, \text{task 2})$. The employment contract is never optimal.*

Proof: Consider first, an employment contract. Suppose that we are in the last period, period 15, in any of the treatments RANDOM, FIXED, or UNFAIR. In an employment contract, since w_b is binding, the self-regarding principal will demand that the agent undertake task 3 because it is revenue maximizing. The cost of the agent is 130 and his outside option is 30. Hence, rational agents will accept an employment contract only if $w_b \geq 160$. These arguments also hold for the particular sales contract, $(w_b, \text{task 3})$. Since both states of the world are equiprobable, so the expected profit of the principal from an employment contract and from the sales contract $(w_b, \text{task 3})$ is $0.5(190 - 160) + 0.5(230 - 160) = 50$.

By contrast, the principal's expected profit from the sales contract (i) $(w_b, \text{task 1})$ is $0.5(160) + 0.5(30) - 30 = 65$, and (ii) $(w_b, \text{task 2})$ is $0.5(80 - 40) + 0.5(200 - 40) - 30 = 70$. Hence, the sales contract $(w_b, \text{task 2})$ is optimal in period 15. We can now solve the game with backward induction.

The sales contract (w_b , task 2) is optimal in every period. In particular, note that in the treatment FIXED, both players know that the principal will act opportunistically in an employment contract in period 15, hence, there is no point for the principal to develop a reputation for non-abusive behavior in period 14. Using backward induction, the same applies, successively, in rounds 13, 12, ..., 1.

Thus, the sales contract (w_b , task 2) that gives the principal an expected profit of 70, dominates an employment contract that offers an expected profit of 50. ■

Suppose now that we allow for a proportion, $\lambda \in (0, 1)$, of self-regarding principals and the remaining proportion $1 - \lambda$ of other-regarding principals who have the Fehr–Schmidt preferences. However, the preferences of each player is private information. This turns the game into a signaling game and one needs to find the perfect Bayesian equilibria of the game, as in Proposition 8.35 above. The presence of other-regarding principals potentially alters the solution because they do not try to exploit their residual control rights in an employment contract by choosing task 3. Unlike Proposition 8.35 that fully characterizes a perfect Bayesian equilibrium, here we offer an informal conjecture that illustrates the potential equilibria in the model.

Conjecture 8.1 *Suppose that there is asymmetric information about the type of the principal; with probability $\lambda \in (0, 1)$, the principal is self-regarding, and with probability $1 - \lambda$, the principal is other-regarding.*

- (a) *In each treatment, and in each period, in one possible perfect Bayesian equilibrium, all principals offer the sales contract (w_b , task 2).*
- (b) *Consider the treatment RANDOM. If a self-regarding agent is sufficiently optimistic that the principal is other-regarding ($\lambda \leq \frac{w-50}{110}$), then he may accept the employment contract. Self-regarding principals always abuse their residual control rights by asking for task 3.*
- (c) *There can be no employment contracts in the treatment UNFAIR.*
- (d) *In the treatment FIXED, we “may” observe employment contracts and long-term contracts between the principal and the agent in which, ex-post, the principal requires the efficient tasks (task 1 in state 1 and task 2 in state 2).*

Basis for Conjecture 8.1: (a) Bayes’ rule does not impose restrictions on out of equilibrium beliefs in such games. So let us impose the extreme beliefs that if the agent observes an employment contract, he assigns the posterior belief that the principal is self-regarding with probability 1. But this means that the agent is sure that the principal will abuse his authority by choosing task 3 in an employment contract. From Proposition 8.36 we know that the sales contract (w_b , task 2) is optimal in every period.

(b) Suppose that both types of principals offer an employment contract in period 15 and a pooling wage w . Agents believe that the self-regarding principal will choose to abuse his authority by choosing the revenue maximizing task 3, while the other-regarding principal will find it optimal to elicit an efficient task (task 1 or task 2). From Table 8.2, the cost of effort for tasks 1 and 2 is identical in states 1 and 2, and task 1 is efficient in state 1 and task 2 in state 2. Thus, a self-regarding agent should accept the employment contract if

$$\lambda [w - 130] + (1 - \lambda) [0.5(w) + 0.5(w - 40)] \geq 30 \Leftrightarrow \lambda \leq \frac{w - 50}{110}. \quad (8.74)$$

Thus, given w , the self-regarding agent should accept the employment contract only if he believes that the probability of encountering a self-regarding principal is low. If the agent is other-

regarding, then we have to proceed as in Proposition 8.36 by allowing extra rents for the inequity-averse agent. In the treatment RANDOM, there is no benefit for a self-regarding principal to try to build a reputation for not abusing residual control rights. Thus, such a principal will always choose to abuse authority. However, for a low enough λ there may be equilibria with employment contracts in each period.

(c) In treatment UNFAIR, task 3 is exogenously imposed with probability 0.8. Thus, even in the best case scenario in which task 1 is chosen in state 1 and task 2 in state 2, the surplus (net of the agent's outside option of 30) in the relation is

$$0.8 [0.5 (190 + 230) - 130] + 0.2 [0.5(160) + 0.5(200 - 40)] - 30 = 66.$$

From Proposition 8.36, we know that this is dominated by a net surplus of 70 from a sales contract (w_b , task 2), hence, employment contracts should never be offered in treatment UNFAIR.

(d) Consider treatment FIXED. The formal proof can be given along the lines of Proposition 8.35. However, note that even self-regarding principals will attempt to pool with other-regarding principals in order to build a reputation for non-exploitative choice of efficient tasks (task 1 in state 1 and task 2 in state 2). The benefit of doing so is that the agent will not update his priors about the type of the principal and accept the principal's private employment offer in the next period. However, if the principal abuses the agent's trust in any period, then the agent rejects the private offer of the principal in the next period and, possibly, updates his posteriors that the type of the principal is self-regarding (this depends on the exact specification of the out of equilibrium beliefs). In the spirit of Kreps and Wilson (1982) and Milgrom and Roberts (1982) (see also Proposition 8.35) principals will try to build a reputation and milk it for a while, then revert back to their preferred contracts in the end. For instance, self-regarding principals will choose task 3 in period 15, irrespective of the treatment and this is foreseen by agents. ■

The empirical results are as follows. In treatment RANDOM, the prediction is that authority will always be abused if players have self-regarding preferences (Conjecture 8.1). Yet, we observe that power is abused only in 51% of the cases (task 3 is chosen) but in the remaining cases one of the efficient tasks, 1 or 2, is chosen. Thus, the purely self-regarding preferences case is rejected.

Figure 8.8 shows the payoffs of the principal and the agent in the sales and employment contracts in treatment RANDOM (left panel) and treatment FIXED (right panel). In each case, on




Figure 8.8 Payoffs of the principal and agent in treatment RANDOM (left panel) and treatment FIXED (right panel).

Source: "Use and abuse of authority: a behavioral foundation of the employment relation," Bartling, B., Fehr, E., and Schmidt, K.M., *Journal of the European Economic Association*, 11(4): 711–41. Copyright © 2013 Journal of the European Economic Association.




Figure 8.9 Share of employment contracts in different treatments (left panel). Frequency of power abuse and share of employment contracts in each treatment (right panel).

Source: "Use and abuse of authority: a behavioral foundation of the employment relation," Bartling, B., Fehr, E., and Schmidt, K.M., *Journal of the European Economic Association*, 11(4): 711–41. Copyright © 2013 Journal of the European Economic Association.

average, the principal earns a higher payoff in the employment contract, relative to the sales contract. In treatment RANDOM, on average, agents receive a higher payoff in sales contracts as compared to employment contracts. This finding is reversed in treatment FIXED and, in an employment contract, agents earn, on average, significantly higher than their outside option of 30 (i.e., they receive positive rents); by contrast agents barely earn their outside option, on average, in an employment contract in treatment RANDOM. In both treatments, principals who abuse their residual control rights in employment contracts earn higher payoffs than those who do not, yet both types of principals earn more relative to a sales contract. Further insights into these results are provided in Figure 8.9.

From the left panel in Figure 8.9, in all treatments, the employment contract is predominantly offered in the beginning. However, over time, agents observe that the self-regarding principals abuse trust by choosing task 3. This, in particular, is the case in treatment RANDOM. Hence, agents who have had an abusive experience, begin in successive rounds to reject the employment contract in favor of the sales contract. In a sales contract, principals typically ask for one of the efficient tasks, task 2. Yet, one still observes that in treatment RANDOM, the share of the employment contract stabilizes at about 40%. What accounts for this high percentage, relative to the predictions of the self-regarding model?

A crucial role is played by the presence of principals with other-regarding preferences. We can see this by comparing the share of employment contracts in the following two treatments: RANDOM and UNFAIR. In the latter, an employment contract is never optimal (Conjecture 8.1(c)). The differences between the two treatments (left panel of Figure 8.9) are statistically significant and employment contracts are almost eliminated in the last few periods in treatment UNFAIR. Thus, it is likely to be the case that in treatment RANDOM, agents expect the presence of a significant fraction of principals with other-regarding preferences.

In contrast to these two treatments, the proportion of employment contracts in treatment FIXED begins from a high level and hovers around the 80–90% mark (there is also an endgame effect). The main reason for the prevalence of the employment contract in treatment FIXED is that principals do not abuse their residual control rights as much, in comparison to the other contracts. This is confirmed from the right panel of Figure 8.9, which shows a statistically significant negative correlation between the share of employment contracts and the abuse of power. This is

the reason why in the right panel of Figure 8.8 (treatment FIXED), on average, the principal and the agent are better off relative to a sales contract. This is despite the fact that the agent has a negative payoff when he faces a principal who abuses his residual control rights.

The main reason that principals abuse their residual control rights relatively less in treatment FIXED is that agents terminate employment relations with abusive principals in the succeeding period. The cost to the principal of this event is that it might lead the principal to be identified as a self-regarding principal (this depends on the precise specification of out of equilibrium beliefs in a formal model). That this will lead to lower subsequent profits for the principal is also confirmed by the data.

Long-term relations do not emerge in a sales contract under the treatment FIXED but do emerge under an employment contract. The reason is that no reputation building is possible under a sales contract as there are no residual control rights that a principal can exploit.

The *efficiency* (actual surplus relative to maximum potential surplus) of employment contracts is significantly greater than sales contracts. This might be an important clue why *vertical integration* takes place in actual practice. The *transaction costs* identified in the traditional industrial organization literature typically originate from contract rigidity and the hold-up problem (Gibbons, 2010). However, the nature of the transaction costs identified in this paper is significantly distinct from the earlier ones. There are three main transaction costs: (i) Some potentially profitable contract offers are rejected by the agents (see also Bartling et al., 2012). (ii) Sales contracts are rigid and do not allow ex-post alteration of tasks. (iii) Some principals do abuse their residual control rights in an employment contract. Indeed, in the treatment FIXED, these transaction costs are lower in an employment contract relative to a sales contract and, hence, efficiency is higher.

Furthermore, these experiments provide exogenous variation in the principal's ability to build a reputation (FIXED versus RANDOM treatments) that alter the costs of integration such as the abuse of power by the principal but not the costs of non-integration (rigidity of task choice under sales contracts). By contrast, existing field evidence provides mainly correlational evidence only (Bartling et al., 2013, p. 715). This should also serve to caution those who might unconditionally espouse the superiority of field evidence relative to lab experiments.

8.5 Extrinsic and intrinsic motivation: theoretical framework

Consider the opening remarks from Ellingsen and Johannesson (2008):

Douglas McGregor's (1960) celebrated management book *The Human Side of Enterprise* argues that managers who subscribe to the conventional economic view that employees dislike work—McGregor labels it Theory X—may create workers who are “resistant, antagonistic, uncooperative”. That is, managerial control and material incentives may trigger the very behaviors that they are designed to avert. Conversely, managers who subscribe to the more optimistic view that employees see their work as a source of self-realization and social esteem—Theory Y—may create workers who voluntarily seek to fulfill the organization's goals.

A major part of the agenda for the theoretical literature, that we consider in this section, is to provide some foundations for Theory Y.

In this section, we consider formal theoretical models that simultaneously examine the effect of incentives on extrinsic and intrinsic motivation. This is an active and productive area of research.¹⁹

8.5.1 Self-signaling and social signaling

In this section, we consider the models due to Bénabou and Tirole (2003, 2006). Consider first, a simplified version of the more general model in Bénabou and Tirole (2003). In this model, incentives signal the amount of trust that a principal has in an agent.

Consider a principal and an agent, both risk neutral, who are engaged in some project. There is asymmetric information about the cost, $c \in [\underline{c}, \bar{c}]$, of the project that the agent will bear if the project is undertaken. The principal knows the actual value of c but the agent only knows the density, f , and the probability distribution of costs, F . For any wage, w , received by the agent and any cost, c , the agent's utility is $w - c$. The agent chooses the effort level, e , to either shirk, $e = 0$, or work hard, $e = 1$. If the project is unsuccessful, then both parties get a payoff of 0. If the agent works hard, then the project is successful with probability $\theta \in (0, 1]$ and the respective payoffs of the principal and the agent are $W > 0$ and $V > 0$; with probability $1 - \theta$, the project is unsuccessful. If the agent shirks, then the project is always unsuccessful. The outside option of the agent is normalized to 0.

The agent does not know the value of c , but he receives a private signal of the cost, c , given by $\sigma \in [0, 1]$ with conditional density $g(\sigma | c)$ and conditional distribution $G(\sigma | c)$. A higher value of the private signal, σ , is *good news* in the sense that it signals a lower value of the cost. This is formally incorporated by using the *monotone likelihood ratio property* (MLRP). For any two values of $\sigma : \sigma_L < \sigma_H$, and cost $c : c_L < c_H$, the MLRP property states that

$$\frac{g(\sigma_H | c_H)}{g(\sigma_L | c_H)} < \frac{g(\sigma_H | c_L)}{g(\sigma_L | c_L)}, \text{ or } \frac{g(\sigma_H | c_H)}{g(\sigma_H | c_L)} < \frac{g(\sigma_L | c_H)}{g(\sigma_L | c_L)}. \quad (8.75)$$

Using Bayes' rule,

$$\Pr(c_L | \sigma_H) = \frac{1}{1 + \frac{g(\sigma_H | c_H) f(c_H)}{g(\sigma_H | c_L) f(c_L)}} \text{ and } \Pr(c_L | \sigma_L) = \frac{1}{1 + \frac{g(\sigma_L | c_H) f(c_H)}{g(\sigma_L | c_L) f(c_L)}},$$

thus,

$$\frac{\Pr(c_L | \sigma_H)}{\Pr(c_L | \sigma_L)} = \frac{1 + \frac{g(\sigma_L | c_H) f(c_H)}{g(\sigma_L | c_L) f(c_L)}}{1 + \frac{g(\sigma_H | c_H) f(c_H)}{g(\sigma_H | c_L) f(c_L)}}. \quad (8.76)$$

Using (8.75), (8.76) we get that $\Pr(c_L | \sigma_H) > \Pr(c_L | \sigma_L)$, which is the precise sense in which a higher value of σ is good news.

Lemma 8.19 *The MLRP in (8.75) implies that*

$$\frac{[1 - G(\sigma_H | c_L)]}{[1 - G(\sigma_L | c_L)]} \geq \frac{[1 - G(\sigma_H | c_H)]}{[1 - G(\sigma_L | c_H)]}.$$

¹⁹ See, for instance, Hwang and Bowles (2011), Ellingsen and Johannesson (2008), Bénabou and Tirole (2003, 2006), and Besley and Ghatak (2005).

The proof of Lemma 8.19 is left as an exercise for the reader.

The principal does not observe the effort level of the agent. If the project is successful, the principal contractually offers a performance based monetary bonus, b , to the agent. This is equivalent to an expected wage offer of $w = \theta b$, contingent on the agent working hard ($e = 1$). The bonus could alternatively be interpreted as the monetary value of working conditions, praise, and friendliness. Thus, in the event that the agent works hard and the project is successful, the respective payoffs of the principal and agent are $W - b$ and $V + b$.

The agent may be willing to work hard if the cost of the project is not too high. The principal benefits if the agent works hard, otherwise each gets a payoff of only 0. Hence, the principal may wish to manipulate incentives in order to signal to the agent that the project has low cost and hope that this will boost the *self-confidence* of the agent and his interest in the task. The agent is aware of the principal's incentives to do so. Thus, he uses, jointly, his private signal, σ , and the bonus payment, b , to make an inference of the cost, c .

The full information benchmark is simple to find. The incentive compatibility condition and the individual rationality conditions are identical in this model. If the agent shirks, then the agent gets zero, which is also the outside option. The risk-neutral agent works hard if the expected payoff exceeds the outside option of zero, i.e., $\theta(V + b) - c \geq 0$. The principal only needs to ensure that this condition binds. Thus, a higher bonus makes it more likely that the agent will work hard.

Now consider the solution under asymmetric information. We derive a perfect Bayesian Nash equilibrium of the model. Under asymmetric information, the agent attempts to infer the cost of the project by observing, jointly, σ and b . Denote by $E[c | \sigma, b]$, the expected level of the cost, conditional on σ, b . Since a higher private signal, σ , is good news about costs, thus, for a fixed b , $E[c | \sigma, b]$ is decreasing in σ .

The payoff of the agent from shirking ($e = 0$) is zero. Hence, the risk-neutral agent works hard ($e = 1$) if the expected payoff from working hard is positive; this is the incentive compatibility condition,

$$\theta(V + b) - E[c | \sigma, b] \geq 0. \quad (8.77)$$

Figure 8.10 gives a diagrammatic representation of the condition in (8.77) for two different levels of bonuses, $b_1 < b_2$. The curve $E[c | \sigma, b]$ is decreasing in σ . The shape is purely for illustrative purposes and the upward shift with respect to the bonus ($b_1 < b_2$) is proved in Proposition 8.37. It is obvious from Figure 8.10 and (8.77) that for any bonus, b , there is a cutoff point $\sigma^*(b)$ such that $E[c | \sigma^*(b), b] = \theta(V + b)$, and

$$\begin{cases} \sigma \geq \sigma^*(b) & \Rightarrow e = 1 \\ \sigma < \sigma^*(b) & \Rightarrow e = 0 \end{cases} \quad (8.78)$$

The probability that $\sigma \geq \sigma^*(b)$ is given by $1 - G(\sigma^*(b))$. Given the best response function of the agent in (8.78), a principal who knows the actual cost to be c , chooses the bonus level, b , to maximize his expected profit

$$E\pi = [1 - G(\sigma^*(b) | c)]\theta(W - b), \quad (8.79)$$

where the expectation is taken over the distribution of σ . One may think of each possible level of cost, $c \in [\underline{c}, \bar{c}]$, as the type of the principal. Thus, an optimal strategy for the principal is a choice




Figure 8.10 Diagrammatic representation of the condition in (8.77).

of b for each $c \in [\underline{c}, \bar{c}]$. Denote the set of optimal values of b by B^* . The results are summarized in the following proposition and we sketch a heuristic proof.

Proposition 8.37

- (i) If $b_1, b_2 \in B^*$ and $b_1 < b_2$, then the cutoff private signal beyond which the agent decides to work hard is relatively lower under b_2 , i.e., $\sigma^*(b_2) < \sigma^*(b_1)$.
- (ii) A higher level of bonus is bad news for the agent. Let b_H and b_L be the optimal levels of bonus when the cost is respectively c_H and c_L , $c_H > c_L$. Then $b_H > b_L$.
- (iii) Higher bonus payments reduce the perceived attractiveness of the task for the agent. For all pairs of signals (σ, σ') , and optimal bonus payments $b_1 < b_2$, $b_1, b_2 \in B^*$, we have $E[c | \sigma, b_1] < E[c | \sigma', b_2]$.

Proof: (i) Suppose that the statement in (a) is false. Then for some $b_1 < b_2$ in B^* we have $\sigma^*(b_2) > \sigma^*(b_1)$. But then, instead of offering b_2 , the principal could offer a lower bonus b_1 . This increases the expected profits of the principal, given in (8.79), in two ways. By reducing σ^* it increases the first term $[1 - G(\sigma^*(b) | c)]$. The second term, $\theta(W - b)$, increases directly when the bonus level falls. This contradicts the original supposition that b_2 is an optimal level of bonus. Thus, we cannot have $\sigma^*(b_2) > \sigma^*(b_1)$.

(ii) Suppose that the bonus b_i , $i = H, L$ is optimal when cost is c_i . Let $\sigma_2 = \sigma^*(b_H)$ and $\sigma_1 = \sigma^*(b_L)$ be the cutoff points for the private signals. Then, we get from (8.79) that

$$[1 - G(\sigma_2 | c_H)] \theta(W - b_H) > [1 - G(\sigma_1 | c_H)] \theta(W - b_L), \quad (8.80)$$

$$[1 - G(\sigma_1 | c_L)] \theta(W - b_L) > [1 - G(\sigma_2 | c_L)] \theta(W - b_H). \quad (8.81)$$

When the actual cost of the principal is c_H the bonus b_H is optimal ((8.80)) and when the actual cost is c_L the bonus b_L is optimal ((8.81)). From (8.80), (8.81) we get

$$\frac{[1 - G(\sigma_1 | c_L)]}{[1 - G(\sigma_2 | c_L)]} > \frac{(W - b_H)}{(W - b_L)} > \frac{[1 - G(\sigma_1 | c_H)]}{[1 - G(\sigma_2 | c_H)]}.$$

Since $c_L < c_H$, Lemma 8.19 implies that $\sigma_1 = \sigma^*(b_L) > \sigma_2 = \sigma^*(b_H)$. Since σ^* is a decreasing function it follows that $b_L \leq b_H$. Thus, a high bonus, b_H , conveys bad news for the agent that costs are high, c_H .

(iii) There cannot be a separating equilibrium. Suppose to the contrary that there is a separating equilibrium. Then each type of the principal, $c \in [\underline{c}, \bar{c}]$, announces a distinct level of bonus. Since the principal's action is fully revealing of his type in a separating equilibrium, hence, there is no need for the agent to use an imperfect private signal, σ . Suppose that we support such an equilibrium by imposing particularly pessimistic out of equilibrium beliefs that if the bonus payment in a separating equilibrium, b_i for some type $c_i \in [\underline{c}, \bar{c}]$ changes to any $b_j \neq b_i$, then the principal is of the worst, or highest cost, type, $\bar{c}$.

But if the agent is going to ignore his private signal, σ , then the principal of type $\bar{c}$, should choose the lowest bonus payment possible and improve his profits. This breaks the separating equilibrium. Thus, the proof of part (b) implies that pooling occurs over intervals. For $b_1 < b_2$, if the principal pays bonus b_1 for costs $c \in [\tilde{c}_0, \tilde{c}_1]$ and bonus b_2 for costs $c \in [\tilde{c}_2, \tilde{c}_3]$, then $\tilde{c}_1 < \tilde{c}_2$. Thus, all cost levels for which a higher bonus is offered are also higher, hence, the expected costs under the higher bonus are also higher, so for all pairs of signals (σ, σ') , $E[c \mid \sigma, b_1] < E[c \mid \sigma', b_2]$. Therefore, higher rewards are given for less attractive tasks (forbidden fruits) and they reduce intrinsic motivation. ■

Bénabou and Tirole (2003) consider intrinsic and extrinsic motivation, Bénabou and Tirole (2006) add to the mix a reputational motivation. For instance, while the Bénabou and Tirole (2003) model deals with situations where one party to a formal or informal contract has private information (child rearing, education, monitoring of employees), a different framework is needed when one makes prosocial decisions such as blood donations, charitable giving, in-kind donations, choosing a fuel efficient car, and voting. Such a framework is provided in Bénabou and Tirole (2006) that we now turn to.

Suppose that an agent can choose some prosocial action a from a set of real-valued numbers, A , at some cost $C(a)$. The action a confers some *intrinsic utility*, u_I , on the agent, per unit of the action, a (e.g., the agent may feel good about the action of donating blood). A principal chooses a monetary reward, $y \geq 0$, per unit of the action, a , for the agent. The *extrinsic utility* derived by the agent from the monetary reward per unit is u_E . For instance, the government (principal) may decide to tax exempt or subsidize the charitable contributions of citizens. Thus, the *private benefit* to the agent from the action, a is

$$P(a, y) = a(u_I + u_E y) - C(a). \quad (8.82)$$

The agent has private information on his *type* or *identity*, which is summarized by $(u_I, u_E) \in R^2$. The other players know that the pair (u_I, u_E) is drawn from the density $f(u_I, u_E)$. Others attempt to learn or infer (u_I, u_E) from the action taken by the agent, a , and the monetary reward given by the principal, y . The reward, y , creates additional difficulties in inferring intrinsic and extrinsic motivations from actions. For instance, society may wish to infer the degree of altruism of a contributor to charity, however, this calculation is made difficult if the tax regime allows charitable contributions to be tax exempt.

Social signaling is incorporated into the model by a concern for reputational payoffs. Denote by $E[u_I \mid a, y]$, $E[u_E \mid a, y]$, respectively, the posterior expectation of others about the agent's intrinsic utility and extrinsic utility. Then the *reputational payoff* is given by

$$R(a, y) = v(E[u_I \mid a, y] - E[u_E \mid a, y]). \quad (8.83)$$

In (8.83), the parameter $v > 0$ represents the *social visibility* of the action a (e.g., the number of people who observe the action) and it captures the extent of *reputational concerns* of the agent. An increase in $E[u_I | a, y]$ improves the reputation payoff of the agent by increasing the social perception of his intrinsic motivation. This term corresponds to the notion of *warm glow* from giving in the literature on charitable donations. On the other hand, an increase in $E[u_E | a, y]$ reduces the reputation payoff of the agent because it increases the social perception of his extrinsic motivation, or greed. In a more general model, the terms $E[u_I | a, y]$, $E[u_E | a, y]$ in (8.83) could be differentially weighted, as in Bénabou and Tirole (2006). Furthermore, there could be individual heterogeneity in these terms which make the problem of inferring motivations from actions even more difficult. We abstract from these considerations here.

The aggregate utility of the agent, U , indexed by the type $(u_I, u_E) \in R^2$, is given by $P(a, y) + R(a, y)$, thus,

$$U = a(u_I + u_E y) - C(a) + v(E[u_I | a, y] - E[u_E | a, y]). \quad (8.84)$$

The agent chooses action a to maximize U , defined in (8.84). Assuming that the problem is well-behaved, the first order condition that is sufficient to determine the optimal action is

$$(u_I + u_E y) + vR_a(a, y) = C'(a), \text{ where } R_a(a, y) = \left(\frac{\partial E[u_I | a, y]}{\partial a} - \frac{\partial E[u_E | a, y]}{\partial a} \right). \quad (8.85)$$

The LHS of (8.85) gives the marginal benefit of increasing a and the RHS, the marginal cost. The marginal benefit has three sources: *intrinsic*, *extrinsic*, and *reputational*. Suppose that we fix a particular value of $a > 0$. Then, from (8.85), the set of people who contribute (called the “high contributors”) satisfy the inequality $(u_I + u_E y) + vR_a \geq C'(a)$, which can alternatively be written as

$$u_E \geq \frac{1}{y} (C'(a) - vR_a(a, y)) - \frac{1}{y} u_I. \quad (8.86)$$

When the inequality in (8.86) binds, the agent chooses a and when it does not bind, the agent chooses an action greater than a . Three different loci of combinations of (u_I, u_E) on which (8.86) binds are shown in Figure 8.11, as I, II, and III; these loci are created by varying the size of rewards, y . In u_I, u_E space, the region above any such locus gives the set of points (u_I, u_E) for which an action greater than a is chosen.

When there are zero rewards ($y = 0$), (8.86) reduces to $u_I^*(a, 0) = C'(a) - vR_a(a, 0)$. This is shown in Figure 8.11 by the vertical line, labeled I. Since there are no extrinsic rewards in this case, the actions of agents do not reveal anything about their extrinsic motivation. All the individuals whose identity or type (u_I, u_E) lies on or to the right of curve I contribute at least an amount a , under zero rewards, while those to the left contribute strictly less than a .

Now suppose that there is a jump in rewards to a strictly positive level ($y > 0$). The slope of the (u_I, u_E) locus is now given by $-\frac{1}{y}$. Two loci are shown in Figure 8.11, corresponding to $y > 0$; these are labeled as II and III. The curve labeled II shows the immediate effect, prior to the adjustments in the reputational concerns arising through the term R_a . Thus, there is no change in the abscissa (see (8.86)) which stays fixed at $u_I^*(a, 0)$. The curve labeled III shows the final effect once the reputational concerns have played out. In the case shown in the figure, reputational




Figure 8.11 Loci of combinations of (u_I, u_E) on which (8.86) binds.

Source: Bénabou and Tirole (2006) with permission from the American Economic Association.

concerns shift the abscissa to the right. The abscissa is now given by $u_I^*(a, y) = C'(a) - vR_a(a, y)$. We now explain these effects in greater detail.

The immediate effect captured by curve II shows the extra set of individuals who now find it worthwhile to contribute at least a , in response to the presence of extrinsic incentives ($y > 0$ rather than $y = 0$). These people (new contributors) are represented by the area $A + B$. This is the classical positive effect of incentives. The reputational effect that is incorporated in curve III is more subtle. The relative intrinsic utility of the “new contributors” must be lower than that of the “high contributors” because the former did not find it worthwhile to contribute in the absence of extrinsic rewards. This reduces the overall average level of intrinsic utility of the set of contributors and the overall reputation of the set of contributors for prosocial behavior. However, the set of non-contributors has now also shrunk because the “new contributors” who have relatively high intrinsic utility among this group have transited to the set of contributors. Hence, the reputation from contributing an amount $a - \Delta a$ also goes down. Thus, the reputational effect for a marginal contributor is ambiguous, hence, the abscissa in (8.86) can go up or down. Curve III is drawn on the supposition that the reputation effect is negative ($R_a(a, y) < R_a(a, 0)$), so the abscissa goes up.

Comparing curves I and III, we are in a position to assess the final effect of incentives. Incentives stimulate contributions from a set of greedy individuals (high u_E relative to u_I). This is shown by the area B in Figure 8.11. However, there is an extra set of marginal contributors shown by the area C . The intrinsic utility u_I of these individuals is higher relative to their extrinsic utility u_E , so (based on (8.86)) they no longer find it optimal to take action a after the introduction of extrinsic incentives. This potentially explains several findings such as the work of Titmuss (1972) in the context of crowding-out effect of incentives for blood donations. Comparing curves I and III, the overall effect on total contributions will depend on the relative size of the areas, B and C as well as the shape of the probability density of individuals, $f(u_I, u_E)$, over these areas.

8.5.2 *Intentions, altruism, and esteem*

In this section, we consider theoretical models of principal–agent relations in which type-based reciprocity plays a key role (Ellingsen and Johannesson, 2008). These are models of two-sided

asymmetric information in which both players can be self-regarding or altruistic. Players hold other players in greater esteem if they are perceived to be altruistic. Furthermore, players value being esteemed more by other players who are perceived to be altruistic. This can explain why, a principal who chooses a kind action in the presence of an unkind action is regarded in greater esteem, because the principal signals clearly to the agent that he is kind. Agents then reciprocate, trusting actions by a kind principal, because the esteem received from a kind principal is valued more highly. Kindness may mean choosing a trusting action when a non-trusting action is available, or choosing a less controlling action when a more controlling action is available. We characterize an intuitive equilibrium of this two-sided signaling game and apply it to explain the experimental findings of McCabe et al. (2003) and Falk and Kosfeld (2006). We also consider the basic intuition behind some of the other related theoretical literature.

Consider two players $i \in N = \{\text{Principal, Agent}\}$. The pure strategy set of player $i \in N$ is denoted by S_i and the mixed strategy set by Σ_i . Consider the class of two-stage games where the principal first chooses a publicly observable strategy from the set Σ_P . The agent then sequentially chooses a strategy from the set Σ_A . There is two-sided asymmetric information about the altruism of the players. Each player belongs to one of two possible, privately known, types $\theta \in \Theta = \{\theta_L, \theta_H\}$, $0 \leq \theta_L < \theta_H < 1$. The prior belief of player $i \in N$ that the opponent is of type θ_H is p_i^0 .

Suppose player i of type θ_i chooses his action $s_i \in S_i$. Let h_i be the history of game, and let the belief of the player that his opponent is of type θ_H be $\mu_i(h_i, \theta_i)$. Let H_i be the set of all possible histories of the game when player i makes his move. Denote by $\bar{h}$, the terminal history of the game; it is summarized by the strategies followed by the two players for the whole game, σ_P, σ_A (defined below). Thus, we may write $\mu_i(\bar{h}, \theta_i) = \mu_i(\sigma_P, \sigma_A, \theta_i)$. We assume that players use Bayes' rule to update their beliefs.

The utility of a risk-neutral player $i \in N$ is given by

$$u_i = y_i + \theta_i y_j + \eta_i, \quad (8.87)$$

where y_i is the monetary payoff of player $i \in N$ and y_j is the monetary payoff of the opponent, player $j \neq i$. The size of θ_i measures the degree of *altruism* that the player feels for the opponents; since we have restricted $\theta_j > 0$, $j = L, H$ we are ruling out spite or envy. The term η_i measures player i 's feelings of being *esteemed* by player j . It is measured as follows. Let θ_{ij} be the esteem with which player i views player j . Conditional on the type of player i , θ_i , this is defined as the expected value of the altruism of player j at the terminal node.

$$\theta_{ij} = E[\theta_j | \theta_i, \bar{h}] = \mu_i(\bar{h}, \theta_i) \theta_H + [1 - \mu_i(\bar{h}, \theta_i)] \theta_L. \quad (8.88)$$

Players hold their opponents in greater esteem if they think they are more altruistic. Thus, there is respect for character virtues.

Let $\chi_j = \chi(\theta_j)$ be the salience of the esteem that an opponent of type θ_j has for player i . Now define η_i as

$$\begin{aligned} \eta_i = & \mu_i(\bar{h}, \theta_i) \chi_H (\text{Expected esteem of a type } \theta_H \text{ opponent for player } i) \\ & + [1 - \mu_i(\bar{h}, \theta_i)] \chi_L (\text{Expected esteem of a type } \theta_L \text{ opponent for player } i) \end{aligned} \quad (8.89)$$

or,

$$\begin{aligned} \eta_i = & \mu_i(\bar{h}, \theta_i) \chi_H \left[\left[\mu_j(\bar{h}, \theta_H) \theta_H + [1 - \mu_j(\bar{h}, \theta_H)] \theta_L \right] \right] \\ & + [1 - \mu_i(\bar{h}, \theta_i)] \chi_L \left[\left[\mu_j(\bar{h}, \theta_L) \theta_H + [1 - \mu_j(\bar{h}, \theta_L)] \theta_L \right] \right]. \end{aligned} \quad (8.90)$$

We are now in a position to evaluate the utility function in (8.87). First, although (8.87) only incorporates altruism (through the second term), the authors claim that richer other-regarding preferences, such as those in Fehr and Schmidt (1999), do not alter their results. Second, the feeling of being esteemed by a player is greater if more virtuous opponents hold them in higher regard. Consider, for instance, the first term on the RHS of (8.89). The expected esteem of a more altruistic player j is weighted by the term $\mu_i(\bar{h}, \theta_i) \chi_H$. Clearly this term is higher if the belief of player i about the altruism of player j , $\mu_i(\bar{h}, \theta_i)$, is higher. Thus, the model is similar in spirit to models of type-based reciprocity, such as that by Levine (1998). Third, an alternative way of dealing with beliefs about esteem and beliefs about beliefs about esteem, require the machinery of psychological game theory. This is discussed in Part 4 of the book.

A strategy for any player $i \in N$ is given by $\sigma_i : \Theta \times [0, 1] \times H_i \rightarrow \Sigma_i$. It specifies a mixed strategy of a player for each possible type of a player, for any beliefs about the opponent, and for any history of the game. Let $\sigma_P^* = (\sigma_P^*(\theta_H), \sigma_P^*(\theta_L))$ and $\sigma_A^* = (\sigma_A^*(\theta_H), \sigma_A^*(\theta_L))$ denote the equilibrium type-contingent strategies of the principal and the agent. Let $\mu_P^* = (\mu_P^*(\theta_H), \mu_P^*(\theta_L))$ and $\mu_A^* = (\mu_A^*(\theta_H), \mu_A^*(\theta_L))$ be the respective equilibrium type-contingent beliefs.

Definition 8.23 (*Intuitive equilibrium*): An equilibrium of the two-stage two-player game is a pair of type-contingent strategies (σ_P^*, σ_A^*) and a pair of type-contingent beliefs (μ_P^*, μ_A^*) such that the following hold.

- (1) Given the beliefs of any player $i \in N$, μ_i^* , the strategy σ_i^* maximizes the utility of player $i = A, P$.
- (2) The beliefs μ_i^* of player $i \in N$ satisfy Bayes' rule whenever possible.
- (3) Off the equilibrium path, beliefs must satisfy the Cho–Kreps intuitive criterion.

Consider the trust game of McCabe et al. (2003) shown in Figure 8.12. In the voluntary trust game (VTG) in Panel A, player 1 (Principal) has two choices: Trust (T) or No Trust (NT). In the involuntary trust game (ITG) in Panel B, player 1 has only one choice, T. Player 2 (Agent) has identical choices in both cases: Reward (R) and Not Reward (N). In classical game theory, the choice of player 2 should be identical in both cases. However, McCabe et al. (2003) find that the frequency of the choice of R for player 2 is significantly higher in the VTG as compared to the ITG. Thus, the principal's choice of trust conveys a better signal of his sociality when the NT option is available.

Let us now explain these experimental results using the theoretical framework developed in this section. Assume that each player has two possible types $\theta \in \Theta = \{\theta_L, \theta_H\}$ such that θ_H is the relatively more altruistic type. The solution is set out in Proposition 8.38. The intuition behind the results is as follows. First, in the VTG, altruistic principals (high θ_H) are more likely to play T because they place a higher value on the payoffs of the agent that arise from playing T rather than NT. Second, in a separating equilibrium where the more altruistic principal plays T but the self-regarding principal plays NT, the action of the principal fully reveals the type. More altruistic agents feel greater esteem from playing the action R, when they know that they are dealing with




Figure 8.12 A trust game that differentiates between voluntary and involuntary trust.
Source: Ellingsen and Johannesson (2008) with permission from the American Economic Association.

altruistic principals (type-based reciprocity). Hence, there is likely to be a separating equilibrium in which the more altruistic agent plays R and the self-regarding agent plays N in the VTG game.

Proposition 8.38 (Ellingsen and Johannesson, 2008): *Consider the trust games shown in Figure 8.12. For an open set of parameters, there exists a unique intuitive equilibrium (Definition 8.23) such that:*

- (i) *An altruistic principal (type θ_H) plays T and a self-regarding principal (type θ_L) plays NT.*
- (ii) *Both types of agents play N in the ITG. Altruistic agents (type θ_H) play R in the VTG.*

Proof: We only prove the existence part. For the proof of uniqueness see Ellingsen and Johannesson (2008).

Consider first the ITG. Both types of player 2s must choose N in this game as part of the stated equilibrium in the proposition. A sufficient condition is that the altruistic player 2 who benefits most by gaining self-esteem prefers to play N over R. Player 1 never gets a chance to update his beliefs because he moves first. Hence, the relevant beliefs for him are captured by his initial belief that player 2 is of type θ_H (denoted by p_1^0). Similarly, since player 1 has no choice but to play T, player 2 does not update his beliefs either, hence, the belief of player 2 that player 1 is of type θ_H is given by p_2^0 .

Using (8.87), (8.88), and (8.90), the utility of player 2 of type θ_H from playing N under the worst possible evaluation from player 1 (i.e., $\mu_1(\bar{h}, \theta_H) = p_1^0 = 0$) is

$$u_2(N, \theta_H, p_1^0 = 0) = 30 + 15\theta_H + \theta_L (p_2^0 \chi_H + (1 - p_2^0) \chi_L).$$

Similarly, the utility of a type θ_H player 2 from playing R under the best possible evaluation from player 1 (i.e., $\mu_1(\bar{h}, \theta_H) = p_1^0 = 1$) is

$$u_2(R, \theta_H, p_1^0 = 1) = 25 + 25\theta_H + \theta_H (p_2^0 \chi_H + (1 - p_2^0) \chi_L).$$

Under these conditions,

$$u_2(N, \theta_H, p_1^0 = 0) > u_2(R, \theta_H, p_1^0 = 1) \Leftrightarrow p_2^0 \chi_H + (1 - p_2^0) \chi_L < \frac{5 - 10\theta_H}{\theta_H - \theta_L}. \quad (8.91)$$

Now consider the VTG. As part of the separating equilibrium, player 1 of type θ_H always plays T in this case, while player 1 of type θ_L always plays NT. The prior beliefs of player 2 are updated in a manner that is consistent with Bayes' rule. Thus, on observing T, player 2 sets $\mu_2(\bar{h}, \theta_H) = 1$, and on observing NT he sets $\mu_2(\bar{h}, \theta_H) = 0$. So, conditional on being given the move $(\mu_2(\bar{h}, \theta_H) = 1)$, we first need to show that the altruistic player 2 will prefer to play R.

$$u_2(R, \theta_H, \mu_2 = 1) = 25 + 25\theta_H + \chi_H\theta_H.$$

In computing the last term, note that consistent with the strategies of player 2, player 1 sets $\mu_1(\bar{h}, \theta_H) = 1$ in (8.90). Similarly, the payoff of the altruistic player 2 from playing N is

$$u_2(N, \theta_H, \mu_2 = 1) = 30 + 15\theta_H + \chi_H\theta_L. \quad (8.92)$$

If player 2 plays N, then, given the strategies, he is revealed to be a type θ_L player, hence, the beliefs of player 1 about him, at the terminal node are $\mu_1(\bar{h}, \theta_H) = 0$. Using this fact in (8.90) gives the last term on the RHS of (8.92), $\chi_H\theta_L$. Thus, the required condition for an altruistic agent to play R in the VTG game is

$$u_2(R, \theta_H, \mu_2 = 1) > u_2(N, \theta_H, \mu_2 = 1) \Leftrightarrow \chi_H > \frac{5 - 10\theta_H}{\theta_H - \theta_L}. \quad (8.93)$$

Now consider a self-regarding agent (type θ_L) in the VTG. Proceeding analogously to the derivations above, such a player prefers to play N over R (as required in the equilibrium strategy profile) if

$$u_2(R, \theta_L, \mu_2 = 1) < u_2(N, \theta_L, \mu_2 = 1) \Leftrightarrow 25 + 25\theta_L + \chi_H\theta_H < 30 + 15\theta_L + \chi_H\theta_L,$$

or, if

$$\chi_H < \frac{5 - 10\theta_L}{\theta_H - \theta_L}. \quad (8.94)$$

Now consider the incentive compatibility conditions of player 1 in the VTG. Player 1 believes that with probability p_1^0 he faces a type θ_H player 2 who will play R, and with probability $1 - p_1^0$, a type θ_L player 2 who will play N. Hence, using (8.87), (8.88), and (8.90), the utility of a type θ_H player 1 from playing T is

$$u_1(T, \theta_H, p_1^0) = p_1^0 [25 + 25\theta_H + \chi_H\theta_H] + (1 - p_1^0) [15 + 30\theta_H + \chi_L\theta_H].$$

Since we are constructing a separating equilibrium, in the first term on the RHS, each of the players assigns probability one that the other is altruistic. Hence, we get the term $\chi_H\theta_H$. In the

second term on the RHS, player 1 assigns probability zero that player 2 is type θ_H ($\mu_1(\bar{h}, \theta_H) = 0$), while player 2 assigns probability 1 that player 1 is type θ_H ($\mu_2(\bar{h}, \theta_L) = 1$); substituting these facts in (8.90) gives the term $\chi_L \theta_H$. Similarly, we can define

$$u_1(NT, \theta_H, p_1^0) = p_1^0 [20 + 20\theta_H + \chi_H \theta_L] + (1 - p_1^0) [20 + 20\theta_H + \chi_L \theta_L].$$

Thus,

$$u_1(T, \theta_H, p_1^0) > u_1(NT, \theta_H, p_1^0) \Leftrightarrow p_1^0 > \frac{5 - 10\theta_H - \chi_L (\theta_H - \theta_L)}{10 - 5\theta_H + (\theta_H - \theta_L) (\chi_H - \chi_L)}. \quad (8.95)$$

By now, the method of proof should be transparent to the reader. Proceeding analogously, a self-regarding player 1 (type θ_L) prefers to play NT (as stated in the equilibrium profile in the proposition) if

$$u_1(T, \theta_L, p_1^0) < u_1(NT, \theta_L, p_1^0) \Leftrightarrow p_1^0 < \frac{5 - 10\theta_L - \chi_L (\theta_H - \theta_L)}{10 - 5\theta_L + (\theta_H - \theta_L) (\chi_H - \chi_L)}. \quad (8.96)$$

We require that the inequalities in (8.93), (8.94), (8.95), and (8.96) hold simultaneously. One set of parameter values for which this is true is $p_1^0 = \frac{2}{5}$, $\chi_L = \theta_L = 0$, $\theta_H = \frac{1}{3}$ and $\chi_H = 6$. Since this statement is robust to small changes in the parameter values, it holds for an open set of parameter values around the stated values. ■

We now consider the effects of *control-aversion* on the part of agents in a principal-agent relation. By choosing incentives that reveal a controlling nature, the agent's effort level may be reduced even when neoclassical theory predicts no such effect.

Consider the Falk and Kosfeld (2006) study in which a principal (player 1) can choose either to control or not control an agent (player 2). The agent is given an endowment of 120 money units and he can transfer any amount $[0, 120]$ back to the principal. For every unit of transfer made by the agent, the experimenter ensures that the principal receives 2 units, hence, the amount received by the principal lies in the interval $[0, 240]$.

Prior to the agent's decision, the principal has two options. (i) Either specify that the agent transfer a minimum amount $\bar{x} = 10$ (so that the principal receives at least $2\bar{x} = 20$). (ii) Or do not specify a minimum transfer. When a minimum transfer $\bar{x}$ is specified, the agent transfers an average amount of 17.5, while in the absence of such minimum transfers it is 23. About 70% of the principals choose to trust the agent by not specifying a minimum transfer. Further, if a minimum amount is exogenously specified, say, by the experimenter, then there is no difference in the amounts transferred. Hence, control-aversion seems to be the most likely explanation of these findings.

Consider now an explanation of these experimental results based on the Ellingsen and Johannesson (2008) framework. The strategy of the agent is $s_A \in [0, 120]$ if the principal chooses to trust (T), and $s_A \in [10, 120]$ if the principal chooses to control (C). The strategy of the principal is a binary choice, $s_P \in \{T, C\}$. Thus, the utility functions of the principal and the agent are, respectively

$$u_P(s_P, \theta_i) = 2s_A + \theta_i(120 - s_A) + \eta_P, \quad (8.97)$$

$$u_A(s_A, \theta_i) = (120 - s_A) + \theta_i s_A + \eta_A. \quad (8.98)$$

As before, we assume that each player belongs to one of two possible, privately known, types $0 \leq \theta_L < \theta_H < 1$. The next proposition describes a separating equilibrium in this game that illustrates the negative effects of control-aversion. We leave the proof as an exercise.

Proposition 8.39 (Ellingsen and Johannesson, 2008): *There exists an open set of parameters such that there are intuitive equilibrium outcomes of the Falk and Kosfeld (2006) game with the following features:*

- (i) *An altruistic principal (type θ_H) plays T, while a self-regarding principal (type θ_L) plays C.*
- (ii) *A self-regarding agent (type θ_L) transfers the lowest amount possible, i.e., 0, when the principal plays T, and 10 when the principal plays C. An altruistic agent (type θ_H) transfers an amount $y \in (10, 120)$. However, the altruistic agent always transfers a greater amount if the principal plays T rather than C.*

Admittedly the results in Propositions 8.38, 8.39 hold for a simple class of games. However, they powerfully illustrate a view of incentives that runs contrary to existing models in neoclassical contract theory, in which incentives are typically, universally, good. These results qualify the existing models in important ways. If a principal mistrusts an agent, when he has an unused alternative that clearly signals trust, then agents may respond with lower effort levels. Furthermore, agents do not like being controlled—controlling them when there is an unused alternative that clearly signals mistrust also reduces effort.

Bénabou and Tirole (2006) also show (see Section 8.5.1), in a model where individuals have private information about multidimensional characteristics, that incentives may undermine effort. As in the blood donation example given above, incentives create an adverse selection problem by attracting individuals who are purely materialistic. Hence, it becomes difficult to signal one's esteem in the presence of such incentives. However, the Ellingsen and Johannesson (2008) model makes the additional prediction that when principals have a choice between two kinds of incentives—one signaling trust and the other signaling mistrust then the former elicits greater effort.

These results are related to those in Sliwka (2007). Here, in addition to the usual two kinds of *steadfast agents*, self-regarding and other-regarding, there is a third type, the *conformists*. The human virtues of conformists are more malleable and less steadfast. They behave like other-regarding agents if they think that there is a sufficiently large number of other-regarding agents. In other words, they are influenced by social norms. Thus, such an agent feels remorse from harming someone if he believes that there are sufficiently many other types of agents who would feel remorse.

The principal can choose either to trust or control the agent. The overall distribution of types of agents is unknown to the agents but the principal has superior information (perhaps he is more experienced). By choosing to trust, when control is available, the principal signals to the agents that the proportion of other-regarding agents is high. This is the signal that the conformists need in order to work hard. Such a transmission mechanism can also be invoked if the principal has a choice between say, a fixed wage scheme and an incentive scheme; the former may be used to signal to the conformists that the frequency of other-regarding agents is high.

Herold (2010) considers a principal-agent relation in which there are two types of agents. A *trustworthy agent* who shares the principal's concern to produce a high output level and a *self-regarding agent* who cares purely for the wage received and the cost of effort. The types of the agents are privately known to them. The principal receives a private signal of the type of the agent. Based on this signal, the principal has two possible types. The trusting principal has a

high signal of the trustworthiness of the agent. The distrusting principal has a low signal of the trustworthiness of the agent. Conditional on the signal, each type offers either a trust contract that is independent of the realization of output (flat wage contract) or an incentive contract in which the wage is conditional on the output level. The principal can signal his type by a choice of the contract—trust or incentive contract.

The effort level of the trusting agent is positively influenced by his belief that the principal is trusting. Furthermore, the principal knows that the trusting agent shares his mission, and requires a lower cost in terms of incentives to produce the same output level. This effect is enhanced if the production function is such that the output depends on the joint efforts of the agent and the principal, and these effort levels are complementary. Hence, a trusting principal gains by separating from a mistrusting principal. In a separating equilibrium, the trusting principal chooses a trust contract and the distrusting principal chooses an incentive contract. An incentive contract, in this framework, is bad news for the trustworthy agent and leads to lower effort levels and output levels.

8.6 Extrinsic and intrinsic motivation: empirical evidence

Motivated by the theoretical models in Section 8.5, we now consider the empirical evidence on intrinsic and extrinsic motivation.

8.6.1 Incentives may signal a hostile intent

In principal–agent interaction, an agent might perceive the following two contractual situations in entirely different ways. In the first, a hostile incentive, say, a fine, is available to a principal and is used. In the second, the hostile incentive is available, but not used by the principal; see, for instance, Section 8.5.2. Empirical tests of these issues are considered by Fehr and Gächter (2001). They modify the basic gift exchange game in Fehr et al. (1997), in the following three potential ways. First, when agents exert inadequate effort (i.e., the actual effort e_a is lower than the non-enforceable, but desired, effort, e_n), they are caught with a probability $p = 0.3$ (in the standard gift exchange game, $p = 1$). Second, in the event that an agent is caught, a fine $f \in [0, f_{\max}]$ may be levied. Third, the contract may specify a *non-binding* bonus, $b_n \in [0, b_{\max}]$, in addition to the binding wage, w , and the non-binding effort, e_n ; thus, this contract is even more incomplete relative to Fehr et al. (1997).

In order to differentiate the effects of various incentives, the following three conditions are considered. (1) *No incentives*, which is similar to the basic gift exchange game. (2) *Negatively framed incentives*, when a fine is possible. (3) *Positively framed incentives*, when a bonus is possible.

The main experimental results are summarized in Figure 8.13. The vertical axis measures the actual effort level, e_a , while the horizontal axis measures rents offered to the worker in different treatments. For a fixed level of rents, the effort levels are highest when no incentives are given. Instead of appealing to the instrumental reciprocity of agents (the incentive treatments) principals can do better by relying on the intrinsic reciprocity of agents (the no-incentive treatment).

One possible explanation of the results is to invoke the role of intentions, as in Section 8.5.2, or to explain the results using psychological game theory (see Part 4 of the book). The no-incentive treatments indicate a trusting principal, thus, agents reciprocate the trust. The incentive treatments on the other hand indicate a lack of trust. In particular, fines might be interpreted as hostile intent. Accordingly, agents may respond with negative reciprocity by putting in a lower level




Figure 8.13 Effort and rents in three cases.

Source: Reprinted from *European Economic Review*, 46: 687–724. Fehr, E. and Falk, A., “Psychological foundations of incentives.” Copyright © 2012, with permission from Elsevier.

of effort. Furthermore, it seems that the framing of incentives is quite important. Effort levels are higher in the positive frame (bonuses) relative to the negative frame (fines). These issues cannot be explained by using self-regarding preferences.

We quote here from Fehr and Gächter (2001). “it may be the case that explicit performance incentives reduce the willingness to voluntarily cooperate.” This possibility may arise because explicit performance incentives may create an atmosphere of threat and distrust. Bewley (1999) reports that “many managers stress that explicit punishment should be rarely used as a way to obtain cooperation ... In particular, explicit incentives may destroy trust—and reciprocity-based incentives, hence, may lead to welfare loss.”²⁰

Fehr and Rockenbach (2003) extend the trust game in the following way. Consider a static game played between an investor and a responder. Each of the players is endowed with 10 MUs (monetary units) and each MU is equivalent to half a German mark. The investor can decide to send (or invest) i MUs, that belong to the set $\{0, 1, 2, \dots, 10\}$ to a responder, via the experimenter. The experimenter triples the amount and passes it on to the responder. The responder then returns r MUs back to the investor. The respective payoffs of the investor and the responder at the end of the game are $10 - i + r$ and $10 + 3i - r$. The investor can also announce some desired transfer, r^* , back from the responder.

The investor can decide to impose a fine on the responder if $r < r^*$. The fine reduces the responder's payoff but does not increase the investor's payoff. Hence, we have three possible cases. In the first, no fine is available. In the second, it is available and is used; this possibly signals unkind intentions to the responder. In the third, it is available and is not used, which signals kind intentions to the responder. Fehr and Rockenbach (2003) find that the total surplus between the investor and the responder and the average payoff of the investor are highest in the third case; see Figure 8.14. Hence, the responder's perception of the principal's intentions (hostile or kind) determines the outcome; kinder intentions produce better outcomes.²¹

²⁰ There is a large psychological literature which suggests that incentives might crowd out intrinsic motivation. Within the Fehr–Schmidt model, intrinsic motivation may be interpreted as reciprocity driven effort.

²¹ There is some evidence that the amounts given in the trust game might be influenced by demographic variables such as age, marital status, health status, etc.; for instance, Fehr and Schmidt (2006). The elderly and the more healthy tend to give back more in the role of a responder.




Figure 8.14 Effect of intentions in the trust game.

Source: Reprinted by permission from Macmillan Publishers Ltd.: *Nature*, Fehr, E. and Rockenbach, B., "Detrimental effects of sanctions on human altruism," 422: 134–40. Copyright © (2003).

In experiments conducted by Dickinson and Villeval (2008), a principal chooses different probabilities of monitoring the agent's effort; these probabilities are observed by the agent. When the effort of the agent directly increases the principal's payoff, then an increase in the monitoring intensity is found to reduce the agent's effort. Similar findings about the potentially negative effect of incentives on effort are reported in other papers that we describe in greater detail elsewhere in this chapter (Fehr et al., 2007; Fehr and Schmidt, 2007).

8.6.2 Control aversion and choking under high incentives

Incentives may signal a desire on the part of principals to "control" agents. Agents may be *averse to control*. This idea goes back to a literature in psychology (Deci and Ryan, 1985; Deci et al., 1999). In this view, incentives conflict directly with the desire of agents for *autonomy* and *independence*. Consider the following example from Ryan and Deci (2000): "Students who are overly controlled not only lose initiative but also learn less well, especially when learning is complex or requires conceptual, creative processing . . . Similarly, studies show children of parents who are more autonomy supportive to be more mastery oriented—more likely to spontaneously explore and extend themselves—than children of parents who are more controlling."

In Falk and Kosfeld (2006), an agent can choose some action $x \in [0, \bar{x}]$ for a principal. Increasing x is costly to the agent but beneficial to the principal. There are two treatments. In the first, the agent's choice is unrestricted—call this a *trust contract*. In the second treatment, the principal sets a lower limit, $\underline{x}$, to the agent's action. Thus, effectively, the agent now must choose x from the interval $(\underline{x}, \bar{x}]$ —call such a contract an *incentive contract*. If the agent had purely self-regarding preferences, then in the trust contract he would choose $x = 0$ and in the incentive contract he would choose $x > \underline{x}$. It turns out that for low levels of $\underline{x}$, the principal's profits under the trust




Figure 8.15 Median choice of action by the agent over when the principal does, and does not, control.

Source: Falk and Kosfeld (2006) with permission from the American Economic Association.

contract are greater than the profits under the incentive contract. Only when $\bar{x}$ is high enough does the incentive contract do as well as the trust contract.

Falk and Kosfeld (2006, p. 1611) summarize their main result as follows: “Explicit incentives backfire and performance is lower if the principal controls, compared to if he trusts.” Their main transmission mechanism (p. 1611–2) is that “agents do not like to be restricted and perceive control as a negative signal of distrust.” Consistent with this hypothesis, principals who try to control the behavior of agents, relative to those who do not, are found to hold more pessimistic beliefs about the performance of the agent.

In order to test this mechanism further, they implement a control treatment in which the experimenter (but not the principal) directly enforces x from the interval $(\underline{x}, \bar{x}]$, where $\underline{x}$ is identical to that chosen in the incentive contract. The control treatment does not signal a lack of trust or the desire to control on part of the principal, because the choice is made by the experimenter. Consistent with the authors’ suggested transmission mechanism, the principal’s profits are higher in the control treatment relative to the profits in the incentive contract.

In a second control treatment, the authors embed their game in a gift exchange game in which the principal chooses the pair $(w, \underline{x})$ where w is a fixed wage. The results are shown in Figure 8.15. As expected, the agent chooses a higher x when w is higher; this is the gift exchange channel. Furthermore, for most wage levels (30 or above in Figure 8.15), the choice of x is lower in the contract $(w, \underline{x})$, $\underline{x} > 0$ (labeled “principal controls” in Figure 8.15) relative to the contract $(w, 0)$ (labeled “principal does not control” in Figure 8.15), which is consistent with the control aversion hypothesis.

The relation between incentives and effort in economics relies on two sequential channels. In the first, performance-based pay increases the worker’s motivation. In the second channel, improved motivation results in greater effort. Ariely et al. (2009a) test for the second channel. They rely on several factors, which indicate that an increase in incentives may have a detrimental effect on effort (p. 2): “These include increased arousal, shifting mental processes from ‘automatic’ to ‘controlled’, narrowing of attention, and preoccupation with the reward itself.” The Yerkes–Dodson law in psychology, based on experiments on rats, posits an optimum level of arousal for each task, with effort falling as the level of arousal differs from the optimal in both directions. The results of Neiss (1988) for human motor performance are consistent with this law. Thus, it is

possible that high incentives over-arouse the subject, focus attention away from the task to the reward, or highlight negative emotions, such as feelings of regret, in case the reward is missed.

Choking refers to emotional responses that make high incentives counterproductive for the sorts of reasons described above (Loewenstein, 2000). The meaning of the term is well understood in common usage. Consider, for instance, the following typically self-effacing quote from Simon Hattenstone in the *Guardian* dated June 22, 2012, prior to the London Olympics:

Britain is no stranger to the choke. Reading the newspapers, or overhearing pub conversations, you might well imagine it's a national pastime. The England football team? Ach, we'll crack up when it comes to penalties. Murray at Wimbledon? Wait till it comes to the crunch. The Olympics? More tears from Paula Radcliffe. Of course, this is an unfair generalization. All those cited have performed at the highest level, and Britain has produced any number of champions. Yet it's undoubtedly true that in a summer in which so many will be playing for the highest stakes, many of the great sporting hopes, from whatever country, will buckle under the pressure.

Ariely et al. (2009a) give six tasks to subjects in an experiment that test memory, motor skills, and creativity. Subjects were randomly assigned to the tasks. There are three treatments, differentiated by the level of incentives: Incentives are very small, medium, or very large. The 6 tasks are simple games. A game called "Packing quarters" measures creativity; this involved fitting nine metal pieces into a wooden frame in a fixed amount of time. Two games, "Simon" and "Recall last three digits" test for memory; these involve recalling a sequence of numbers. Motor skills are tested using three games: "Labyrinth," "Dart Ball," and "Roll-Up." The results are shown in Figure 8.16.

The measure of performance that is used in these six tasks is share of earnings relative to maximum possible earnings. Figure 8.16 shows the average performance in each of the games (Panel(a) pools across all games). It is obvious that relatively high monetary incentives can have perverse effects on performance, but there was no noticeable difference between the low and the medium incentive conditions.

These results used a between-subjects treatment in rural India. A potential criticism of these results is that the subjects might not have played these kinds of games before. Hence, in a second set of experiments, a within-subjects treatment is implemented for MIT students, using games that the subjects would be familiar with. Subjects perform two possible tasks—adding a fixed set of numbers within a certain amount of time that measures cognitive skills, and repeatedly pushing a small sequence of keys to proxy for a physical task. The results are shown in Figure 8.17. The results of the cognitive task are similar to the earlier, between-subjects, treatment and high incentives are detrimental to effort. However, the results of the physical task are the opposite; high incentives are effective in this case.

Bartling et al. (2012) propose a new framework to test for the trade-off between control aversion caused by limited discretion on the one hand, and the misuse of discretion by workers on the other. Suppose that an employer hires an employee by offering a contract (w_b, e_n) , where w_b is the binding wage level and e_n is the non-binding effort level (or *requested effort* that cannot be contractually enforced). The payoffs of the employer and the employee are, respectively, $\pi = be_a - w_b$ and $u = w_b - e_a$, where e_a is the actual effort level, and $b > 0$ is the marginal productivity of effort.

The employer can offer one of two possible contracts. In the *full discretion contract* (contract-F), the employee chooses $e_a \in \{1, 2, \dots, 10\}$ and the productivity parameter $b = 5$. By contrast, in the *limited discretion contract* (contract-L) the employee chooses $e_a \in \{3, 4, \dots, 10\}$ and $b = 4$.




Figure 8.16 Mean earnings as a percentage of maximum earnings in different games.

Source: Ariely, D., Gneezy, U., Loewenstein, G., and Mazar, N. "Large stakes and big mistakes," *Review of Economic Studies* (2009) 76(2): 451–69, by permission of Oxford University Press.




Figure 8.17 Mean earnings as a percentage of maximum earnings in different games.

Source: Ariely, D., Gneezy, U., Loewenstein, G., and Mazar, N. "Large stakes and big mistakes," *Review of Economic Studies* (2009) 76(2): 451–69, by permission of Oxford University Press.

These numbers illustrate the trade-off between *efficiency* and *control*; contract-F is more efficient because b is higher, but the employee may shirk. The reason b is lower under contract-L is that the employee's flexibility is reduced; for instance, control may be associated with inflexible hours that do not suit the employee. The payoff functions and the technology is known to all players and the outside option of the employee is zero.

The employers and employees are matched randomly over 15 periods. There are two treatments. In the *base treatment*, the employer does not receive any information about the past effort levels of the agent. In the *screening treatment*, the employer is told about the effort choice of the employee, e_a , in the last three rounds but not the type of contracts (contract-F or contract-L) that were offered or the requested effort level, e_n .²²

Consider the solution under self-regarding preferences. In the last round, the employee puts in the minimum effort possible under each of the two contracts ($e = 1$ under contract-F and $e = 3$ under contract-L). The employer offers a wage to ensure that the individual rationality constraint of the employee binds, so $w - e = 0$. It is straightforward to check that the profit of the employer under contracts F and L is, respectively, 4 and 9. Thus, the employer optimally chooses contract-L. By backward induction, this is the solution in each round and it is identical under the base treatment and the screening treatment.

The solution under other-regarding preferences of the employer and the employee is along the lines discussed in Section 8.4; see for instance, Proposition 8.35. In this case, employers (self-regarding or other-regarding) will attempt to tap into the reciprocity of other-regarding agents. Self-regarding employees will tend to mimic the actions of the other-regarding employees in order to get rents. The results depend on the proportion of self-regarding and other-regarding employees, and the magnitudes of the inequity parameters of other-regarding players (as in Proposition 8.35).

While we do not prove these results here, under reasonable restrictions, one gets the following results.²³ In the base treatment, employers choose contract-L, i.e., they choose to control. In the screening treatment, however, employers offer contracts conditional on the previous levels of efforts of the employees (screening). When the previous effort levels are high, the employers offer contract-F and a high wage level in order to tap into the employees' reciprocity. For low previous effort levels, they offer the contract-L. The opportunity to build a reputation also induces employees to put in higher effort levels. Such an incentive is even higher if employers compete for workers, putting more upward pressure on wages.

In principle, one could have highly variable contracts. For instance, contract-F or contract-L, any positive wage level, and any, feasible, requested effort level. However, predominantly, two types of jobs are offered by the employers that we may call, respectively, as the *trust strategy* and the *control strategy*. The trust strategy has the following four fold features: (contract-F, high w_b , high e_n , high employee surplus). By contrast, the control strategy has the following four fold features: (contract-L, low w_b , low e_n , low employee surplus).²⁴ The trust strategy is based on tapping into the other-regarding employees' reciprocity, while the control strategy is designed to minimize shirking by self-regarding employees.

²² For instance, employers may have access to reference letters of potential employees that may not provide all possible information.

²³ Readers can consult Section A1 of the online appendix in Bartling et al. (2012) for the details.

²⁴ The employee surplus is defined as $\min \left\{ 1, \frac{w_b - e_n}{b e_n - e_n} \right\}$.

Which of the two strategies, trust and control is more profitable? The answer, as shown in the upper row of Figure 8.18, depends on the reputation of the employees. Let us group employees into three groups by their average past effort levels, $\bar{e}$, over three rounds: *Low reputation* if $\bar{e} < 3.5$, *medium reputation* if $3.5 \leq \bar{e} < 6.5$, and *high reputation* if $6.5 \leq \bar{e}$.

In the base treatment, contract-L is significantly more profitable than contract-F for employers, conditional on workers of all three types of reputations. In the base treatment, employers who offer contract-F, make losses for all levels of wages. The outcome is identical in the screening treatment when workers are of low reputation. However, the opposite is true for workers of medium and high reputation. In this case, under a screening contract, contract-F is relatively more profitable. The main reason for these results can be seen in the second row of Figure 8.18. The slope of the effort–wage response is too flat in the base treatment to justify the cost of a higher wage payment. By contrast, in the screening treatment, the effort–wage response is sufficiently steep to justify higher wage payments. Indeed, in an attempt to build a reputation, even self-regarding employees may mimic other-regarding employees and respond with a high effort choice. Regression results confirm the insights from Figure 8.18.

Recall the results of Falk and Kosfeld (2006) and Ariely et al. (2009a), above. These papers show that greater control reduces effort. This finding is not replicated in the second row of Figure 8.18. In fact, other than the high reputation employees in the screening treatment, one gets the opposite result, namely, that average effort is higher in contract-L. However, the high reputation employees also receive higher wages in a screening contract. Hence, these results neither support nor reject the *hidden costs of control*.




Figure 8.18 Average profits and average effort levels conditional on the employees' reputations.

Source: Bartling et al. (2012) with permission from the American Economic Association.




Figure 8.19 Share of contract-F in each treatment.

Source: Bartling et al. (2012) with permission from the American Economic Association.

It turns out that most employers in the two treatments converge towards the optimal contract over time. In other words, when the reputation of the employees is medium or high, employers choose to control in the base treatment and trust in the screening treatment. The results are shown in Figure 8.19. However, not all employers optimize. There is greater optimization in the base treatment. For instance, in the screening treatment, 45% of the employers fail to optimize when faced with employees of medium to high reputations.

A similar trend is found with respect to the employees. A significant percentage of self-regarding employees do not attempt to mimic the other-regarding employees by responding to high wages with high effort levels. Hence, they forgo an opportunity to improve their payoffs. This leads to *segmented labor markets*. Some employees work hard, are offered good jobs (the trust strategy of the employers) and have higher payoffs relative to others who shirk and get offered bad jobs (the control strategy). Despite the presence of a positive fraction of these non-optimizing employees, the opportunity to screen workers increases the profits of the employers. The reason is that there is a sufficiently large fraction of employees who take the opportunity of screening to build reputations for hard work and they are rewarded by good jobs. These features of the results are strengthened even further when competition is introduced in a separate treatment.

8.6.3 Moral disengagement

Implicit incentives such as bonuses may engage the *moral frame*, while *explicit incentives* may sometimes turn off the moral frame. Bowles and Polanía-Reys (2012, p. 373) write: “market-like incentives trigger what psychologists term ‘moral disengagement’ (Bandura, 1991), a process that occurs because ‘people can switch their ethicality on and off’ (Shu et al. 2011).” In a similar vein,

Schotter et al. (1996, p. 38) write: “While we are not claiming that people change their nature when they function in markets, it may be that the competition inherent in markets and the need to survive offers justifications for actions that, in isolation, would be unjustifiable . . . The morality of economic agents embedded in a market context may therefore be quite different from their morality in isolation.”

Hoffman et al. (2008) consider two treatments. In the first treatment, subjects are randomly assigned to the roles of proposers and responders in a standard ultimatum and dictator game. In the second treatment, subjects have to earn the right to be proposers. This is achieved by ranking the subjects, based on the number of correct answers to a general knowledge quiz. The top half of the ranked students is then assigned the role of a proposer and the bottom half is assigned the role of a responder. Further, proposers are told that they are “sellers” and responders are told that they are “buyers.”

In this “exchange version” of the ultimatum game, subjects play the game shown in Figure 8.20. Note that this game is identical to a standard ultimatum game, although it is framed differently, hence, standard frame-invariant economic theory predicts that the results should be identical in the two games. However, in the exchange version of the ultimatum game, sellers (or proposers) make lower offers relative to the standard ultimatum game and this difference is statistically significant. Buyers (or responders) are also less likely to reject offers in the exchange version of the game. Similar results carry over to the dictator game when a standard dictator game is compared with its exchange version: dictators make significantly lower offers. One plausible interpretation of these results is that the “moral frame” of the ultimatum and the dictator games is altered when these are framed as exchange games.

In Schotter et al. (1996), subjects in the experiment play two rounds of an ultimatum game and subjects are split into three groups. The first group forms the set of proposers in round 1. The second and third groups are, respectively, the responders in rounds 1 and 2. The proposers in round 1 play a tournament in the sense that they are ranked by the amounts that they could keep for themselves successfully (i.e., in cases where their offer was accepted). The top half of the proposers in round 1, who win the tournament, then go on to play the ultimatum game in round 2. Thus, as in Hoffman et al. (2008), proposers have to earn the right to play in round 2. The competitive tournament environment produces lower offers by proposers in round 2, i.e., the presence of incentives induces more self-regarding outcomes.

Cardenas et al. (2000) conduct experiments in three rural areas in Colombia among communities that live on the edge of forests. As part of their lives, these communities harvest a common resource, wood, from the forest. Subjects in the experiments engage in a common pool resource

		Seller chooses											
		PRICE											
Buyer chooses to	BUY	\$0	\$1	\$2	\$3	\$4	\$5	\$6	\$7	\$8	\$9	\$10	
		\$0	\$1	\$2	\$3	\$4	\$5	\$6	\$7	\$8	\$9	\$10	Seller profit
		\$10	\$9	\$8	\$7	\$6	\$5	\$4	\$3	\$2	\$1	\$0	Buyer profit
	NOT BUY	\$0	\$0	\$0	\$0	\$0	\$0	\$0	\$0	\$0	\$0	\$0	Seller profit
		\$0	\$0	\$0	\$0	\$0	\$0	\$0	\$0	\$0	\$0	\$0	Buyer profit

Figure 8.20 Buyer/seller profits as a function of buyer/seller choices.

Source: Reprinted from Hoffman, E., McCabe, K., and Smith, V., “Preferences and property rights in ultimatum and dictator games,” in Charles R. Plott and Vernon L. Smith (eds.) *Handbook of Experimental Economics Results, Volume 1*. Copyright © 2008, with permission from Elsevier.

game. As is typical in these problems, the efficient solution involves a degree of resource conservation, while in the Nash solution there is an over-exploitation of resources on account of free-riding.

The authors consider two groups of subjects. All subjects first play nine rounds of the common pool resource game. As expected, there is over-exploitation of the common resource. In the subsequent nine rounds, one group of subjects can communicate (Com), while the other cannot (Reg). Further, group Reg is subjected to government incentives to prevent over-exploitation in the form of a quota. Communication in group Com is designed to mimic the real-world alternative to regulation in such problems, namely, local cooperative efforts.²⁵ It was found that group Reg, who faced government incentives made choices close to the Nash equilibrium, while group Com made much more efficient choices.

Henrich et al. (2010) consider results for three games, the dictator game (DG), the ultimatum game, and the third party punishment game (TPG). The TPG is a three player game in which the first two players play a dictator game, while the third player can punish the proposer in the dictator game for low offers, by levying a fine at a personal cost. The subjects in the experiments are drawn from 15 diverse societies that include Amazonian tribes, African hunter-gatherers, and US students. Fines were not uncommon. Across all societies, the average offer in DG is 37%, while in the TPG it is 32%. In 4 out of the 6 societies where the offers in DG and the TPG are significantly different in a statistical sense, the offers are significantly lower in the TPG. Thus, the incentives embodied in the fines actually reduce offers. This crowding-out effect of incentives disappeared if one restricted attention to offers made by dictators who subscribed to one of the world's religions such as Islam or Christianity (as compared to other local or traditional religions).

Gneezy and Rustichini (2000a) explore the implications of introducing a fine in private day care centers in Haifa. They choose 10 day care centers that are observed over 20 weeks. In the first 4 weeks, the number of parents who came late to pick up their children was observed for all 10 centers. In 6 out of the 10 centers, a fine of 10 Israeli Shekels for picking up children late was announced in week 5. The remaining 4 centers were the control group that continued to operate without a fine. The 6 centers where the fine was levied between periods 5–16, had the fine removed in the beginning of week 17. Parents were observed again for the remaining 4 weeks to see if incentives have long-lasting effects.

The results are shown in Figure 8.21. There was an immediate jump in the number of parents arriving late after the introduction of the fine. The *reciprocity relation* between the parents and day care center staff seems to have altered with the introduction of the fine to an *exchange relation* (a form of switching-off of ethicality). Those parents who had a high opportunity cost of their time, seemed to be willing to pay the fine to pick up their children late (although this was not formally tested). There is no reason to believe that the opportunity cost altered after the imposition of the fine. It was also found that from weeks 17–20, when the fine was removed, the parents in day care centers where the fine had been levied, continued to come late and their late arrival stabilized at the end of the period. Thus, incentives may have a long-lasting effect on reciprocity relations. Similar results are found by Holmås et al. (2010) who find that disincentives in the form of fines for prolonging hospital stay in Norway had the opposite effect.

One possible explanation for the negative effect of incentives is that incentives may be too low. On the other hand, higher incentives may have a positive effect. This line of reasoning is pursued in Gneezy and Rustichini (2000b).²⁶ In the first experiment, the IQ experiment, university

²⁵ See, for instance, Ledyard (1995) in the context of public goods.

²⁶ See also Heyman and Ariely (2004) who find that giving no incentives is better than giving small incentives; the latter reduce motivation and effort.




Figure 8.21 Average number of parents who arrive late.

Source: Gneezy, U. and Rustichini, A., 2000a "A fine is a price," *Journal of Legal Studies*, 29: 1–17, © 2000. The University of Chicago Press, with permission of the University of Chicago Press.

subjects with an average age of 23 are asked GMAT style questions. Four treatments, with 40 students each, are implemented that differ in the degree of incentives offered; the respective rewards for a correct answer in these treatments are 0, 0.1, 1, and 3 New Israeli Shekels (NIS). In a second experiment, the donation experiment, students were grouped in twos and had to collect money for a good cause. There were three treatments here. In treatment 1a, students were given no monetary incentives but were told that the results of their efforts would be made public and that society wishes for them to collect as much money as possible. In treatment 1b, students were promised an incentive of 1% of the amount they collected and in treatment 1c, an incentive of 10% of the money collected.

The results are as follows. In the IQ experiment, the number of correct answers in the four treatments is respectively 28.4, 23.075, 34.7, 34.1. The number of correct answers with low incentives (monetary incentives 0, 0.1) is significantly lower as compared to high incentives (monetary incentives 1, 3). There is a significant marginal reduction in the number of correct answers when one moves from zero incentives (intrinsic motivation) to very low incentives (0.1). There is no statistical difference in the number of correct answers when incentives are already high (1, 3).

In the donations experiment, the average collections in the three treatments, 1a, 1b, and 1c, respectively, is NIS 238.67, 153.67, 219.33. There is a statistically significant reduction in collections when one goes from treatment 1a (intrinsic motivation) to treatment 1b (1% monetary payments). There is a statistically significant increase in collections as one moves from low incentives in treatment 1b to higher incentives in treatment 1c (10% monetary payments). However, even with high incentives, the collections are not higher than those that arise from purely intrinsic motivation.

Frey and Oberholzer-Gee (1997) showed that Swiss residents were more likely to agree to nuclear waste disposal in their communities out of a sense of civic duty as compared to a situation where they were offered monetary compensation. The monetary compensation may also potentially indicate to the residents that the potential risks of the plant are high.

In an unusual and interesting experiment, Falk and Szech (2013) explore if markets erode morality when subjects could choose between earning money and saving the life of a mouse. In the individual (or non-market) treatment, subjects had to choose between receiving 10 euros,

or saving the life of a young, healthy mouse, who could live for two more years. One may presume that in most cultures, the unjustified and intentional taking of a life must be considered immoral. In the bilateral treatment, a seller and a buyer bargain over the division of 20 euros. Should bargaining conclude successfully, then the life of the mouse is lost, but the two players get to share 20 euros. The multilateral treatment is similar to the bilateral treatment except that there are several buyers and sellers who simultaneously bid in a double auction format.

The authors conjecture three possible channels through which morality may be eroded in markets. First, in the bilateral and the multilateral treatments, responsibility for killing the mouse is more diffused relative to the individual treatment. Hence, guilt may be shared, and shared even more widely in the multilateral treatment, relative to the bilateral treatment. Second, the observation that others are willing to engage in trade (at the moral cost of the mouse's life) may loosen one's own morality in favor of the self-interest motive. Third, markets may draw attention to a non-moral frame by focusing the attention of individuals on negotiations, bargaining, and competition.

In the individual treatment, 45.9% of the subjects are willing to accept 10 euros, rather than save the mouse's life. In the bilateral and the multilateral treatments, 72.2% and 75.9%, respectively, of the subjects are willing to enter into a bargaining agreement that potentially gives them 10 euros, rather than save the mouse's life. These percentages are significantly higher relative to the individual treatment. Further, in the individual treatment, 72.2% and 75.9% of the subjects are willing to have the mouse killed, but only in return for much higher sums of money, respectively, 47.50 and 50 euros. Hence, morality would appear to decay in markets. It would be of interest to examine the precise channels through which these effects operate.²⁷

Ellman and Pezanis-Christou (2010) use insights from social psychology to explore the effect on the ethicality of decisions by varying the organizational form within firms and the levels of communication among players. In their three-player variant of the dictator game, a third player is passive. The first two players, who are members of the firm, decide whether to take an action that improves the profits of the firm, but harms the third player. In the *vertical organizational form*, the two-player firm has one boss who makes the decision and the other player is a subordinate. The subordinate cannot take a decision but can quit, which shuts down the firm. In the *horizontal organizational form*, both players in the firm take a decision by consensus.

When responsibility for the decision is diffused, such as in the horizontal organizational form, then each of the players feels a lower level of responsibility for an unethical decision. Indeed, the authors find that with unrestricted communication between the two players in the firm, vertical organizational forms exhibit more ethical behavior. The authors also recover an important role for communication. Communication may enhance ethicality for several reasons. For instance, communication may aid in transmitting norms of ethicality. Subordinates in horizontal structures may also feel a greater sense of social responsibility when they are able to transmit their views to superiors. Communication may also create an *observer effect*, whereby when one is observed by others one exhibits greater sociality (recall the evidence from dictator games in Section 5.2).

8.6.4 *Morality may be difficult to signal in the presence of incentives*

Consider the work of Bénabou and Tirole (2006) in Section 8.5.1 that addresses issues of altruism and social participation in several diverse contexts such as charitable contributions, voting, and

²⁷ Falk and Szech (2014) already make headway in this direction by considering alternative treatments that differ in the extent to which any one subject is pivotal in the decision to kill the mouse.

blood donation. Positive actions by an individual, such as higher charitable contributions, may improve one's own self-image (self-signaling) and could also be a medium through which one engages in social signaling to others such as family, friends, relatives, and colleagues. Individuals may be privately informed of their degree of altruism and they may derive utility from positive signals of their altruism to themselves and to others. Incentives may, however, dilute the reputational value of the signal by making it noisier. Hence, incentives may be counterproductive in eliciting greater effort.

In a pioneering study, Titmuss (1971) found that individuals are more likely to donate blood on a voluntary basis as compared to the case where they are compensated to donate blood; this became known as the *crowding-out effect of incentives*. Mellström and Johannesson (2008) use Swedish subjects to test the crowding-out effect, taking account of gender differences among subjects. Three treatments are considered. In the first, there are no incentives for giving blood and in the second there is a fixed monetary incentive of SEK 50 (the equivalent of US\$7) for donating blood. The results are shown in Figure 8.22.

Across all subjects in experiments, as we move from treatment 1 to treatment 2, the percentage of subjects who decide to donate blood falls from 43% to 33% but this difference is not statistically significant. However, there are noticeable gender differences. There is no crowding-out for men, but for women there is a statistically significant crowding-out effect and blood donations fall from 52% to 30%. A third treatment offers the kind of incentives that the *Economic Journal* gives to its referees. Subjects can choose to receive SEK 50 for their blood donations or have SEK 50 contributed, on their behalf, to a Swedish charity. In this case, when individuals choose the option to contribute to charity, there is no crowding out. This result is consistent with the Bénabou and Tirole (2006) explanation because contributions to charity may allow one to use contributions to signal one's type.

Lacetera et al. (2012) consider a very large sample of 500,000 blood donations to the American Red Cross (ARC) in Ohio between 2006 and 2008. The ARC collects blood in so called *blood drives*; each blood drive is a single event that solicits blood donations and there are an average




Figure 8.22 Percentage of blood donors under varying levels of incentives.

Source: "Crowding out in blood donation: was Titmuss right?" Mellström, C. and Johannesson, M., *Journal of the European Economic Association*, 6: 845–63. Copyright © 2008 European Economic Association.

of 36 donors per blood drive. They find a positive effect of non-monetary incentives (blankets, T-shirts, mugs, coupons) on blood donations. Higher incentives induce greater blood donations on two accounts: An increase in the turnout and an increase in the number of units donated. A significant part of this increase in donations, about 45%, is caused by substitution from other spatially and temporally separated blood drives. Thus, one should exercise caution in estimating the overall effect of incentives on blood donations.

Reich et al. (2006) find no effect of incentives, in the form of T-shirts, on blood donation. Goette and Stutzer (2008) find a positive effect of incentives that take the form of a free lottery ticket worth USD 5, although the bulk of the effect came from occasional donors, while committed donors exhibited no effect. In their review, Goette et al. (2010) opine that non-monetary incentives are more likely to induce greater blood donations, which is consistent with the signaling hypothesis. Furthermore, while survey evidence suggests that health related incentives (e.g., free cholesterol test for donors) are most effective, randomized trials show that they are not. In general, the evidence does not allow one to make conclusive policy implications. As Goette et al. (2010) say: "Overall, the results provide no support for the predictions from theories of motivational crowding. They provide moderate support for the view that incentives help to tip the balance in favor of more blood donations. However, several studies also found no effects, and this type of heterogeneity remains to be explained."

Angrist and Lavy (2009) consider incentives given to high school grade 12 students in Israel (Bagrut certification) to improve grades. Successful students go on to university, while the unsuccessful ones have an uncertain future. Hence, the economic gains for the students from doing well in these examinations would seem to be very high. However, anti-dropout measures for children taking these exams and measures to improve their grades have not been very successful. Given the widespread trend towards offering financial incentives to high school students for better grade performance, the authors wonder if financial incentives could work. They find that incentives are effective in the case of girls but not boys. Furthermore, incentives work only for those girls whose predicted grades are higher relative to the average. These marginal students would seem to have the most to gain from putting in extra effort or adjust their test-taking strategy. Five years after the intervention, these girls were also more likely to enroll in higher education, suggesting that incentives may have longer term effects.

Charitable donations may be publicly revealed or not. Under private information, donations are anonymous and carry no signal about the goodness or the altruism of the donor. Public information can take several forms, such as publicly revealing the list of donors, or giving donors special privileges, or simply giving tokens such as wristbands that publicly identify the donors. Ariely et al. (2009b) conjecture that in the private information case, incentives for donations (e.g., tax exempt donations) may enhance donations. However, in the public information case, if donations are made to signal one's image to others (or to oneself), then incentives deteriorate such signals. In this case, incentives may be counterproductive.

This conjecture is verified in their empirical results. Figure 8.23 shows donations to either a good cause (as determined by an individual) or to the Red Cross, which was determined to be a good cause by a majority of the subjects. Hence, the latter is particularly revealing for the signaling value of one's donations in the public information case. The vertical axis is the number of presses; subjects were asked to press buttons and the experimenter made donations to the chosen charity that were directly proportional to the number of presses. In each case, incentives reduce prosocial behavior in the public information case but increased such behavior in the private information case.




Figure 8.23 The effect of incentives on prosocial behavior in different treatments. Vertical bars are standard errors of means.

Source: Ariely et al. (2009b) with permission from the European Economic Association.

Carpenter and Myers (2010) consider the decision to volunteer as a firefighter and the decision to respond to a call. They find that subjects who are more altruistic, as identified by their performance in a dictator game, are more likely to be volunteer firefighters and also put in more training hours. However, altruism is not necessarily associated with the decision to respond to calls. The decision to respond to calls is more strongly correlated with incentives that boost one's public image such as a vanity license plate. For decision makers who are motivated by such reputational concerns, extrinsic motivation in the form of monetary incentives are ineffective.

Gneezy et al. (2012) conduct three field experiments to examine the effect of incentives on the degree of self-regarding behavior in markets. These experiments are similar in nature to the Prucker and Sausgruber (2013) experiments that we considered in Chapter 7. In the first experiment, an amusement park sells photos taken during a ride to potential customers. Potential customers can *pay the price that they want* (PWYW pricing). Under purely self-regarding preferences they should not pay at all. However, 8.39% of the customers choose to buy under PWYW pricing. In a second treatment (PWYW + charity), customers are told that half of the sales proceeds will go towards charitable contributions. In this case 4.49% buy their photos but they pay a price \$5.33 relative to a price of \$0.92 in under PWYW pricing. Indeed, the PWYW and PWYW + charity pricing scheme are profitable to the amusement park. The data supports the hypothesis that people buy less under PWYW + charity because they avoid the possibility that they will pay too little and harm their self-image.²⁸

In another field experiment, customers in a restaurant in Vienna could directly, and publicly, pay the owner or, under PWYW pricing, pay privately in an envelope. Under self-regarding preferences, customers should not pay under PWYW pricing, at least in a static game. The direct payment to the restaurant owner rules out any prosocial signaling. However, under PWYW pricing, one may self-signal prosocial behavior, or other desirable human virtues, by paying-up. Indeed, under PWYW pricing, customers paid more relative to direct public payments to the owner. A major conclusion of this study is that the desire to maintain a favorable self-image induces departures from self-regarding outcomes in markets.

²⁸ For a similar experiment and results see also Gneezy et al. (2010).

8.6.5 *Economic environment, learning, and incentives*

We have reviewed the empirical evidence in Henrich et al. (2001) in Chapter 5, which showed that participation in a market environment enhanced sociality among individuals from diverse cultures. Perhaps individuals learn to alter their preferences in such environments because doing so is advantageous to them. Furthermore, the effects of incentives can be long lasting; some examples of this are as follows. The association between educational incentives and long-term enrollment in higher education (Angrist and Lavy, 2009); the link between fines and lateness in the arrival of parents in after care centers in Israel (Gneezy and Rustichini, 2000a). We now consider some additional, related results.

Bohnet et al. (2001) consider the effects of the legal environment on contracts. They consider a two-player game in which a principal decides whether to enter or not into a contract with an agent; this is a measure of the *trust* of the principal. The principal does not know if the agent will perform; the probability that the agent performs is a measure of his *trustworthiness*. If the agent does not perform, then he is penalized with some probability that depends on the strength of the law enforcement institutions. The agent may experience psychological costs (such as guilt, or aversion to lying) from not performing or he may be completely self-regarding. In classical enforcement theory, formulated under purely self-regarding preferences, there is a monotonic relation between the strength of law enforcement and the trustworthiness of the agent. However, in the presence of human virtues, and malleable preferences, this monotonic relation may not hold.

Clearly, when the enforcement institutions are weak, the principal will only offer a contract if his belief about the trustworthiness of the agent is high. A low belief in the trustworthiness of the agent protects the principal from being exploited under these circumstances. A high degree of belief in trustworthiness may ensure that, over time, more trustworthy agents are hired (the authors formalize this in terms of an evolutionary argument). Hence, in an environment with weak enforcement, trustworthiness may be crowded-in. With medium enforcement, the probability of law enforcement is high enough so that it is worthwhile for the principal to enter into a contract, based on expected utility calculations. In this case, the self-regarding agents will maximize their payoffs by not performing and the honest agents will maximize their payoff by performing. Thus, with medium enforcement, trustworthiness may be crowded-out. When the enforcement institutions are strong, then non-performance is always detected and all types of agents perform. The predictions on the relation between the strength of the legal institutions and the degree of trustworthiness are verified in the experiments conducted by the authors.

In a field study, Meier (2007) examines the effect of temporary incentives and the long-term effect on behavior once the incentives are removed. His subject pool is Swiss students who can contribute to one of two possible social funds. It is announced in a treatment group that the contributions of those who contribute to both funds will be matched to varying degrees (in one sub-treatment, the match is 25% and in another it is 50%). It is found that incentives in the form of matching contributions increase donations to the social funds. However, once the incentive is removed, (1) the average contributions in the treatment group are significantly lower than the control group in which no incentives were given, and (2) most of the subjects in the treatment group who had originally contributed to one of the two social funds, stop contributing altogether to any fund.

Reeson and Tisdell (2008) run three treatments of a public goods game without punishments, over 12 rounds. Rounds 1–4 comprise the initial stage, rounds 5–8 the middle stage, and rounds




Figure 8.24 Total public goods contributions over successive rounds when the incentives are varied.

Source: Reprinted from *Journal of Economic Behavior and Organization* 68(1), Andrew F. Reeson and John G. T. Sdell, "Institutions, motivations and public goods: an experiments test of motivational crowding," 273–81. Copyright (2008), with permission from Elsevier.

9–12 the final stage. The control treatment is the standard public goods game experiment run over 12 rounds. In the *moral suasion treatment*, a standard text is read out to the subjects, which says that everyone will be better off if they contributed to the public good. This text was read out at the beginning of each stage, i.e., in rounds 1, 5, and 9. Finally the *regulation treatment* is identical to the moral suasion treatment except that in the middle stage, i.e., rounds 5–8, in addition to moral suasion, the experimenter announces a minimum required contribution of 100 cents to the public good. The results are shown in Figure 8.24 that shows the average contributions to the public good in cents in the three treatments.

The Nash equilibrium that is consistent with free-riding is constructed in such a way that it is 600 cents. Contributions are significantly higher in the moral suasion treatment relative to the control treatment in the initial stage. Consistent with the dominant finding in public goods games, the level of contributions declines over successive rounds in the initial stage. Reading out the moral suasion text at the beginning of round 5 boosts contributions in the middle stage to a level similar to that of the initial stage (the authors call this the “ground-hog day” effect). In the middle stage, contributions are highest in the regulation treatment and these are significantly higher than the control treatment. However, once the regulations are removed in round 9, contributions in the regulation treatment drop below the level of the moral suasion treatment and are significantly lower relative to the initial stage. Thus, not only does moral suasion play an important role, it dominates incentives, in the form of regulation, in the long run.

Burks et al. (2009) consider the effect of economic incentives on Swiss and American bicycle couriers. When distinguished by their compensation, bicycle messengers belong to two main groups. (1) The first group receives performance based pay or commission payment in which

couriers typically receive 40–45% of the revenues they generate. (2) The second group is made up of two subgroups that do not receive performance-based pay. They either receive a fixed hourly wage or they are members of cooperatives that simply share the total revenues.

Members of both groups are asked to play a sequential prisoner's dilemma game, which really is a standard prisoner's dilemma game except that the first player makes a choice that is observed by the second player who then sequentially makes his choice. In this game, respondents are asked to choose their strategy (cooperate or defect) in the role of the first player and the second player.²⁹ They are also asked what percentage of the other players are likely to cooperate. Each respondent is then randomly paired with another, their roles as the first player and the second player are also randomly chosen. The stated responses of both players in their respective roles are then implemented.

The outcomes are used to classify respondents into one of three groups. *Egotists* always defect, *altruists* always cooperate, and *conditional cooperators* are second players who cooperate when the first player plays cooperate and defect when the first player plays defect.³⁰ The same game is also played with students. The results are as follows.

On average, player 1s who played the strategy “cooperate,” estimated that 67% of players 2 will also cooperate. This was an underestimate—the actual figure was 86%. Thus, players may have other-regarding preferences. Respondents who are exposed to performance-based pay, exhibit significantly lower levels of the strategy “cooperate” relative to those who are not exposed to performance-based pay. They also expect others to cooperate less. The experimental classification of respondents into egotists, altruists, and conditional cooperators is a good predictor of the answers to survey questions that further probe their behavior at work and off work. There are significantly more egotists in the student population relative to the population of bicycle couriers.

8.6.6 *Crowding-in effect of incentives*

It is possible that, in certain cases, there might be a synergy between incentives and other-regarding preferences. For instance, punishments by one's peers may trigger feelings of guilt and shame, thus pushing individuals towards a better social outcome. It can often, however, be difficult to isolate the channels that are involved. For instance, punishment in public goods games may contribute to achieving an outcome close to the first best; see Section 5.4.3 above. However, it is also found that non-monetary punishments such as social ostracization can enhance contributions in public goods games (Masclot et al., 2003). Do non-contributors contribute more in the presence of punishment simply because they anticipate being punished and so they contribute, or because the punishment activates feeling of guilt, shame, lower self-evaluation, or even stigma if punishments are publicly implemented?

Herrmann et al. (2008), which was considered in Section 5.4.6, show that in public goods games with punishments, cooperation is higher when there is already a strong social norm of cooperation. In the presence of such norms, punishments trigger feelings of shame rather than resentment against the punishers. When such social norms are absent, then there is a crowding-out effect of punishment. In the absence of social norms of contributing, non-contributors who

²⁹ The *strategy method* is used to enhance the number of observations. In this method, the second player is asked what choice they would make for each possible observed choice of the first player (cooperate or defect).

³⁰ There is a fourth category of players, who the authors term as “Wingnuts” who do not appear to have understood the game. Since these are only 3% of the total, they are excluded from the sample.

have been punished in the past turn on the contributors in subsequent rounds by engaging in antisocial punishments.³¹

Xiao and Houser (2011) ask if the negative effect of incentives arise because the punishments are privately implemented. They consider a public goods game in which the punishments are publicly announced. Yet the anonymity of the target of the punishment is maintained in order to rule out any effects arising from “naming and shaming” of the non-cooperators. They announce a probability with which any round of the public goods game is monitored. In the monitored round, the individual whose contribution is the lowest is punished in proportion to the difference between the average contribution and the lowest contribution. They conjecture that the transmission channel through which incentives work is that public punishments help establish the norm of “contribute at least as much as the average of one’s group members.” Contributions are found to be much higher in the treatment with public punishment relative to the baseline case of private punishment.

Carpenter et al. (2009) reconsider the public goods game with punishments. In the strangers treatment, subjects are randomly assigned to a new group at the beginning of each round. They find that being punished seems to trigger feelings of shame and guilt among non-contributors, who then contribute more. The efficacy of punishments depends on there being a high enough proportion of players who exhibit indirect reciprocity and a group size that is not too large.

Cappelen et al. (2015) consider if giving behavior in dictator games is motivated by intrinsic motivation or extrinsic motivation. In different treatments, the authors vary the morality associated with giving, for instance, by varying entitlements and needs; differences in giving between these treatments reflect changes in intrinsic morality in different moral frames. Extrinsic motivation is altered by varying the information given to the receiver about the dictator’s giving behavior. Such motivation is caused by factors such as shame and guilt that, however, are better captured using psychological game theory that we develop in Part 4. The two main experimental findings are as follows. First, intrinsic motivation, which reflects human virtues, is fundamentally important. Second, intrinsic motivation crowds-in extrinsic motivation; in the absence of intrinsic motivation, extrinsic motivation is irrelevant.

Tyran and Field (2006) consider two kinds of sanctions on non-contributors in a public goods game. Sanctions imposed by the experimenter and sanctions imposed by a referendum among group members. Although both kinds of sanctions raised the cost of free-riding, sanctions imposed by the experimenter were not effective, while those imposed by a referendum were effective. One possible explanation is that the referendum conveys a signal of the moral disapproval by one’s peers, so it is more effective in triggering feelings of guilt and shame.³²

Fryer (2011) conducted experiments in four school districts in the US in which an amount of USD 6.3 million was distributed as incentive payments in 261 schools. In Dallas and Columbia, incentives were given for “inputs” such as attendance, good behavior, and reading a book, followed by a quiz to test the understanding of students. In New York City and Chicago, incentive payments were made for “outputs” such as exam grades. If students understand the production

³¹ Evidence for antisocial punishments are also found in Hopfensitz and Reuben (2009) using a social dilemma game such as the sequential prisoners’ dilemma and the trust game. They find that punishment is effective only when it induces feelings of shame and guilt among the punished.

³² For a theoretical model of the formation of a sanctioning institution (e.g., a court, or an arbitration board) that enables punishment of non-cooperators in a public goods game see Kosfeld et al. (2009). Institutional punishments may also evoke feelings of guilt and shame.

function that links inputs to outputs and if behavioral factors such as self-control are absent, then incentives based on outputs may be superior. The reason is that incentives based on inputs may distort the use of inputs, while output-based incentives allow users to optimally combine inputs. The results of the study are not easy to interpret. However, no link is found between intrinsic motivation and incentives. An important finding of the experiments is that incentives for output are ineffective, while incentives for inputs are effective, in a statistical sense, for some measures of performance in some of the districts. This suggests that we need to reconsider the underlying model of incentives in neoclassical economics.

Levitt et al. (2012) consider the effect of short-term educational incentives that are not confounded by issues of time discounting. They find that the effects of financial incentives, framed as gains, are mixed. However, there are stronger effects of financial incentives framed as losses.³³ Among younger students, non-financial incentives such as awards and trophies are as effective as financial incentives. There are also significant treatment effects such as age-specific effects (older children respond more to financial incentives) and gender effects (boys are more responsive to short-term incentives).³⁴ When incentives are delayed in the experiments (e.g. students are paid a month later, then the beneficial effects of incentives are substantially curtailed. In actual practice the rewards to education, such as jobs, salaries, and perks (which are the corresponding incentives in the real world) follow with substantial delay, suggesting that there could be underinvestment in education.)

A regular exercise regime coupled with a healthy diet, taken in moderation, has well-known benefits. These include higher self-esteem, better body image, a lower likelihood of being depressed and a lower risk factor for several health problems such as cardiovascular disease. Hence, it may be argued that there are already high incentives in place to induce a regular exercise regime, yet such a regime is not universally observed. Charness and Gneezy (2009) ask if financial incentives can aid in developing a habit to exercise?

In Study 1, they first give a handout that describes the benefit of exercise. Then they distribute the subjects into three groups. In the first group, the control group, no financial incentives are given. In the second group, the *one time group*, \$25 are given for exercising once a week. The third group, the *eight time group*, starts off with a payment of \$25 for exercising once a week, followed by a payment of \$100 for exercising eight more times in the next four weeks. In study 2, subjects were given \$75 to show up, have biometric readings taken, and receive a handout that described the benefits from exercise. At this time, subjects were divided into three groups and invited to come back two more times with the promise of a payment of \$50 each time as well as the recording of biometric readings. In the first group, the control group, there was no requirement to exercise. In group 2, the one time group, participants were asked to attend the gym once in the intervening period, while in group 3, the eight time group, participants were asked to attend the gym eight times in the intervening period.

The results are shown in Figure 8.25. The vertical axis shows visits to the gym per week. The period “before” records gym visits in the eight weeks before the intervention and the period “after” records gym visits in the seven weeks after the intervention. There is no statistical difference

³³ The framing effect of incentives jives well with the importance of framing considered elsewhere in the book. For instance, in the context of social preferences, Liberman et al. (2004) find that the manner in which a prisoner’s dilemma game is framed (a cooperative or competitive game) significantly affects the prosocial behavior of participants.

³⁴ The gender specific effects of incentives were also found in some of the studies considered above that examine educational incentives, e.g., Fryer (2011) and Angrist and Lavy (2009).




Figure 8.25 Number of weekly visits to the gym. Error bars reflect one standard error.
Source: Charness and Gneezy (2009) with permission of The Econometric Society.

between the control group and the one-time group. However, the eight-times group, the high incentive group, reveals significantly higher gym attendance in the “after” period. There is also a significant improvement in the biometric measures for this group. The positive effect of incentives in this group is driven largely by people who did not attend the gym regularly in the “before” period.

Exercises for Part 2

1. Prove Lemma 6.4.
2. (Fehr and Schmidt, 1999) Consider the ultimatum game with responder competition when all agents are other-regarding. Suppose that there are n players. Player 1 is the proposer and the remaining players are the responders, whose Fehr–Schmidt preference parameters are commonly given by (α, β) . The proposer offers a share $s \in [0, 1]$ that is observed by all responders who simultaneously choose to accept or reject. If more than one responder accepts, then one among this set is chosen randomly as the winner. The winning responder gets s , the proposer gets $1 - s$, and all other responders get 0. Suppose that $\beta < \frac{n-1}{n}$. Show that there exists a subgame perfect equilibrium in which all responders accept any $s \geq 0$ and the proposer offers $s = 0$. Show also that the highest offer that can be sustained as a subgame perfect equilibrium is given by $s^* = \frac{\alpha}{(1-\beta)(n-1)+2\alpha+\beta} < \frac{1}{2}$.
3. (Fehr and Schmidt, 1999) Prove Proposition 6.5.
4. (Fehr and Schmidt, 1999) Consider the public goods game in Section 5.4. Suppose that all contributors have Fehr–Schmidt preferences. Show that:
 - (a) If $r + \beta_i < 1$ for player i , then it is a dominant strategy for that player to choose $g_i = 0$.
 - (b) Let k denote the number of players with $r + \beta_i < 1$, $k \in [1, n]$. If $\frac{k}{n-1} > \frac{r}{2}$, then there is a unique equilibrium with $g_i = 0$ for all i .
 - (c) If $\frac{k}{n-1} < \frac{r+\beta_j^{-1}}{\alpha_j+\beta_j}$ for all players $j = 1, 2, \dots, n$ and $r + \beta_j > 1$, then there exist other equilibria with positive contribution levels. In these equilibria, all players with $r + \beta_i < 1$ choose $g_i = 0$, while all other players contribute $g_j = [0, y]$. Further, $\frac{r+\beta_j^{-1}}{\alpha_j+\beta_j} < \frac{r}{2}$.
5. (Fehr and Schmidt, 1999) Prove Proposition 6.6.
6. In a public goods game experiment that is repeated over several rounds, contributions stabilize at a very low level. Does this provide support for self-regarding preferences?
7. Consider the three-player experiments by Güth and van Damme (1998), which combine a ultimatum and a dictator game. There is a single proposer (player 1) and two responders, an active responder (player 2), and a passive responder (player 3). The proposer makes an offer of the form (x_1, x_2, x_3) such that the proposer's share is x_1 , the active responder's share is x_2 , and the passive responder's share is x_3 . Player 2 can (as in ultimatum games) either accept or reject the offer. Player 3 does not have any actions. The main empirical finding is that the proposer allocates very low amounts to player 3 and the rejection decision of player 2 does not seem to be significantly influenced by the amount offered to player 3. What are the implications of these results for maximin preferences (as in Charness and Rabin, 2002) and for models of inequity aversion such as the Fehr–Schmidt model and the Bolton–Ockenfels model?
8. (al-Nowaihi and Dhami, 2015) Consider the prisoner's dilemma game.

	C	D
C	2, 2	0, x
D	x , 0	1, 1

Each player has two choices: cooperate (C) or defect (D). The payoff-matrix shows material payoffs of the two players and $x > 2$. In terms of the material payoffs, (C, C) is not a Nash equilibrium. Now suppose that players are inequity averse and have Fehr–Schmidt preferences.

- (a) Suppose that $x = 3$. What is the lower bound on β that achieves the cooperative outcome (C, C) ?
- (b) Now suppose that $x = 10$. How does your answer to (a) change? Is the lower bound on β consistent with the experimental evidence? What can you conclude about the ability of Fehr–Schmidt preferences to explain the empirical results that in the static prisoner’s dilemma games, we observe the outcome (C, C) about 50% of the time?
- (c) “It has been suggested that players achieve a higher degree of cooperation in prisoner’s dilemma games when the game is framed as a community game (cooperative frame) rather than as a Wall Street game (competitive frame). It has also been shown that preplay communication improves cooperation in the prisoner’s dilemma game.” Can models of inequity aversion explain these results? What additional features might you need to explain this finding?
9. There is a positive level of giving in dictator games. Which of the three main types of theories—inequity aversion, intention-based reciprocity, or type-based reciprocity—can explain this finding?
 10. Equal splits (50–50) are often observed in the ultimatum game. Can these equal splits be rationalized in terms of the Bolton–Ockenfels (2000), ERC model? What restrictions are needed on the parameters α, β in the Fehr and Schmidt (1999) model to account for the empirical finding of equal splits? Can the model of Rotemberg (2008) (see Section 6.8.2) account for the equal splits, and if so, what is the difference in the transmission channel relative to the Fehr–Schmidt model.
 11. Charness (2004) conducts two treatments in a gift exchange game. In the random choice condition, the proposer’s choice is randomly determined. In the human choice condition, a human proposer makes the choice. What is the prediction of intention-based theories? What is the prediction of type-based reciprocity theories about the relative correlation between wage and effort in the two different treatments? The empirical finding in these experiments was that there is a significant positive correlation between wages and effort in the random choice condition. This correlation was only slightly lower in the human choice condition. What are the implications of these experiments for the three main types of theories—inequity aversion, intention-based reciprocity, or type-based reciprocity?
 12. (a) Eckel and Gintis (2010, p. 3) make the following comment: “Binmore and Shaked’s defense of the neoclassical analysis on the grounds that ‘all possible divisions of the money in the ultimatum game are Nash equilibrium outcomes’ is transparently faulty.” Do you agree with the comment?
 (b) Evaluate the following quotes from Eckel and Gintis (2010):
 (b1) (p. 4) “Binmore and Shaked also argue that other-regarding behavior disappears or is greatly attenuated when the stakes are high. It is not clear why this is a criticism.”
 (b2) On p. 11–12, Eckel and Gintis quote Binmore and Shaked, thus: “Behavioral economists sometimes claim that neoclassical economists hold that people are self-regarding . . . But no such [selfishness] axiom appears in neoclassical economics textbooks . . . When utility functions of various kinds are fitted to data obtained in laboratory experiments, neoclassical economics is therefore in no danger of being refuted.” Evaluate this claim in the context of a general equilibrium neoclassical model.
 13. (Dhami and al-Nowaihi, 2013) Prove Proposition 6.10.
 14. (Dhami and al-Nowaihi, 2013) Prove Propositions 6.11, 6.12.
 15. (Dhami and al-Nowaihi, 2010a) Prove Proposition 6.13.
 16. (Dhami and al-Nowaihi, 2010a) Prove Lemmas 6.6 and 6.7.

17. (Dhami and al-Nowaihi, 2010a) Consider the model in Section 6.5. (a) Prove Proposition 6.14.
(b) Show that if utility is quasilinear

$$\tilde{u}(c, 1 - l) = c - \frac{1}{2}l^2,$$

then a majority prefers the tax rate that is optimal for the median-skill voter. In particular, show that assumption A1 holds and assumption A2 in Proposition 6.14 is not needed in this case.

18. (Dhami and al-Nowaihi, 2010a) Prove Proposition 6.15.
19. (Dhami and al-Nowaihi, 2010a) Prove Proposition 6.16.
20. (Bartling et al., 2013) Consider the model of vertical integration in Section 8.4.4. How does this model offer at least a partial justification of the role of trade unions and labor market legislation?
21. (Itoh, 2004) Prove Propositions 8.26 and 8.27.
22. (Englmaier and Leider, 2012) Prove Proposition 8.30.
23. (Fehr et al., 2007) Prove Lemma 8.11. [Hint: Use the same method that you used in Lemma 8.10].
24. (Fehr et al., 2007) Prove Proposition 8.32. [Hint: Set up the Kuhn–Tucker conditions for the other-regarding principal's optimization problem and consider the various combinations of the complementary slackness conditions.]
25. (Fehr et al., 2007) Prove Proposition 8.31. [Hint: For each of the contracts, solve, using backward induction, the two-stage game in which the principal first chooses the contractual parameters followed by the agent's choice of effort].
26. Prove Lemma 8.19. [Hint: Rewrite the MLRP in (8.75) as $g(\sigma_H | c_H)g(\sigma_L | c_L) < g(\sigma_L | c_H)g(\sigma_H | c_L)$. Then integrate both sides twice, first over σ_L from σ_L to 1 and then over σ_H from σ_H to 1.
27. (Bénabou and Tirole, 2003) (a) Show that the conclusions of Proposition 8.37 are unaltered if the outside option of the agent is $\bar{U} > 0$.
(b) Consider the setting of the Bénabou–Tirole (2003) model in Section 8.5.1. Suppose now that costs are public information but that the principal has private information about the agent's ability to perform the task, θ , that is drawn from the distribution $F(\theta)$. The principal observes θ perfectly. The agent observes an imperfect signal of θ given by $\sigma \in [0, 1]$ that has conditional distribution function $G(\sigma | \theta)$ and conditional density $g(\sigma | \theta)$ that satisfies the monotone likelihood ratio property (see (8.75)). In all other respects we have the model given in Section 8.5. The principal does not observe the agent's effort e but pays a performance contingent reward, b , in the event that the project is successful. Show that all the results in Proposition 8.37 continue to hold with appropriate changes in notation and terminology.
28. Can you explain the following empirical findings using standard models of neoclassical economics in a plausible manner using reasonable assumptions? If not, what assumptions might you need on human sociality or non-standard human motivations to explain them.
(a) (Restivo and van de Rijt, 2012) Contributors to Wikipedia were randomly chosen to be given Barnstar rewards, a form of status-reward that was publicly displayed. Over the next three months, the individuals who were given the reward were 60% more productive as compared to a control group that was not given the reward.

(b) (Kosfeld and Neckermann, 2011) In a single data entry task, subjects in a gift exchange experiment were told that the top two performers will receive a congratulatory card and a personal thank-you note from the managing director. This simple intervention increased performance by 12%, which is identical to the hypothetical effect of an increase in wage of 35–72%, based on output elasticity measured in previous gift exchange games.

(c) (Ashraf et al., 2014) In a field experiment in Zambia, barbers were recruited by the public health organization to sell condoms. Four treatments varied in their reward structure. In a control treatment, there were no rewards; in two treatments, the margins offered were 10% and 90% of sales; in the star treatment there was a non-pecuniary reward in the form of a star stamped on a publicly displayed chart. The star treatment was supposed to capture the esteem associated with a sense of public duty. After a year it was found that barbers in the star treatment had sold twice as many condoms as in any other treatment.

29. Use the material in this chapter to explain the interaction between other-regarding preferences and incentives in the following description of events from Bowles and Polanía-Reys (2012, p. 369). “On December 1, 2001, the Boston Fire Department terminated its policy of unlimited paid sick days, replacing it with a fifteen-day sick day limit; pay would be docked for firemen exceeding the limit. The firemen responded to the new incentives: those calling in sick on Christmas and New Year’s Day increased tenfold over the previous year. The Fire Commissioner retaliated by canceling their holiday bonus checks (Belkin 2002). The firemen were unimpressed: the year following they claimed 13,431 sick days; up from 6,432 the previous year (Greenberger 2003). Many of the firemen, apparently angered by the new system, abused it or abandoned their previous ethic of serving the public even when injured or not feeling well.”
30. (Ariely et al. 2009b) Consider the following example. “An individual is considering buying a new hybrid car. This car is more expensive than an equivalent car operating with a standard gasoline engine, but the hybrid car helps in preserving the environment. Driving a car which is clearly a hybrid automobile would probably add to one’s positive image, especially for those who live in a community that values environmentally friendly technologies. Suppose now that the government gives a large tax benefit to those who decide to purchase a hybrid car (and everybody knows about this). The tax incentive, of course, reduces the price of hybrid cars and therefore should make the hybrid car more attractive for the individual.” Will the tax incentive be necessarily successful in boosting sales of the hybrid car among all groups of individuals?
31. (Ellingsen and Johannesson, 2008) Prove Proposition 8.39.

REFERENCES FOR PART 2

- Abbink, K., Brandts, J., Herrmann, B., and Orzen, H. (2010). Intergroup conflict and intra-group punishment in an experimental contest game. *American Economic Review* 100(1): 420–47.
- Abbink, K., Irlenbusch, B., and Renner, E. (2000). The moonlighting game. An experimental study on reciprocity and retribution. *Journal of Economic Behavior and Organization* 42(2): 265–77.
- Abeler, J., Altmann, S., Kube, S., and Wibrall, M. (2010). Gift exchange and workers' fairness concerns: when equality is unfair. *Journal of the European Economic Association* 8(6): 1299–324.
- Abeler, J., Becker, A., and Falk, A. (2014). Truth-telling: a representative assessment. Forthcoming in *Journal of Public Economics*.
- Ackert, L. F., Martinez-Vazquez, J., and Rider, M. (2007). Social preferences and tax policy design: some experimental evidence. *Economic Inquiry* 45(3): 487–501.
- Akerlof, G. A. (1982). Labor contracts as partial gift exchange. *Quarterly Journal of Economics* 97(4): 543–69.
- Akerlof, G. A. and Kranton, R. E. (2000). Economics and identity. *Quarterly Journal of Economics* 115(3): 715–53.
- Akerlof, G. A. and Kranton, R. E. (2005). Identity and the economics of organizations. *Journal of Economic Perspectives* 19(1): 9–32.
- Akerlof, G. A. and Kranton, R. E. (2008). Identity, supervision, and work groups. *American Economic Review* 98(2): 212–17.
- Akerlof, G. A. and Yellen, J. L. (1990). The fair wage-effort hypothesis and unemployment. *Quarterly Journal of Economics* 105(2): 255–83.
- Alexander, R. D. (1987). *The Biology of Moral Systems*. New York: Aldine De Gruyter.
- al-Nowaihi, A. and Dhami, S. (2015). Evidential equilibria: heuristics and biases in static games of complete information. *Games* 6(4): 637–77.
- Andersen, S., Ertaç, S., Gneezy, U., Hoffman, M., and List, J. A. (2011). Stakes matter in ultimatum games. *American Economic Review* 101(7): 3427–39.
- Andreoni, J. (1988). Why free ride?: strategies and learning in public goods experiments. *Journal of Public Economics* 37(3): 291–304.
- Andreoni, J. and Bernheim, B. D. (2009). Social image and the 50–50 norm: a theoretical and experimental analysis of audience effects. *Econometrica* 77(5): 1607–636.
- Andreoni, J. and Miller, J. H. (1993). Rational cooperation in the finitely repeated prisoner's dilemma: experimental evidence. *Economic Journal* 103(418): 570–85.
- Andreoni, J. and Miller, J. H. (2002). Giving according to GARP: an experimental test of the consistency of preferences for altruism. *Econometrica* 70(2): 737–53.
- Andreoni, J. and Petrie, R. (2008). Beauty, gender and stereotypes: evidence from laboratory experiments. *Journal of Economic Psychology* 29: 73–93.
- Andreoni, J. and Vesterlund, L. (2001). Which is the fair sex? Gender differences in altruism. *Quarterly Journal of Economics* 116(1): 293–312.
- Angrist, J. and Lavy, V. (2009). The effects of high stakes high school achievement rewards: evidence from a randomized trial. *American Economic Review* 99(4): 1384–414.
- Ariely, D., Bracha, A., and Meier, S. (2009b). Doing good or going well? Image motivation and monetary incentives in behaving prosocially. *American Economic Review* 99(1): 544–55.
- Ariely, D., Gneezy, U., Loewenstein, G., and Mazar, N. (2009a). Large stakes and big mistakes. *Review of Economic Studies* 76(2): 451–69.
- Ashraf, N., Bandiera, O., and Jack, B. K. (2014). No margin, no mission? A field experiment on incentives for public services delivery. *Journal of Public Economics* 120: 1–17.
- Axelrod, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.
- Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science* 211(4489): 1390–6.
- Babcock, L. and Loewenstein, G. (1997). Explaining bargaining impasse: the role of self-serving biases. *Journal of Economic Perspectives* 11(1): 109–26.
- Bandiera, O., Barankay, I., and Rasul, I. (2005). Social preferences and the response to incentives: evidence from personnel data. *Quarterly Journal of Economics* 120(3): 917–62.

- Bandura, A. (1991). Social cognitive theory of moral thought and action. In W. M. Kurtines, J. Gewirtz, and J. L. Lamb (eds.), *Handbook of Moral Behavior and Development: Volume I, Theory*. Hillsdale, Lawrence Erlbaum and Associates, pp. 45–103.
- Baran, N. M., Sapienza, P., and Zingales, L. (2010). Can we infer social preferences from the lab? Evidence from the trust game. NBER. Working Paper 15654.
- Barmettler, F., Fehr, E., and Zehnder, C. (2012). Big experimenter is watching you! Anonymity and prosocial behavior in the laboratory. *Games and Economic Behavior* 75(1): 17–34.
- Bartling, B. (2011). Relative performance or team evaluation? Optimal contracts for other-regarding agents. *Journal of Economic Behavior and Organization* 79(3): 183–93.
- Bartling, B., Fehr, E., and Schmidt, K. M. (2012). Screening, competition, and job design: economic origins of good jobs. *American Economic Review* 102(2): 834–64.
- Bartling, B., Fehr, E., and Schmidt, K. M. (2013). Use and abuse of authority: a behavioral foundation of the employment relation. *Journal of the European Economic Association* 11(4): 711–42.
- Bartling, B. and Fischbacher, U. (2012). Shifting the blame: on delegation and responsibility. *Review of Economic Studies* 79(1): 67–87.
- Bartling, B. and von Siemens, F. A. (2010). The intensity of incentives in firms and markets: moral hazard with envious agents. *Labour Economics* 17(3): 598–607.
- Bateson, M., Nettle, D., and Roberts, G. (2006). Cues of being watched enhance cooperation in a real-world setting. *Biology Letters* 2: 412–14.
- Battigalli, P. and Dufwenberg, M. (2007). Guilt in games. *American Economic Review Papers and Proceedings* 97(2): 170–76.
- Battigalli, P. and Dufwenberg, M. (2009). Dynamic psychological games. *Journal of Economic Theory* 144(1): 1–35.
- Belkin, D. (2002). Boston firefighters sick—or tired of working 15-day allowance seen fueling call-ins. *Boston Globe* 18 Jan. B1.
- Bellemare, C., Kröger, S., and Van Soest, A. (2008). Measuring inequity aversion in a heterogeneous population using experimental decisions and subjective probabilities. *Econometrica* 76(4): 815–39.
- Bellemare, C. and Shearer, B. S. (2007). Gift exchange within a firm: evidence from a field experiment. IZA, Discussion Paper 2696.
- Bénabou, R. and Tirole, J. (2003). Intrinsic and extrinsic motivation. *Review of Economic Studies* 70(3): 489–20.
- Bénabou, R. and Tirole, J. (2006). Incentives and prosocial behaviour. *American Economic Review* 96(5): 1652–78.
- Bénabou, R. and Tirole, J. (2011). Identity, morals, and taboos: beliefs as assets. *Quarterly Journal of Economics* 126(2): 805–55.
- Benjamin, D. J., Choi, J. J., and Strickland, A. J. (2010). Social identity and preferences. *American Economic Review* 100(4): 1913–28.
- Benz, M. and Meier, S. (2008). Do people behave in experiments as in the field?—evidence from donations. *Experimental Economics* 11(3): 268–81.
- Berg, J., Dickhaut, J., and McCabe, K. (1995). Trust, reciprocity, and social history. *Games and Economic Behavior* 10(1): 122–42.
- Bernhard, H., Fehr, E., and Fischbacher, U. (2006). Group affiliation and altruistic norm enforcement. *American Economic Review* 96(2): 217–21.
- Besley, T. and Ghatak, M. (2005). Competition and incentives with motivated agents. *American Economic Review* 95(3): 616–36.
- Bewley, T. (1999). *Why Wages Don't Fall During a Recession*. Cambridge, MA: Harvard University Press.
- Bewley, T. (2005). Fairness, reciprocity and wage rigidity. In H. Gintis, S. Bowles, R. Boyd, and E. Fehr (eds.), *Moral Sentiments and Material Interests* Cambridge, MA: The MIT Press, pp. 303–38.
- Billig, M. and Tajfel, H. (1973). Social categorization and similarity in intergroup behaviour. *European Journal of Social Psychology* 3(1): 27–52.
- Binmore, K. (ed.) (1998). *Just Playing: Game Theory and the Social Contract, Volume 2*. Cambridge, MA: MIT Press.
- Binmore, K. G., Samuelson, L., and Gale, J. (1995). Learning to be imperfect: the ultimate game. *Games and Economic Behavior* 8(1): 56–90.
- Blanco, M., Engelmann, D., and Normann, H. T. (2011). A within-subject analysis of other-

- regarding preferences. *Games and Economic Behavior* 72(2): 321–38.
- Blount, S. (1995). When social outcomes aren't fair: the effect of causal attributions on preferences. *Organizational Behavior and Human Decision Process* 63(2): 131–44.
- Bochet, O., Page, T., and Putterman, L. (2006). Communication and punishment in voluntary contribution experiments. *Journal of Economic Behavior and Organization* 60(1): 11–26.
- Bohnet, I., Frey, B. S., and Huck, S. (2001). More order with less law: on contractual enforcement, trust, and crowding. *American Political Science Review* 95(1): 131–44.
- Bohnet, I., Harmgart, H., Huck, S., and Tyran, J.-R. (2005). Learning trust. *Journal of the European Economic Association* 3(2–3): 322–9.
- Bohnet, I. and Kübler, D. (2005). Compensating the cooperators: is sorting in the prisoner's dilemma possible? *Journal of Economic Behavior and Organization* 56(1): 61–76.
- Bolle, F. and Kritikos, A. (1998). Self-centered inequality aversion versus reciprocity and altruism. Mimeo, Europa-Universität Viadrina.
- Bolton, G. E. (1991). A comparative model of bargaining: theory and evidence. *American Economic Review* 81(5): 1096–136.
- Bolton, G. E., Brandts, J., Katok, E., Ockenfels, A., and Zwick, R. (2008). Testing theories of other-regarding behavior: a sequence of four laboratory studies. In C. R. Plott and V. L. Smith (eds.), *Handbook of Experimental Economics Results* Amsterdam: North Holland, pp. 488–99.
- Bolton, G. E., Brandts, J., and Ockenfels, A. (1998). Measuring motivations for the reciprocal responses observed in a simple dilemma game. *Experimental Economics* 1(3): 207–19.
- Bolton, G. E., Brandts, J., and Ockenfels, A. (2005). Fair procedures: evidence from games involving lotteries. *Economics Journal* 115(506): 1054–76.
- Bolton, G. E., Greiner, B., and Ockenfels, A. (2013). Engineering trust: reciprocity in the production of reputation information. *Management Science* 59(2): 265–85.
- Bolton, G. E., Katok, E., and Ockenfels, A. (2004). How effective are electronic reputation mechanisms? *Management Science* 50(11): 1587–602.
- Bolton, G. E. and Ockenfels, A. (1998). Strategy and equity: an ERC-analysis of the Guth-van Damme Game. *Journal of Mathematical Psychology* 42(2): 215–26.
- Bolton, G. E. and Ockenfels, A. (2000). ERC: a theory of equity, reciprocity, and competition. *American Economic Review* 90(1): 166–93.
- Bolton, G. E. and Ockenfels, A. (2005). A stress test of fairness measures in models of social utility. *Economic Theory* 25(4): 957–82.
- Bolton, G. E. and Ockenfels, A. (2006). Measuring efficiency and equity motives: a comment on “Inequality aversion, efficiency, and maximin preferences in simple distribution experiments”. *American Economic Review* 96(5): 1906–11.
- Bolton, G. E. and Ockenfels, A. (2014). Does laboratory trading mirror behavior in real world markets? Fair bargaining and competitive bidding on eBay. *Journal of Economic Behavior and Organization* 97: 143–54.
- Bolton, G. E. and Zwick, R. (1995). Anonymity versus punishment in ultimatum bargaining. *Games and Economic Behavior* 10(1): 95–121.
- Bolton, G. E., Zwick, R., and Katok, E. (1998). Dictator game giving: rules of fairness versus acts of kindness. *International Journal of Game Theory* 27(2): 269–99.
- Bowles, S. and Gintis, H. (2003). Origins of human cooperation. In P. Hammerstein (ed.), *Genetic and Cultural Evolution of Cooperation*. Cambridge, MA: MIT Press, pp. 429–43.
- Bowles, S. and Gintis, H. (2013). *A Cooperative Species: Human Reciprocity and Its Evolution*. Princeton, NJ: Princeton University Press.
- Bowles, S., Gintis, H., and Osborne, M. (2001). The determinants of earnings: a behavioral approach. *Journal of Economic Literature* 39(4): 1137–76.
- Bowles, S. and Polania-Reyes, S. (2012). Economic incentives and social preferences: substitutes or complements? *Journal of Economic Literature* 50(2): 368–25.
- Brandts, J. and Charness, G. (2004). Do labour market conditions affect gift exchange? Some experimental evidence. *Economic Journal* 114(497): 684–708.
- Brandts, J., Saijo, T., and Schram, A. (2004). How universal is behavior? A four country comparison of spite and cooperation in voluntary contribution mechanisms. *Public Choice* 119: 381–424.
- Brandts, J. and Schram, A. (2001). Cooperation and noise in public goods experiments: applying the contributions function approach. *Journal of Public Economics* 79(2): 399–427.

- Brown, M., Falk, A., and Fehr, E. (2004). Relational contracts and the nature of market interactions. *Econometrica* 72(3): 747–80.
- Brown, M., Falk, A., and Fehr, E. (2012). Competition and relational contracts: the role of unemployment as a disciplinary device. *Journal of the European Economic Association* 10(4): 887–907.
- Burks, S., Carpenter, J., and Goette, L. (2009). Performance pay and worker cooperation: evidence from an artefactual field experiment. *Journal of Economic Behavior and Organization* 70(3): 458–69.
- Burnham, T., McCabe, K., and Smith, V. L. (2000). Friend-or-foe intentionality priming in an extensive form trust game. *Journal of Economic Behavior and Organization* 43(1): 57–73.
- Cabrales, A., Miniaci, R., Piovesan, M., and Ponti, G. (2010). Social preferences and strategic uncertainty: an experiment on markets and contracts. *American Economic Review* 100(5): 2261–78.
- Cadsby, C. B. and Maynes, E. (1998). Gender and free riding in a threshold public goods game: experimental evidence. *Journal of Economic Behavior and Organization* 34(4): 603–20.
- Camerer, C. F. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton University Press.
- Camerer, C. F. (2015). The promise and success of lab-field generalizability in experimental economics: a critical reply to Levitt and List. In G. R. Fréchette and A. Schotter (eds.) *Handbook of Experimental Economic Methodology*. Oxford: Oxford University Press, pp. 249–95.
- Camerer, C. F. and Thaler, R. H. (1995). Anomalies: ultimatums, dictators, and manners. *Journal of Economic Perspectives* 9(2): 209–19.
- Camerer, C. F. and Weighlt, K. (1988). Experimental tests of a sequential equilibrium reputation model. *Econometrica* 56(1): 1–36.
- Cameron, L. A. (1999). Raising the stakes in the ultimatum game: experimental evidence from Indonesia. *Economic Inquiry* 37(1): 47–59.
- Cappelen, A. W., Sørensen E. Ø., and Tungodden, B. (2013). When do we lie? *Journal of Economic Behavior and Organization* 93: 258–65.
- Cappelen, A. W., Halvorsen, T., Sorensen, E., and Tungodden, B. (2015). Face-saving or fair-minded: what motivates moral behavior? Mimeo, Norwegian School of Economics.
- Cardenas, J. C., Stranlund, J. K., and Willis, C. E. (2000). Local environmental control and institutional crowding-out. *World Development* 28(10): 1719–33.
- Carpenter, J. P. (2007). Punishing free-riders: how group size affects mutual monitoring and the provision of public goods. *Games and Economic Behavior* 60(1): 31–51.
- Carpenter, J. P., Bowles, S., Gintis, H., and Hwang, S. H. (2009). Strong reciprocity and team production: theory and evidence. *Journal of Economic Behavior and Organization* 71(2): 221–32.
- Carpenter, J., Corrolly, C., and Myers, C. (2008). Altruistic behavior in a representative dictator experiment. *Experimental Economics* 11: 282–98.
- Carpenter, J. and Myers, C. (2010). Why volunteer? Evidence on the role of altruism, image, and incentives. *Journal of Public Economics* 94(11–12): 911–20.
- Carpenter, J. P. and Seki, E. (2011). Do social preferences increase productivity? Field experimental evidence from fishermen in Toyama Bay. *Economic Inquiry* 49(2): 612–30.
- Carpenter, J. P., Verhoogen, E., and Burks, S. (2005). The effect of stakes in distribution experiments. *Economics Letters* 86(3): 393–8.
- Charness, G. (2000). Responsibility and effort in an experimental labor market. *Journal of Economic Behavior and Organization* 42(3): 375–84.
- Charness, G. (2004). Attribution and reciprocity in an experimental labor market. *Journal of Labor Economics* 22(3): 665–88.
- Charness, G., Cobo-Reyes, R., and Jiménez, M. (2008). An investment game with third-party intervention. *Journal of Economic Behavior and Organization* 68(1): 18–28.
- Charness, G., Du, N., and Yang, C. L. (2011). Trust and trustworthiness reputations in an investment game. *Games and Economic Behavior* 72(2): 361–75.
- Charness, G. and Dufwenberg, M. (2006). Promises and partnership. *Econometrica* 74(6): 1579–601.
- Charness, G., Frechette, G. R., and Kagel, J. H. (2004). How robust is laboratory gift exchange? *Experimental Economics* 7(2): 189–05.
- Charness, G. and Gneezy, U. (2009). Incentives to exercise. *Econometrica* 77(3): 909–31.

- Charness, G. and Rabin, M. (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics* 117(3): 817–69.
- Charness, G., Rigotti, L., and Rustichini, A. (2007). Individual behavior and group membership. *American Economic Review* 97(4): 1340–52.
- Chaudhuri, A., Graziano, S., and Maitra, P. (2006). Social learning and norms in a public goods experiment with intergenerational advice. *Review of Economic Studies* 73(2): 357–80.
- Chaudhuri, A. and Paichayontvijit, T. (2006). Conditional cooperation and voluntary contributions to a public good. *Economics Bulletin* 3(8): 1–14.
- Chen, R. and Chen, Y. (2011). The potential of social identity for equilibrium selection. *American Economic Review* 101(6): 2562–89.
- Chen, Y. and Li, S. X. (2009). Group identity and social preferences. *American Economic Review* 99(1): 431–57.
- Cherry, T. L., Frykblom, P., and Shogren, J. F. (2002). Hardnose the dictator. *American Economic Review* 92(4): 1218–21.
- Childs, J. (2012). Gender differences in lying. *Economics Letters* 114(2): 147–9.
- Cinyabuguma, M., Page, T., and Putterman, L. (2005). Cooperation under the threat of expulsion in a public goods experiment. *Journal of Public Economics* 89(8): 1421–35.
- Cleave, B. L., Nikiforakis, N., and Slonim, R. (2012). Is there selection bias in laboratory experiments? The case of social and risk preferences. *Experimental Economics* 16(3): 372–82.
- Coase, R. (1937). The nature of the firm. *Economica* 4(16): 386–05.
- Cohn, A., Fehr, E., and Goette, L. (2012). Fair wages and effort: evidence from a field experiment. Mimeo, Department of Economics, University of Zurich.
- Cohn, A., Fehr, E., Herrmann, B., Schneider, F. (2011). Social comparison in the workplace: evidence from a field experiment. IZA, Discussion Paper 5550.
- Cohn, A., Maréchal, M. A., and Noll, T. (2013). Bad boys: how criminal identity affects rule violation. Mimeo, Department of Economics, University of Zurich.
- Conrads, J., Irlenbusch, B., Rilke, R. M., and Walkowitz, G. (2013). Lying and team incentives. *Journal of Economic Psychology* 34: 1–7.
- Cox, J. C. (2000). Trust and reciprocity: implications of game triads and social contexts. University of Arizona. Working Paper 00–11.
- Cox, J. C. (2004). How to identify trust and reciprocity. *Games and Economic Behavior* 46(2): 260–81.
- Cox, J. C., Friedman, D., and Gjerstad, S. (2007). A tractable model of reciprocity and fairness. *Games and Economic Behavior* 59(1): 17–45.
- Crawford, V. (1997). Theory and experiment in the analysis of strategic interaction. In D. Kreps and K. Wallis (eds.), *Advances in Economics and Econometrics: Theory and Applications, Seventh World Congress, Volume I*. Cambridge: Cambridge University Press, pp. 206–42.
- Croson, R. and Gneezy, U. (2009). Gender differences in preferences. *Journal of Economic Literature* 47(2): 448–74.
- Dal Bó, P. (2002). Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. Brown University, Department of Economics. Working Paper 2002–21.
- Dal Bó, P., Foster, A., and Putterman, L. (2010). Institutions and behavior: experimental evidence on the effects of democracy. *American Economic Review* 100(5): 2205–29.
- Daly, M. and Wilson, M. (1988). Evolutionary social psychology and family homicide. *Science* 242(4878): 519–24.
- Dana, J., Cain, D. M., and Dawes, R. M. (2006). What you don't know won't hurt me: costly (but quiet) exit in a dictator game. *Organizational Behavior and Human Decision Processes* 100(2): 193–201.
- Dana, J., Weber, R. A., and Kuang, J. X. (2007). Exploiting moral wriggle room: experiments demonstrating an illusory preference for fairness. *Economic Theory* 33: 67–80.
- Davis, D. D. and Holt, C. A. (1993). *Experimental Economics*. Princeton, NJ: Princeton University Press.
- Dawes, C. T., Fowler, J. H., Johnson, T., McElreath, R., and Smirnov, O. (2007). Egalitarian motives in humans. *Nature* 446: 794–6.
- Dawes, R. M. (1980). Social dilemmas. *Annual Review of Psychology* 31: 169–93.
- Dawes, R. M. and Thaler, R. H. (1988). Anomalies: cooperation. *Journal of Economic Perspectives* 2(3): 187–97.

- Dawkins, R. (1989). *The Selfish Gene*. 2nd edition. Oxford: Oxford University Press.
- de Quervain, D. J.-F., Fischbacher, U., Treyer, V., et al. (2004). The neural basis of altruistic punishment. *Science* 305(5688): 1254–8.
- Deci, E. L. (1975). *Intrinsic Motivation*. New York: Plenum Press.
- Deci, E. L., Koestner, R., and Ryan, R. M. (1999). A meta-analytic review of experiments examining the effects of extrinsic rewards on intrinsic motivation. *Psychological Bulletin* 125(6): 627–68.
- Deci, E. L. and Ryan R. M. (1985). *Intrinsic Motivation and Self-Determination in Human Behavior*. Berlin: Springer.
- DellaVigna, S. (2009). Psychology and economics: evidence from the field. *Journal of Economics Literature* 47(2): 315–72.
- Dhami, S. and al-Nowaihi, A. (2010a). Existence of a Condorcet winner when voters have other-regarding preferences. *Journal of Public Economic Theory* 12(5): 897–22.
- Dhami, S. and al-Nowaihi, A. (2010b). Redistributive policy with heterogeneous social preferences of voters. *European Economic Review* 54(6): 743–59.
- Dhami, S. and al-Nowaihi, A. (2013). Stochastic dominance for Fehr-Schmidt preferences. University of Leicester, Discussion Papers in Economics. Working Paper 13/09.
- Dickenson, D. and M.-C. Villeval (2008). Does monitoring decrease work effort? The complementarity between agency and crowding-out theories. *Games and Economic Behavior* 63(1): 56–76.
- Dohmen, T. and Falk, A. (2011). Performance pay and multidimensional sorting: productivity, preferences, and gender. *American Economic Review* 101(2): 556–90.
- Dohmen, T., Falk, A., Huffman, D., and Sunde, U. (2009). Homo reciprocans: survey evidence on behavioural outcomes. *Economic Journal* 119(536): 592–12.
- Dreber, A. and Johannesson, M. (2008). Gender differences in deception. *Economics Letters* 99(1): 197–9.
- Duesenberry, J. S. (1949). *Income, Saving, and the Theory of Consumer Behavior*. Cambridge, MA: Harvard University Press.
- Dufwenberg, M., Heidhues, P., Kirchsteiger, G., et al. (2011). Other-regarding preferences in general equilibrium. *Review of Economic Studies* 78(2): 613–39.
- Dufwenberg, M. and Kirchsteiger, G. (2004). A theory of sequential reciprocity. *Games and Economic Behavior* 47(2): 268–98.
- Eckel, C. C. and Gintis, H. (2010). Blaming the messenger: notes on the current state of experimental economics. *Journal of Economic Behavior and Organization* 73(1): 109–19.
- Eckel, C. C. and Grossman, P. J. (1996). Altruism in anonymous dictator games. *Games and Economic Behavior* 16(2): 181–91.
- Eckel, C. C. and Petrie, R. (2011). Face value. *American Economic Review* 101(4): 1497–513.
- Egas, M. and Riedl, A. (2008). The economics of altruistic punishment and the maintenance of cooperation. *Proceedings of the Royal Society B: Biological Sciences* 275(1637): 871–8.
- Ellingsen, T. and Johannesson, M. (2004). Promises, threats and fairness. *Economic Journal* 114(495): 397–420.
- Ellingsen, T. and Johannesson, M. (2008). Pride and prejudice: the human side of incentive theory. *American Economic Review* 98(3): 990–1008.
- Ellman, M. and Pezanis-Christou, P. (2010). Organizational structure, communication and group ethics. *American Economic Review* 100(5): 2478–91.
- Engelmann, D. and Fischbacher, U. (2009). Indirect reciprocity and strategic reputation building in an experimental helping game. *Games and Economic Behavior* 67(2): 399–407.
- Engelmann, D. and Strobel, M. (2004). Inequality aversion, efficiency, and maximin preferences in simple distribution experiments. *American Economic Review* 94(4): 857–969.
- Engelmann, D. and Strobel, M. (2006). Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: reply. *American Economic Review* 96(5): 1918–23.
- Englmaier, F. and Leider, S. (2012). Contractual and organizational structure with reciprocal agents. *American Economic Journal: Microeconomics* 4(2): 146–83.
- Englmaier, F., Strasser, S., and Winter, J. (2014). Worker characteristics and wage differentials: evidence from a gift-exchange experiment. *Journal of Economic Behavior and Organization* 97: 185–203.

- Englmaier, F. and Wambach, A. (2010). Optimal incentive contracts under inequity aversion. *Games and Economic Behavior* 69(2): 312–28.
- Erat, S. and Gneezy, U. (2012). White lies. *Management Science* 58(4): 723–33.
- Falk, A. (2007). Gift exchange in the field. *Econometrica* 75(5): 1501–11.
- Falk, A., Fehr, E., and Fischbacher, U. (2003). On the nature of fair behaviour. *Economic Inquiry* 41(1): 20–6.
- Falk, A., Fehr, E., and Fischbacher, U. (2005). Driving forces behind informal sanctions. *Econometrica* 73(6): 2017–30.
- Falk, A., Fehr, E., and Fischbacher, U. (2008). Testing theories of fairness and reciprocity—intentions matter. *Games and Economic Behavior* 62(1): 287–303.
- Falk, A., Fehr, E., and Zehnder, C. (2006). Fairness perceptions and reservation wages—the behavioral effects of minimum wage laws. *Quarterly Journal of Economics* 121(4): 1347–81.
- Falk, A. and Fischbacher, U. (2006). A theory of reciprocity. *Games and Economic Behavior* 54(2): 293–315.
- Falk, A. and Kosfeld, M. (2006). The hidden costs of control. *American Economic Review* 96(5): 1611–30.
- Falk, A., Meier, S., and Zehnder, C. (2013). Do lab experiments misrepresent social preferences? The case of self-selected student samples. *Journal of the European Economic Association* 11(4): 839–52.
- Falk, A., and Szech, N. (2013). Morals and markets. *Science* 340(6133): 707–11.
- Falk, A. and Szech, N. (2014). Organization, dif-fused pivotality and immoral outcomes. Mimeo, University of Bonn.
- Fehr, E. (2009). On the economics and biology of trust. *Journal of the European Economic Association* 7(2–3): 235–66.
- Fehr, E. and Falk, A. (1999). Wage rigidity in a competitive incomplete contract market. *Journal of Political Economy* 107(1): 106–34.
- Fehr, E. and Falk, A. (2002). Joseph Schumpeter lecture: psychological foundations of incentives. *European Economic Review* 46(4–5): 687–24.
- Fehr, E. and Fischbacher, U. (2002). Why social preferences matter—the impact of non-self-regarding motives on competition, cooperation and incentives. *Economic Journal* 112(478): C1–C33.
- Fehr, E. and Fischbacher, U. (2003). The nature of human altruism. *Nature* 425: 785–91.
- Fehr, E. and Fischbacher, U. (2004). Third-party punishment and social norms. *Evolution and Human Behavior* 25(2): 63–87.
- Fehr, E. and Fischbacher, U. (2005). Human altruism: proximate patterns and evolutionary origins. *Analyse & Kritik* 27(1): 6–47.
- Fehr, E., Fischbacher, U., and Gächter, S. (2002). Strong reciprocity, human cooperation, and the enforcement of social norms. *Human Nature* 13(1): 1–25.
- Fehr, E., Fischbacher, U., and Tougareva, E. (2002). Do high stakes and competition remove reciprocal fairness? Evidence from Russia. University of Zurich, Institute for Empirical Research in Economics, Working paper 120.
- Fehr, E. and Gächter, S. (1999). Cooperation and punishment in public goods experiments. University of Zurich, Institute for Empirical Research in Economics, Working Paper 10.
- Fehr, E. and Gächter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review* 90(4): 980–94.
- Fehr, E. and Gächter, S. (2001). Do incentive contracts crowd out voluntary cooperation? University of Zurich, Institute for Empirical Research in Economics, Working Paper 34.
- Fehr, E. and Gächter, S. (2002). Altruistic punishment in humans. *Nature* 415: 137–40.
- Fehr, E., Gächter, S., and Kirchsteiger, G. (1997). Reciprocity as a contract enforcement device: experimental evidence. *Econometrica* 65(4): 833–60.
- Fehr, E., Goette, L., and Zehnder, C. (2009). A behavioral account of the labor market: the role of fairness concerns. *Annual Reviews: Economics* 1: 355–84.
- Fehr, E., Kirchsteiger, G., and Riedl, A. (1993). Does fairness prevent market clearing? An experimental investigation. *Quarterly Journal of Economics* 108(2): 437–59.
- Fehr, E., Kirchsteiger, G., and Riedl, A. (1998). Gift exchange and reciprocity in competitive experimental markets. *European Economic Review* 42(1): 1–34.
- Fehr, E., Klein, A., and Schmidt, K. M. (2007). Fairness and contract design. *Econometrica* 75(1): 121–54.
- Fehr, E., Krehmelmer, S., and Schmidt, K. M. (2008). Fairness and the optimal allocation of

- ownership rights. *Economic Journal* 118(531): 1262–84.
- Fehr, E. and Leibbrandt, A. (2011). A field study on cooperativeness and impatience in the Tragedy of the Commons. *Journal of Public Economics* 95(9–10): 1144–55.
- Fehr, E. and List, J. A. (2004). The hidden costs and returns of incentives: trust and trustworthiness among CEOs. *Journal of the European Economic Association* 2(5): 743–71.
- Fehr, E., Naef, M., and Schmidt, K. M. (2006). Inequality aversion, efficiency, and maximin preferences in simple distribution experiments: comment. *American Economic Review* 96(5): 1912–17.
- Fehr, E. and Rockenbach, B. (2003). Detrimental effects of sanctions on human altruism. *Nature* 422: 137–40.
- Fehr, E. and Schmidt K. M. (1999). A theory of fairness, competition and cooperation. *Quarterly Journal of Economics* 114(3): 817–68.
- Fehr, E. and Schmidt, K. M. (2006). The economics of fairness, reciprocity and altruism: Experimental evidence and new theories. In S.-C. Kolm and J. M. Ythier (eds.), *Handbook of the Economics of Giving, Altruism and Reciprocity, Volume I*. Elsevier.
- Fehr, E. and Schmidt, K. M. (2007). Adding a stick to the carrot? The interaction of bonuses and fines. *American Economic Review* 97(2): 177–81.
- Fehr, E. and Schneider, F. (2010). Eyes are on us, but nobody cares: are eye cues relevant for strong reciprocity? *Proceedings of the Royal Society B: Biological Sciences* 277: 1315–23.
- Fehr, E., Tougareva, E., and Fischbacher, U. (2014). Do high stakes and competition undermine fair behaviour? Evidence from Russia. *Journal of Economic Behavior and Organization*, in press.
- Fischbacher, U. and Föllmi-Heusi, F. (2013). Lies in disguise: an experimental study on cheating. *Journal of the European Economic Association* 11(3): 525–47.
- Fischbacher, U., Fong, C. M., and Fehr, E. (2009). Fairness, errors and the power of competition. *Journal of Economic Behavior and Organization* 72(1): 527–45.
- Fischbacher, U. and Gächter, S. (2010). Social preferences, beliefs, and the dynamics of free riding in public good experiments. *American Economic Review* 100(1): 541–56.
- Fischbacher, U., Gächter, S., and Fehr, E. (2001). Are people conditionally cooperative? Evidence from a public goods experiment. *Economics Letters* 71(3): 397–04.
- Fisman, R., Kariv, S., and Markovits, D. (2007). Individual preferences for giving. *American Economic Review* 97(5): 1858–76.
- Forsythe, R., Horowitz, J. L., Savin, N. E., and Sefton, M. (1994). Fairness in simple bargaining games. *Games and Economic Behavior* 6(3): 347–69.
- Fréchette, G. R. (2015). Laboratory experiments: Professionals versus students. In G. R. Fréchette and A. Schotter (eds.) *Handbook of Experimental Economic Methodology*. Oxford: Oxford University Press, pp. 360–90.
- Fréchette, G. R., Kagel, J. H., and Lehrer, S. F. (2003). Bargaining in legislatures: an experimental investigation of open versus closed amendment rules. *American Political Science Review* 97(2): 221–32.
- Frey, B. S. and Oberholzer-Gee, F. (1997). The cost of price incentives: an empirical analysis of motivation crowding-out. *American Economic Review* 87(4): 746–55.
- Fryer, R. G. (2011). Financial incentives and student achievement: evidence from randomized trials. *Quarterly Journal of Economics* 126(4): 1755–98.
- Gächter, S. and Falk, A. (2002). Reputation and reciprocity: consequences for the labour relation. *Scandinavian Journal of Economics* 104(1): 1–26.
- Gächter, S. and Herrmann, B. (2009). Reciprocity, culture, and human cooperation: previous insights and a new cross-cultural experiment. *Philosophical Transactions of the Royal Society B* 364(1518): 791–806.
- Gächter, S., Nosenzo, D., and Sefton, M. (2012). The impact of social comparisons on reciprocity. *Scandinavian Journal of Economics* 114(4): 1346–67.
- Gächter, S., Renner, E., and Sefton, M. (2008). The long run benefits of punishment. *Science* 322(5907): 1510.
- Gächter, S. and Thöni, C. (2005). Social learning and voluntary cooperation among like-minded people. *Journal of the European Economic Association* 3(2–3): 303–14.
- Geanakoplos, J., Pearce, D., and Stacchetti, E. (1989). Psychological games and sequential

- rationality. *Games and Economic Behavior* 1(1): 60–79.
- Ghiselin, M. T. (1974). *The Economy of Nature and the Evolution of Sex*. Berkeley, CA: University of California Press.
- Gibbons, R. (2010). Transaction cost economics: past, present, and future? *Scandinavian Journal of Economics* 112(2): 263–88.
- Gibson, R., Tanner, C., and Wagner, A. F. (2013). Preferences for truthfulness: heterogeneity among and within individuals. *American Economic Review* 103(1): 532–48.
- Gintis, H. (2000). Strong reciprocity and human sociality. *Journal of Theoretical Biology* 206(2): 169–79.
- Gintis, H. (2009). *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences*. Princeton, NJ: Princeton University Press.
- Gintis, H., Bowles, S., Boyd, R., and Fehr, E. (2003). Explaining altruistic behavior in humans. *Evolution and Human Behavior* 24(3): 153–72.
- Gintis, H. and Fehr, E. (2012). The social structure of cooperation and punishment. *Behavioral and Brain Sciences* 35(1): 28–9.
- Gintis, H., Smith, E. A., and Bowles, S. (2001). Costly signaling and cooperation. *Journal of Theoretical Biology* 213(1): 103–19.
- Gneezy, A., Gneezy, U., Nelson, L. D., and Brown, A. (2010). Shared social responsibility: a field experiment in pay-what-you-want pricing and charitable giving. *Science* 329(5989): 325–27.
- Gneezy, A., Gneezy, U., Riener, G., and Nelson, L. D. (2012). Pay-what-you-want, identity, and self-signaling in markets. *Proceedings of the National Academy of Sciences of the United States of America* 109(19): 7236–240.
- Gneezy, U. (2005). Deception: the role of consequences. *American Economic Review* 95(1): 384–94.
- Gneezy, U. and List, J. A. (2006). Putting behavioral economics to work: testing for gift exchange in labor markets using field experiments. *Econometrica* 74(5): 1365–84.
- Gneezy, U., Rockenbach, B., and Serra-Garcia, M. (2013). Measuring lying aversion. *Journal of Economic Behavior and Organization* 93: 293–300.
- Gneezy, U. and Rustichini, A. (2000a). A fine is a price. *Journal of Legal Studies* 29(1): 1–17.
- Gneezy, U. and Rustichini, A. (2000b). Pay enough or don't pay at all. *Quarterly Journal of Economics* 115(3): 791–810.
- Goerg, S. J., Kube, S., and Zultan, R. (2010). Treating equals unequally: incentives in teams, workers' motivation and production technology. *Journal of Labor Economics* 28(4): 747–72.
- Goette, L. F., Huffman, D., and Meier, S. (2012). The impact of social ties on group interactions: evidence from minimal groups and randomly assigned real groups. *American Economic Journal: Microeconomics* 4(1): 101–15.
- Goette, L. F. and Stutzer, A. (2008). Blood donations and incentives: evidence from a field experiment. IZA, Discussion Paper 3580.
- Goette, L. F., Stutzer, A., and Frey, B. M. (2010). Prosocial motivation and blood donations: a survey of the empirical literature. *Transfusion Medicine and Hemotherapy* 37(3): 149–54.
- Greenberger, S. S. (2003). Sick day abuses focus of fire talks. *Boston Globe* 17 Sep. B7.
- Grossman, S. J. and Hart, O. D. (1983). An analysis of the principal-agent problem. *Econometrica* 51(1): 7–45.
- Grossman, S. J. and Hart, O. D. (1986). The costs and benefits of ownership: a theory of vertical and lateral integration. *Journal of Political Economy* 94(4): 691–719.
- Guala, F., Mittone, L., and Ploner, M. (2013). Group membership, team preferences, and expectations. *Journal of Economic Behavior and Organization* 86: 183–90.
- Gunnthorsdottir, A., Houser, D., and McCabe, K. (2007). Disposition, history and contributions in public goods experiments. *Journal of Economic Behavior and Organization* 62(2): 304–15.
- Gürerk, Ö., Irlenbusch, B., and Rockenbach, B. (2006). The competitive advantage of sanctioning institutions. *Science* 312(5770): 108–11.
- Güth, W., Schmittberger, R., and Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization* 3(4): 367–88.
- Güth, W. and Tietz, R. (1990). Ultimatum bargaining behavior: a survey and comparison of experimental results. *Journal of Economic Psychology* 11(3): 417–49.
- Güth, W. and van Damme, E. (1998). Information, strategic behavior and fairness in ultimatum bargaining: an experimental study. *Journal of Mathematical Psychology* 42(2–3): 227–47.

- Gylfason, H. F., Arnardottir, A. A., and Kristinsson, K. (2013). More on gender differences in lying. *Economics Letters* 119(1): 94–6.
- Halali, E., Bereby-Meyer, Y., and Ockenfels, A. (2013). Is it all about the self? The effect of self-control depletion on ultimatum game proposers. *Frontiers in Human Neuroscience* 7: 240.
- Haley, K. J. and Fessler, D. M. T. (2005). Nobody's watching? Subtle cues affect generosity in an anonymous economic game. *Evolution and Human Behavior* 26(3): 245–56.
- Hamermesh, D. and Biddle, J. (1994). Beauty and the labor market. *American Economic Review* 84(5): 1174–94.
- Hamilton, W. D. (1964). The genetical evolution of social behavior. *Journal of Theoretical Biology* 7(1): 1–52.
- Hannan, R. L., Kagel, J. H., and Moser, D. V. (2002). Partial gift exchange in experimental labor markets: impact of subject population differences, productivity differences and effort request on behavior. *Journal of Labor Economics* 20(4): 923–51.
- Hargreaves Heap, S. P., Tan, J. H. W., and Zizzo, D. J. (2013). Trust, inequality and the market. *Theory and Decision* 74(3): 311–33.
- Hargreaves-Heap, S. P. and Varoufakis, Y. (2002). Some experimental evidence on the evolution of discrimination, cooperation and the perception of fairness. *Economic Journal* 112(481): 679–703.
- Hargreaves-Heap, S. P., Verschoor, A., and Zizzo, D. J. (2012). A test of the experimental method in the spirit of Popper. *Journal of Economic Methodology* 19(1): 63–76.
- Hargreaves Heap, S. P. and Zizzo, D. J. (2009). The value of groups. *American Economic Review* 99(1): 295–323.
- Harrison, G. W. and McCabe, K. A. (1996). Expectations and fairness in a simple bargaining experiment. *International Journal of Game Theory* 25(3): 303–27.
- Hart, O. (1995). *Firms, Contracts, and Financial Structure*. Oxford: Oxford University Press.
- Henrich, J., Boyd, R., Bowles, S., et al. (2001). Cooperation, reciprocity and punishment in fifteen small-scale societies. *American Economic Review* 91: 73–78.
- Henrich, J., Ensminger, J., McElreath, R., et al. (2010). Markets, religion, community size, and the evolution of fairness and punishment. *Science* 327(5972): 1480–84.
- Herold, F. (2010). Contractual incompleteness as a signal of trust. *Games and Economic Behavior* 68(1): 180–91.
- Herrmann, B. and Thöni, C. (2009). Measuring conditional cooperation: a replication study in Russia. *Experimental Economics* 12(1): 87–92.
- Herrmann, B., Thöni, C., and Gächter, S. (2008). Antisocial punishment across societies. *Science* 319(5868): 1362–67.
- Heyman, J. and Ariely, D. (2004). Effort for payment: a tale of two markets. *Psychological Science* 15(11): 787–93.
- Hoffman, E., McCabe, K. A., Shachat, K., and Smith, V. L. (1994). Preferences, property rights, and anonymity in bargaining games. *Games and Economic Behavior* 7(3): 346–80.
- Hoffman, E., McCabe, K. A., and Smith, V. L. (1996). On expectations and the monetary stakes in ultimatum games. *International Journal of Game Theory* 25(3): 289–301.
- Hoffman, E., McCabe, K., and Smith, V. L. (2008). Preferences and property rights in ultimatum and dictator games. In: C. R. Plott and V. L. Smith (eds.), *Handbook of Experimental Economics Results, Volume 1*. Amsterdam: North-Holland, pp. 417–22.
- Hoffman, E. and Spitzer, M. L. (1982). The Coase theorem: some experimental tests. *Journal of Law and Economics* 25(1): 73–98.
- Holmås, T. H., Kjerstad, E., Lurås, H., and Straume, O. R. (2010). Does monetary punishment crowd out pro-social motivation? A natural experiment on hospital length of stay. *Journal of Economic Behavior and Organization* 75(2): 261–7.
- Holmstrom, B. and Milgrom, P. (1991). Multitask principal-agent analyses: linear contracts, asset ownership and job design. *Journal of Law, Economics and Organization* 7: 24–52.
- Hopfensitz, A. and Reuben, E. (2009). The importance of emotions for the effectiveness of social punishment. *Economic Journal* 119(540): 1534–59.
- Hopkins, E. (2008). Inequality, happiness and relative concerns: what actually is their relationship? *Journal of Economic Inequality* 6(4): 351–72.
- Huck, S., Lünser, G. K., and Tyran, J.-R. (2012). Competition fosters trust. *Games and Economic Behavior* 76(1): 195–209.

- Huck, S., Müller, W., and Normann, H.-T. (2001). Stackelberg beats Cournot—on collusion and efficiency in experimental markets. *Economic Journal* 111(474): 749–65.
- Hurkens, S. and Kartik, N. (2009). Would I lie to you? On social preferences and lying aversion. *Experimental Economics* 12(2): 180–92.
- Hwang, S.-H. and Bowles, S. (2011). A note on optimal incentives with state-dependent preferences. Sogang University, Research Institute for Market Economy, Working Paper 1118.
- Iriberry, N. and Rey-Biel, P. (2013). Elicited beliefs and social information in modified dictator games: what do dictators believe other dictators do? *Quantitative Economics* 4: 515–47.
- Isaac, R. M., McCue, K., and Plott, C. R. (1985). Public goods provision in an experimental environment. *Journal of Public Economics* 26(1): 51–74.
- Itoh, H. (2004). Moral hazard and other-regarding preferences. *Japanese Economic Review* 55(1): 18–45.
- Iturbe-Ormaetxe, I., Ponti, G., Tomás, J., and Ubeda, L. (2011). Framing effects in public goods: prospect theory and experimental evidence. *Games and Economic Behavior* 72(2): 439–47.
- Jex, S. M., and Britt, T. W. (2008). *Organizational Psychology: A Scientist-Practitioner Approach*. 2nd edition. New Jersey: John Wiley & Sons.
- Johnson, N. D. and Mislin, A. A. (2011). Trust games: a meta-analysis. *Journal of Economic Psychology* 32(5): 865–89.
- Kagel, J. H., Kim, C., and Moser, D. (1996). Fairness in ultimatum games with asymmetric information and asymmetric payoffs. *Games and Economic Behavior* 13(1): 100–10.
- Kahneman, D., Knetsch, J. L., and Thaler, R. H. (1986). Fairness as a constraint on profit seeking: entitlements in the market. *American Economic Review* 76(4): 728–41.
- Karlan, D. (2005). Using experimental economics to measure social capital and predict financial decisions. *American Economic Review* 95(5): 1688–99.
- Kartik, N. (2009). Strategic communication with lying costs. *Review of Economic Studies* 76(4): 1359–95.
- Kartik, N., Ottaviani, M., and Squintani, F. (2007). Credulity, lies, and costly talk. *Journal of Economic Theory* 134(1): 93–116.
- Keser, C. and Van Winden, F. (2000). Conditional cooperation and voluntary contributions to public goods. *Scandinavian Journal of Economics* 102(1): 23–39.
- Knoch, D., Pascual-Leone, A., Meyer, K., Treyer, V., and Fehr, E. (2006). Diminishing reciprocal fairness by disrupting the right prefrontal cortex. *Science* 314(5800): 829–32.
- Kocher, M. G., Cherry, T., Kroll, S., Netzer, R. J., et al. (2008). Conditional cooperation on three continents. *Economics Letters* 101(3): 175–8.
- Koford, K. and Penno, M. (1992). Accounting, principal-agent theory, and self-interested behavior. In N. E. Bowie and R. E. Freeman (eds.), *Ethics and Agency Theory: An Introduction*. New York: Oxford University Press, pp. 127–42.
- Kosfeld, M. and Neckermann, S. (2011). Getting more work for nothing? Symbolic awards and worker performance. *American Economic Journal: Microeconomics* 3(3): 86–99.
- Kosfeld, M., Okada, A., and Riedl, A. (2009). Institution formation in public goods games. *American Economic Review* 99(4): 1335–55.
- Kreps, D. and Wilson, R. (1982). Reputation and imperfect information. *Journal of Economic Theory* 27: 280–312.
- Krueger, A. B. and Mas, A. (2004). Strikes, scabs, and tread separations: labor strife and the production of defective bridgestone/firestone tires. *Journal of Political Economy* 112(2): 253–89.
- Kube, S., Maréchal, M. A., and Puppe, C. (2006). Putting reciprocity to work positive versus negative responses in the field. University of St. Gallen, Economics, Discussion Paper 2006–27.
- Kube, S., Maréchal, M. A., and Puppe, C. (2012). The currency of reciprocity: gift exchange in the workplace. *American Economic Review* 102(4): 1644–62.
- Kube, S., Maréchal, M. A., and Puppe, C. (2013). Do wage cuts damage work morale? Evidence from a natural field experiment. *Journal of the European Economic Association* 11(4): 853–70.
- Kurzban, R. and Houser, D. (2004). Experiments investigating cooperative types in human groups: a complement to evolutionary theory and simulations. *Proceedings of the National Academy of Sciences of the United States of America* 102(5): 1803–7.
- Lacetera, N., Macis, M., and Slonim, R. (2012). Will there be blood? Incentives and displacement

- effects in pro-social behaviour. *American Economic Journal: Economic Policy* 4(1): 186–223.
- Lambert, P. J. (2001). *The Distribution and Redistribution of Income*. 3rd edition. Manchester: Manchester University Press.
- Ledyard, J. O. (1995). Public goods: a survey of experimental research. In J. H. Kagel and A. E. Roth (eds.), *Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press, pp. 111–94.
- Lee, D. and Rupp, N. G. (2007). Retracting a gift: how does employee effort respond to wage reductions. *Journal of Labor Economics* 25(4): 725–61.
- Leibbrandt, A. (2012). Are social preferences related to market performance? *Experimental Economics* 15(4): 589–603.
- Lemieux, T., MacLeod, W. B., and Parent, D. (2009). Performance pay and wage inequality. *Quarterly Journal of Economics* 124(1): 1–49.
- Levine, D. K. (1998). Modeling altruism and spitefulness in experiment. *Review of Economic Dynamics* 1(3): 593–622.
- Levitt, S. D. and List, J. A. (2007). What do laboratory experiments measuring social preferences reveal about the real world? *Journal of Economic Perspectives* 21(2): 153–74.
- Levitt, S. D. and List, J. A. (2008). Homo economicus evolves. *Science* 319: 909–10.
- Levitt, S. D., List, J. A., Neckermann, S., and Sadoff, S. (2012). The behavioralist goes to school: leveraging behavioral economics to improve educational performance. NBER, Working Paper 18165.
- Liberman, V., Samuels, S. M., and Ross, L. (2004). The name of the game: predictive power of reputations versus situational labels in determining prisoner's dilemma game moves. *Personality and Social Psychology Bulletin* 30(9): 1175–85.
- List, J. A. (2004). Young, selfish, and male: field evidence on social preferences. *Economic Journal* 114(492): 121–49.
- List, J. A. (2006). The behavioralist meets the market: measuring social preferences and reputation effects in actual transactions. *Journal of Political Economy* 114: 1–37.
- List, J. A. (2007). On the interpretation of giving in dictator games. *Journal of Political Economy* 115(3): 482–93.
- List, J. A. and Cherry, T. L. (2000). Learning to accept in ultimatum games: evidence from an experimental design that generates low offers. *Experimental Economics* 3(1): 11–29.
- List, J. A. and Cherry, T. L. (2008). Examining the role of fairness in high stakes allocation decisions. *Journal of Economic Behavior and Organization* 65(1): 1–8.
- Loewenstein, G. F. (2000). Emotions in economic theory and economic behavior. *American Economic Review* 90(2): 426–32.
- Loewenstein, G. F., Bazerman, M., and Thompson, L. (1989). Social utility and decision making in interpersonal contexts. *Journal of Personality and Social Psychology* 57(3): 426–41.
- Lotz, S., Schlösser, T., Cain, D. M., and Fetchenhauer, D. (2013). The (in)stability of social preferences: using justice sensitivity to predict when altruism collapses. *Journal of Economic Behavior and Organization* 93: 141–8.
- Mas, A. (2006). Pay, reference points, and police performance. *Quarterly Journal of Economics* 121(3): 783–821.
- Mas, A. (2008). Labor unrest and the quality of production: evidence from the construction equipment resale market. *Review of Economic Studies* 75(1): 229–258.
- Masclet, D., Noussair, C., Villeval, M.-C., and Tucker, S. (2003). Monetary and non-monetary punishment in the voluntary contributions mechanism. *American Economic Review* 93(1): 366–80.
- Mas-Collel, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. New York: Oxford University Press.
- Matusaka, J. G. (2005). Direct democracy works. *Journal of Economic Perspectives* 19(2): 185–206.
- Maximiano, S., Sloof, R., and Sonnemans, J. (2007). Gift exchange in a multi-worker firm. *Economic Journal* 117(522): 1025–50.
- Mazar, N., Amir, O., and Ariely, D. (2008). The dishonesty of honest people: a theory of self-concept maintenance. *Journal of Marketing Research* 45(6): 633–44.
- McCabe, K. A., Rassenti, S. J., and Smith, V. L. (1998). Reciprocity, trust, and payoff privacy in extensive form bargaining. *Games and Economic Behavior* 24(1–2): 10–24.
- McCabe, K. A., Rigdon, M. L., and Smith, V. L. (2003). Positive reciprocity and intentions in trust games. *Journal of Economic Behavior and Organization* 52(2): 267–75.

- McDermott, R. (2009). Psychological approaches to identity: definitions, measurement and experimentation. In R. Abdelal, Y. M. Herrera, A. I. Johnson, and R. McDermott (eds.), *Measuring Identity: A Guide for Social Science Research* Cambridge: Cambridge University Press, pp. 345–68.
- McGregor, D. (1960). *The Human Side of Enterprise*. New York: McGraw-Hill.
- Meier, S. (2007). Do subsidies increase charitable giving in the long run? Matching donations in a field experiment. *Journal of the European Economic Association* 5(6): 1203–22.
- Mellström, C. and Johannesson, M. (2008). Crowding out in blood donation: was Titmuss right? *Journal of the European Economic Association* 6(4): 845–63.
- Milgrom, P. and Roberts, J. (1982). Predation, reputation, and entry deterrence. *Journal of Economic Theory* 27: 280–312.
- Mobius, M. and Rosenblat, T. (2006). Why beauty matters. *American Economic Review* 90(1): 222–35.
- Muller, L., Sefton, M., Steinberg, R., and Vesterlund, L. (2008). Strategic behavior and learning in repeated voluntary contribution experiments. *Journal of Economic Behavior and Organization* 67(3–4): 782–93.
- Neilson, W. S. (2006). Axiomatic reference-dependence in behavior toward others and toward risk. *Economic Theory* 28(3): 681–92.
- Neiss, R. (1988). Reconceptualizing arousal: psychological states in motor performance. *Psychological Bulletin* 103(3): 345–66.
- Nikiforakis, N. and Normann, H.-T. (2008). A comparative statics analysis of punishment in public good experiments. *Experimental Economics* 11(4): 358–69.
- Ockenfels, A. and Werner, P. (2012). Hiding behind a small cake in a newspaper dictator game. *Journal of Economic Behavior and Organization* 82(1): 82–5.
- Ockenfels, A. and Werner, P. (2014). Beliefs and ingroup favoritism. *Journal of Economic Behavior and Organization*, in press.
- Offerman, T. (2002). Hurting hurts more than helping helps. *European Economic Review* 46(8): 1423–37.
- Okada, A. and Riedl, A. (2005). Inefficiency and social exclusion in a coalition formation game: experimental evidence. *Games and Economic Behavior* 50(2): 278–311.
- Ostrom, E., Walker, J., and Gardner, R. (1992). Covenants with and without a sword: self-governance is possible. *American Political Science Review* 86(2): 404–17.
- Owens, M. F. and Kagel, J. H. (2010). Minimum wage restrictions and employee effort in incomplete labor markets: an experimental investigation. *Journal of Economic Behavior and Organization* 73(3): 317–26.
- Page, T., Putterman, L., and Unel, B. (2005). Voluntary association in public goods experiments: reciprocity, mimicry, and efficiency. *Economic Journal* 115(506): 1032–53.
- Parrett, M. (2006). An analysis of the determinants of tipping behavior: a laboratory experiment and evidence from restaurant tipping. *Southern Economic Journal* 73(2): 489–514.
- Persson, T. and Tabellini, G. E. (2000). *Political Economics: Explaining Economic Policy*. Cambridge, MA: MIT Press.
- Pruckner, G. J. and Sausgruber, R. (2013). Honesty on the streets: a field study on newspaper purchasing. *Journal of the European Economic Association* 11(3): 661–79.
- Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review* 83(5): 1281–302.
- Reeson, A. F. and Tisdell, J. G. (2008). Institutions, motivations and public goods: an experimental test of motivational crowding. *Journal of Economic Behavior and Organization* 68(1): 273–81.
- Rege, M. and Telle, K. (2004). The impact of social approval and framing on cooperation in public good situations. *Journal of Public Economics* 88(7–8): 1625–44.
- Reich, P., Roberts, P., Laabs, N., et al. (2006). A randomized trial of blood donor recruitment strategies. *Transfusion* 46(7): 1090–6.
- Restivo, M. and van de Rijt, A. (2012). Experimental study of informal rewards in peer production. *PLoS ONE* 7(3): e34358.
- Rey-Biel, P. (2008). Inequity aversion and team incentives. *Scandinavian Journal of Economics* 110(2): 297–320.
- Riedl, A. and Vyrastekova, J. (2002). Social preference in three-player ultimatum game experiments. Tilburg University, Center for Economic Research, Discussion Paper 2002–5.

- Rotemberg, J. J. (2008). Minimally acceptable altruism and the ultimatum game. *Journal of Economic Behavior and Organization* 66(3–4): 457–76.
- Roth, A. E. and Erev, I. (1995). Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior* 8(1): 164–12.
- Roth, A. E., Prasnikar, V., Okuno-Fujiwara, M., and Zamir, S. (1991). Bargaining and market behavior in Jerusalem, Ijubljana, Pittsburgh, and Tokyo: an experimental study. *American Economic Review* 81(5): 1068–95.
- Rothschild, M. and Stiglitz, J. E. (1970). Increasing risk: I. A definition. *Journal of Economic Theory* 2(3): 225–43.
- Rothschild, M. and Stiglitz, J. E. (1971). Increasing risk: II. Its Economic Consequences. *Journal of Economic Theory* 3(1): 66–84.
- Rustagi, D., Engel, S., and Kosfeld, M. (2010). Conditional cooperation and costly monitoring explain success in forest commons management. *Science* 330(6006): 961–5.
- Ryan, R. M. and Deci, E. L. (2000). Intrinsic and extrinsic motivations: classic definitions and new directions. *Contemporary Educational Psychology* 25(1): 54–67.
- Sánchez-Pagés, S. and Vorsatz, M. (2007). An experimental study of truth-telling in a sender–receiver game. *Games and Economic Behavior* 61(1): 86–112.
- Schmitt, M., Baumert, A., Gollwitzer, M., and Maes, J. (2010). The justice sensitivity inventory: factorial validity, location in the personality facet space, demographic pattern, and normative data. *Social Justice Research* 23(2–3): 211–38.
- Schmitt, M., Neumann, R., and Montada, L. (1995). Dispositional sensitivity to befallen injustice. *Social Justice Research* 8(4): 385–407.
- Schotter, A., Weiss, A., and Zapater, I. (1996). Fairness and survival in ultimatum and dictatorship games. *Journal of Economic Behavior and Organization* 31(1): 37–56.
- Sefton, M., Shupp, R., and Walker, J. M. (2007). The effect of rewards and sanctions in provision of public goods. *Economic Inquiry* 45(4): 671–90.
- Segal, U. and Sobel, J. (2007). Tit for tat: foundations of preferences for reciprocity in strategic settings. *Journal of Economic Theory* 136(1): 197–216.
- Seinen, I. and Schram, A. (2001). Social status and group norms: indirect reciprocity in a helping experiment. University of Amsterdam, Department of Economics and Econometrics, Working Paper.
- Selten, R. and Ockenfels, A. (1998). An experimental solidarity game. *Journal of Economic Behavior and Organization* 34(4): 517–39.
- Serra-Garcia, M., van Damme, E., and Potters, J. (2013). Lying about what you know or about what you do? *Journal of the European Economic Association* 11(5): 1204–29.
- Shapiro, C. and Stiglitz, J. (1984). Equilibrium unemployment as a worker discipline device. *American Economic Review* 74(3): 433–44.
- Shih, M., Pittinsky, T. L., and Ambady, N. (1999). Stereotype susceptibility: identity salience and shifts in quantitative performance. *Psychological Science* 10(1): 80–3.
- Shu, L. L., Gino, F., and Bazerman, M. H. (2011). Dishonest deed, clear conscience: cheating leads to moral disengagement and motivated forgetting. *Personality and Social Psychology Bulletin* 37(3): 330–49.
- Simon, H. A. (1951). A formal theory of the employment relationship. *Econometrica* 19(3): 293–305.
- Sliwka, D. (2007). Trust as a signal of a social norm and the hidden costs of incentive schemes. *American Economic Review* 97(3): 999–012.
- Slonim, R. and Roth, A. E. (1998). Learning in high stakes ultimatum games: an experiment in the Slovak Republic. *Econometrica* 66(3): 569–96.
- Sloof, R. and Sonnemans, J. (2011). The interaction between explicit and relational incentives: an experiment. *Games and Economic Behavior* 73(2): 573–94.
- Smith, V. L. (1962). An experimental study of competitive market behavior. *Journal of Political Economy* 70(2): 111–37.
- Sobel, J. (2005). Independent preferences and reciprocity. *Journal of Economic Literature* 43(2): 392–436.
- Soetevent, A. R. (2005). Anonymity in giving in a natural context—a field experiment in 30 churches. *Journal of Public Economics* 89(11–12): 2301–23.
- Sonnemans, J., Schram, A., and Offerman, T. (1999). Strategic behavior in public good games:

- when partners drift apart. *Economics Letters* 62(1): 35–41.
- Stoop, J., Noussair, C., and van Soest, D. (2012). From the lab to the field: public good provision with fishermen. *Journal of Political Economy* 120(6): 1027–56.
- Tajfel, H., Billig, M. G., Bundy, R. P., and Flament, C. (1971). Social categorization and inter-group behavior. *European Journal of Social Psychology* 1(2): 149–78.
- Tajfel, H. and Turner, J. (1979). An integrative theory of intergroup conflict. In W. G. Austin and S. Worchel (eds.), *The Social Psychology of Intergroup Relations*. Monterey, CA: Brooks Cole.
- Tajfel, H. and Turner, J. (1986). The social identity theory of intergroup behavior. In W. G. Austin and S. Worchel (eds.), *The Psychology of Intergroup Relations*. Chicago: Nelson-Hall, pp. 7–24.
- Thaler, R. H. (1988). Anomalies: the ultimate game. *Journal of Economic Perspectives* 2(4): 195–206.
- Titmuss, R. M. (1971). *The Gift Relationship: From Human Blood to Social Policy*. New York: Pantheon Books.
- Trivers, R. L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology* 46(1): 35–57.
- Turner, J. C., Hogg, M. A., Oakes, P. J. Reicher, S. D., and Wetherell, M. S. (1987). *Rediscovering the Social Group: A Self-Categorization Theory*. Oxford and New York: Wiley & Sons.
- Tyran, J.-R. and Feld, L. P. (2006). Achieving compliance when legal sanctions are non-deterrent. *Scandinavian Journal of Economics* 108(1): 135–56.
- Tyran, J.-R. and Sausgruber, R. (2006). A little fairness may induce a lot of redistribution in a democracy. *European Economic Review* 50(2): 469–85.
- Vanberg, C. (2008). Why do people keep their promises? An experimental test of two explanations. *Econometrica* 76(6): 1476–80.
- van der Weele, J. J., Kulisa, J., Kosfeld, M., and Friebe, G. (2014). Resisting moral wiggle room: how robust is reciprocal behavior? *American Economic Journal: Microeconomics* 6(3): 256–64.
- van Dijk, F., Sonnemans, J., and van Winden, F. (2002). Social ties in a public good experiment. *Journal of Public Economics* 85(2): 275–99.
- Veblen, T. (1899). *The Theory of the Leisure Class. An Economic Study of Institutions*. Random House.
- Volk, S., Thöni, C., and Ruigrok, W. (2012). Temporal stability and psychological foundations of cooperation preferences. *Journal of Economic Behavior and Organization* 81(2): 664–76.
- Vrij, A. (2000). *Detecting Lies and Deceit: The Psychology of Lying and the Implications for Professional Practice*. New York: John Wiley & Sons.
- Vrij, A. (2008). Detecting lies and deceit: pitfalls and opportunities. In *Wiley Series in the Psychology of Crime, Policing and Law*. Chichester: John Wiley & Sons.
- Williams, G. C. (1966). *Adaptation and Natural Selection: A Critique of Some Current Evolutionary Thought*. Princeton, NJ: Princeton University Press.
- Winter, E. (2004). Incentives and discrimination. *American Economic Review* 94(3): 764–73.
- Xiao, E. and Houser, D. (2011). Punish in public. *Journal of Public Economics* 95(7–8): 1006–17.
- Yamagishi, T. (1986). The provision of a sanctioning system as a public good. *Journal of Personality and Social Psychology* 51(1): 110–16.
- Yamagishi, T. and Kiyonari, T. (2000). The group as the container of generalized reciprocity. *Social Psychology Quarterly* 63(2): 116–32.
- Zizzo, D. J. (2010). Experimenter demand effects in economic experiments. *Experimental Economics* 13(1): 75–98.

PART 3

Behavioral Time Discounting

INTRODUCTION TO PART 3

Decisions that have a time dimension are central to many economic problems facing firms, governments, and individuals, although our focus will primarily be on individual decisions. Examples of individual decisions that are considered in Part 3 include the amount and type of savings; timing of retirement; investment in liquid versus illiquid assets; purchase of an annual, or pay-as-you-go health club membership. There is a long history of economic thought devoted to analyzing the complex issues associated with temporal choices. The classical economists focused squarely on the psychological foundations of such decisions (Frederick et al., 2002; Loewenstein, 2007, Ch. 3).

The basic idea in most modern theories of time preference is the following. Delaying a reward can increase its size, e.g., due to the accrual of interest payments. However, delayed outcomes are less valuable to impatient individuals. This trade-off determines the optimal temporal choice of individuals. The basic unit of analysis is preferences over sets of outcome–time pairs, or sequences of consumption levels. The subset of models that are separable in time and outcome and allow for impatience is sometimes called *discounted utility* (DU) models.

Chapter 9 considers a special member of the class of DU models, the *exponential discounted utility* (EDU) model. The EDU model owes its origin entirely to a technical question: it is the unique member of the DU class that gives rise to *time consistent preferences*. However, the psychological foundations of the EDU model are fairly weak. A single parameter, the *discount rate*, captures all relevant psychological considerations. It is no surprise that a very large literature on the anomalies of the EDU model has sprung up. This would not have surprised its inventor (Samuelson, 1937), who argued that there was no presumption that individuals maximized a functional of the form proposed by EDU. Furthermore, he also cautioned against using it for welfare analysis. But his caveats were either forgotten or ignored.

Among the anomalies of the EDU model that we consider, are the following. The most serious is the rejection of constant discounting (*common difference effect*). In a famous example of the common difference effect that violates constant discounting, Richard Thaler proposed the following example. Most people would prefer one apple today to two apples tomorrow, but they prefer two apples in 51 days to one apple in 50 days. Clearly, if the discount rate were constant across any interval of fixed size, as in the EDU model, we would not observe this preference reversal.

Observed discount rates increase as a future reward is brought forward towards the present. Experiments on animals and humans show that discount rates are hyperbolically declining when plotted against time (Ainslie, 1975, 1992). This has led to the development of the *hyperbolic discounting utility model*. Losses are more salient than gains, in the sense that they are discounted less (*the sign effect*), greater magnitudes are discounted less than smaller ones (*the magnitude effect*), discounting over an interval depends on how that interval is subdivided (*subadditive discounting*), choices may appear to be cyclical (*apparent intransitivity of preferences*), and choices may depend on the *spread and shape of the consumption sequence* for two consumption sequences that have the same exponentially discounted utility.

Chapter 10 considers alternative explanations of the anomalies of the EDU model, and exposit a range of behavioral models of time discounting that have emerged to explain these anomalies. We explain the *magnitude effect* and the *sign effect* by putting appropriate restrictions on the

utility function. We also introduce the class of *simple increasing elasticity* utility functions that are able to operationalize these restrictions.

The explanation of the *common difference effect* has been a major endeavor in behavioral time discounting. *Hyperbolic discounting* is the leading explanation. Individuals who use hyperbolic discounting are more impatient over shorter horizons, giving rise to present-biased preferences. We consider the *general hyperbolic discounting function*, as well as the *quasi-hyperbolic* or (β, δ) form, that is popular in applied work. *Subadditivity*, the finding that splitting a time interval into subintervals reduces the extent of discounting is our second explanation of the common difference effect. We also consider the *Read–Scholten discount function*, which is one of the most flexible discount functions available. Finally, we consider the role of uncertainty and the interaction of risk preferences and time preferences in explaining the common difference effect.

In any choice between outcome–time pairs, there are two *attributes*, *outcome* and *time*. We consider three different attribute-based models. These are *the model of vague time preferences*, *the model of similarity relations*, and *the trade-off model of intertemporal choice*. These models typically prescribe a hierarchical choice among outcome–time pairs, based on primary, secondary, or tertiary criteria. However, the choices may be quite sensitive to the order in which the criteria are applied. These models will need further developments, but they can potentially explain phenomena such as the common difference effect, and apparent intransitivity of preferences.

The final model that we consider is a general theory of time discounting, the *reference time theory*, RT. RT combines generalized hyperbolic discounting with subadditive discounting. It also introduces a reference time in addition to a reference outcome level. RT can explain all the main anomalies of the EDU model and can incorporate recent developments in time discounting such as uncertainty and attribute models. Preferences in RT are complete and transitive, yet it is able to explain apparent intransitivity of preferences by shifts in the reference point for time.

While Chapter 10 highlights the present bias of humans, Chapter 11 considers applications of present-biased preferences, with particular emphasis on models of quasi-hyperbolic discounting. In the presence of self-control problems, it becomes critical to specify one's degree of awareness about future self-control problems. The convenient device of humans as *multiple selves* postulates any individual as a succession of multiple selves, one for each time period. All selves exhibit hyperbolic discounting (i.e., relatively greater impatience at shorter horizons), but may not be fully aware of their future self-control problems. The current self may then wish to take some costly actions that can mitigate the self-control problems of future selves. For instance, tying one's resources in illiquid assets to reduce the chances that future selves will splurge on consumption.

We begin by considering applications in *behavioral macroeconomics*. In particular, we consider a model of life-cycle consumption with hyperbolic discounting and limited awareness of future self-control problems. We derive a *behavioral Euler equation*. This framework can address several puzzles in the neoclassical model. These include the following. Why does consumption track income so closely? Why do individuals undersave for retirement? Why is there a sharp drop in consumption at retirement? Why do individuals hold illiquid assets and credit card debt simultaneously? Why have national savings rates been declining? Why is marginal propensity to consume asset-specific? Why is high consumption variability not associated with high consumption growth? We also clarify the relation between intertemporal substitution and risk.

Retirement is another important decision that falls in the domain of present-biased preferences. The present bias of future selves might induce them to retire too early relative to the preference of the current self, who is more patient at longer horizons. One solution is for the current self to withhold savings for the future selves that can potentially reduce the benefits of early retirement. However, should this not work, then the current self might be resigned to save too much

in order to counteract the early retirement decision of future selves. Employment legislation such as a mandatory retirement age may be an important commitment device that solves the problem of the current selves.

Chapter 11 also presents several other applications. We consider features of temporal reward and cost profiles that lead, in the presence of limited awareness, to *procrastination*, or *preproportionation*. Market competition between present-biased agents may reduce procrastination for low costs of procrastination, but give rise to multiple equilibria for intermediate costs. Current selves may, through harmful addictions such as smoking, create a negative externality for their own future selves. Sales taxes on cigarettes may then lead to an improvement in the welfare of individuals by reducing the stock of harmful habits. This is a variant of the standard *Pigouvian taxes* argument.

Consider *investment goods* that have a current cost and a future benefit such as gym memberships. Current selves may be willing to pay high gym membership fees in return for a lower price/visit in the future if they underestimate their future self-control problems. Indeed, in this case, it is optimal for the firm to charge a unit price below marginal costs. In the opposite case of *leisure goods*, there are current benefits, but future costs (e.g., credit card borrowing). We show that in this case, the opposite result holds, i.e., price exceeds marginal costs. In each of these cases, the variation in the price serves as a commitment device to reduce future self-control problems. We close the chapter with a range of applications of hyperbolic discounting that illustrate the rich domain of problems to which it can be applied.

CHAPTER 9

The Evidence on Temporal Human Choice

9.1 Introduction

In this chapter, we discuss the *discounted utility model* (DU) with particular emphasis on its most popular variant in economics, the *exponential discounted utility model* (EDU). Our main aim is to outline the basics of EDU and consider the empirical evidence on the model.

In Section 9.2 we discuss the DU model, due to Samuelson (1937), and define the concepts of *discount function*, *discount rate*, and *discount factor* in discrete and in continuous time. The DU model is often too general for making precise predictions. For instance, choices under the DU model could be *time inconsistent*, i.e., decision makers may wish to alter previously made plans. In the class of DU models, the EDU model uniquely gives *time consistent choices* (Strotz, 1956). EDU is the dominant model of time discounting in economics and it has proved extremely useful for many purposes. However, its psychological foundations are fairly limited. In fact, a single parameter, the *discount rate*, encapsulates all psychological considerations. Samuelson (1937, p. 156) himself realized that it is completely arbitrary to assume that individuals behave in accordance with the EDU model and he also expressed reservations about its use for welfare comparisons (p. 161): “In conclusion, any connection between utility as discussed here and any welfare concept is disavowed.” Despite Samuelson’s own reservations, the EDU model is still the dominant model used in welfare comparisons in economics.

Section 9.3 discusses the EDU model and outlines its properties. These properties include *time-invariant instantaneous utility function*, *consumption independence*, *independence of discounting from consumption*, *constant discounting*, *stationarity of preferences*, and *time consistency*. Section 9.3.2 shows that EDU is the unique functional form that gives rise to time consistent choices in the class of DU models (Strotz, 1956). Section 9.3.3 gives a brief note on the assumption of constant marginal utility in eliciting the discount rate under EDU. Section 9.3.4 briefly outlines the axiomatic foundations of EDU (Fishburn and Rubinstein, 1982).

Section 9.4 discusses the *anomalies* of EDU. Think of an anomaly as violating an axiom or an implication of EDU.¹ The anomalies of the EDU model pose a serious challenge to its continued use. Under EDU, the discount rate between two consecutive time periods (assuming discrete

¹ We begin Section 9.4 with a small note on the distinction between *refutations* (rejection of the axioms of EDU) and *anomalies* (rejection of the auxiliary assumptions of EDU). But in line with the practice in this literature, we use the term anomaly for both.

time), remains the same, irrespective of the location of these periods. This derives from the *stationarity axiom* of EDU (Fishburn and Rubinstein, 1982). However, this assumption is routinely violated. A violation of the stationarity axiom is known as the *common difference effect*. Consider the following well-known example.

Example 9.1 (*Thaler's (1981) apples example*): It is reasonable to assume that a person will prefer one apple today to two apples tomorrow. The same person will, however, in all likelihood, prefer two apples in 51 days to one apple in 50 days.

Such a pattern of preference reversals is not consistent with the assumption of stationarity in EDU. Stationarity implies that the discount rate is constant. By contrast, the evidence shows that when the time delay before a future reward is received is short, the per-period discount rate is high and the decision maker appears very impatient. However, when the time delay is long, the per-period discount rate is low and the decision maker appears more patient. This can explain the puzzle in Example 9.1. In particular, not only is the discount rate declining with delay, evidence has been given to show that it declines in a hyperbolic manner (Ainslie, 1975, 1992); this holds for humans and close primate relatives such as chimps. In this case, the decision maker is said to be using *hyperbolic discounting*. This is a leading explanation for *present-biased preferences* that is able to explain a wide range of important phenomena.

Other anomalies of EDU that we consider in Section 9.4, include the following. (i) *Sign effect*: Losses are discounted less as compared to gains; thus, losses are temporally more salient than gains (Thaler, 1981). (ii) *Magnitude effect*: Outcomes of greater magnitudes are discounted less than smaller ones. Thus, elicited discount rates for larger outcomes are lower (Thaler, 1981). (iii) *Subadditive discounting*: Discounting over an interval depends on how that interval is sub divided (Read, 2001). Sub dividing an interval reduces salience. (iv) *Intransitivity of preferences or cyclical preferences*: Time preferences often appear intransitive (Roelofsma and Read, 2000), although controlling for reference point, context, or frame, such intransitivities may be reduced (Gintis, 2009). (v) *Shapes of consumption profile*: Under EDU, if two consumption profiles have identical discounted utility, then their shape is irrelevant. However, keeping discounted utility fixed, individuals may have a preference for rising consumption profiles (Loewenstein and Sicherman, 1991), and individuals may violate *consumption independence*² (Loewenstein and Prelec, 1992; Mayer and Russell, 2005).

9.2 The discounted utility model

Suppose that time $t \in \Gamma = \{0, 1, 2, \dots, T\}$, where $0 < T \leq \infty$. Denote by $c_t \in C \subset R$, the real-valued consumption of an individual at time t ; c_t could be a scalar, or a vector that denotes a bundle of consumption goods. Let $\mathbf{c}_0 = (c_0, c_1, \dots, c_T)$ denote the *temporal consumption profile* of the individual at time $t = 0$. Denote by $C \subset R^{T+1}$, the set of all bounded consumption profiles. One popular utility representation of preferences over C is the *discounted utility model* (DU). It requires the existence of a utility function $U : C \rightarrow \mathbb{R}$ such that,

$$U(\mathbf{c}_0) = \sum_{j=0}^T D(j)u(c_j); \mathbf{c}_0 \in C. \quad (9.1)$$

² Under EDU, if one compares two consumption profiles and if at any time t , the outcomes in both profiles are identical then one may ignore the outcome at t in a comparison of the two profiles (consumption independence).

In DU models, $u : C \rightarrow \mathbb{R}$ is typically a time-invariant, bounded, instantaneous utility function (sometimes referred to as *felicity*). $D(t)$ is known as the *discount function* at time t . The continuous time analogue of (9.1) over the time interval $[0, T]$ can be found by letting each time period be divided into n equal intervals and letting $n \rightarrow \infty$. In this case,

$$U(c_0) = \int_{t=0}^T D(t)u(c_t) dt; c_0 \in C. \quad (9.2)$$

We assume that $D : [0, T] \rightarrow (0, 1]$.

The discount function, D , is often assumed to be decreasing in order to reflect *impatience*, i.e., individuals prefer to receive the same reward earlier rather than later. Hence, $D(\tau') > D(\tau)$ for $\tau' < \tau$. When D is differentiable, then, in the continuous case, impatience is captured by the condition $D'(t) < 0$.³ Further, it is also standard practice to normalize $D(0) = 1$.

Definition 9.1 (*Discount Rate*): The discount rate or the “pure rate of time preference” is defined as the rate at which the discount function declines over time. In discrete time, the discount rate at time t , $\theta(t)$, is given by

$$\theta(t) = - \left[\frac{D(t) - D(t-1)}{D(t)} \right]; t \in \Gamma. \quad (9.3)$$

Dividing each time period into n equal intervals and letting $n \rightarrow \infty$, the discount rate in continuous time is given by

$$\theta(t) = - \frac{D'(t)}{D(t)}; t \in [0, T]. \quad (9.4)$$

The multiplication by -1 ensures that $\theta(t) > 0$, thus, its magnitude reflects how fast the discount function falls; $\theta(t)$ has no units. A higher discount rate implies a greater impatience to consume in the current period.⁴

9.3 The exponential discounted utility model (EDU)

A special case of the DU model is the *exponential discounted utility* model (EDU). In the discrete case, the discount function under EDU takes the form

$$D(t) = \frac{1}{(1 + \theta)^t} = \delta^t, \text{ where } \delta = \frac{1}{1 + \theta}, t \in \Gamma, \theta \in \mathbb{R}, \quad (9.5)$$

where θ is the discount rate over a single time period (in many cases, we refer to it as the annual discount rate) and δ is known as the *discount factor*. $\theta > 0$ captures the common assumption of impatience in the EDU model. Letting each time period be divided into n equal intervals, the continuous analogue of (9.5) can be found using $\lim_{n \rightarrow \infty} \left(1 + \frac{\theta}{n}\right)^{nt}$, which gives⁵

³ In the special case of $D'(t) = 0$ the individual gives equal weight to the instantaneous utility at each date, i.e., there is no impatience.

⁴ The *discount factor* is defined as $\phi(t) = \left[\frac{D(t)}{D(t-1)} \right]$, where $t \in \Gamma$.

⁵ Here we use the well-known formula $e = \lim_{n \rightarrow \infty} \left(1 + \frac{1}{n}\right)^n$.

$$D(t) = e^{-\theta t}. \quad (9.6)$$

Under EDU, using (9.1), (9.5), the intertemporal utility of the consumer when time is discrete, is given by

$$U(c_0) = \sum_{t=0}^T \delta^t u(c_t), \text{ where } \delta = \frac{1}{1+\theta}, c_0 \in C. \quad (9.7)$$

Using (9.2), (9.6), the continuous time analogue of (9.7) is

$$U(c_0) = \int_{t=0}^T e^{-\theta t} u(c_t) dt. \quad (9.8)$$

Our choice between the discrete and the continuous case will be based largely on pedagogical grounds, and on the norm in a particular literature in order to ease transition from the book to the relevant papers.

The EDU model allows one to compare distinct consumption profiles. So, for instance, under EDU, the discrete consumption profile $c^a = (c_0^a, c_1^a, \dots, c_T^a)$ is preferred to the consumption profile $c^b = (c_0^b, c_1^b, \dots, c_T^b)$, if and only if, $U(c^a) \geq U(c^b)$, or

$$\sum_{t=0}^T \delta^t u(c_t^a) \geq \sum_{t=0}^T \delta^t u(c_t^b). \quad (9.9)$$

If the individual is asked to choose a consumption profile from some set $A \subset C$, then it is assumed that he chooses the consumption profile $c^* \in A$ that satisfies the condition $U(c^*) \geq U(c')$ for all $c' \in A$.

9.3.1 *Some salient properties of the EDU model*

We now highlight some salient properties of the EDU model and introduce some of the terminology used in the debates on the EDU model.

1. *Stationarity of the felicity:* The instantaneous utility, u , is time invariant.
2. *Consumption independence:* Using (9.7), the marginal rate of substitution between two consecutive time periods t and $t+1$ is $\frac{-u'(c_t)}{\delta u'(c_{t+1})}$, which is independent of any other time periods.
3. *Utility independence:* The numerical values associated with distinct consumption profiles determine which is preferred (e.g., see (9.9)). Provided that the numerical value is identical, there is no preference associated with particular shapes of the consumption profiles, for instance, increasing, or decreasing.
4. *Independence of discounting from consumption:* The same discount function applies to an individual's consumption of any two distinct goods (say, oranges and apples). Furthermore, EDU does not specify how and why the discount function could depend on individual-specific traits or behaviors.
5. *Constant discounting:* The discount rate between any two consecutive periods is θ , irrespective of the time-location of the periods. This is known as the assumption of constant discounting.
6. *Diminishing marginal utility and positivity of the discount factor:* For most practical purposes, it is assumed that the felicity, u , is concave (diminishing marginal utility), and the

discount rate is strictly positive, i.e., $\theta > 0$ (this reflects impatience). Diminishing marginal utility creates an incentive to postpone consumption; however, a positive discount rate pushes in the direction of current consumption. This is the basic trade-off in EDU that determines the choice between current and future consumption.

7. *Stationarity and the common difference effect*: Suppose that we begin with a consumption profile $c^a = (c_0^a, c_1^a, \dots, c_T^a)$. Perturb this profile by altering the individual's consumption by x and $-y$ at, respectively, time periods 0 and $t > 0$. This gives us a new consumption profile $c^b = (c_0^a + x, \dots, c_{t-1}^a, c_t^a - y, c_{t+1}^a, \dots, c_T^a)$. Using (9.9), profile c^b is at least as good as profile c^a if and only if

$$u(c_0^a + x) + \delta^t u(c_t^a - y) \geq u(c_0^a) + \delta^t u(c_t^a),$$

or,

$$u(c_0^a + x) - u(c_0^a) \geq \delta^t [u(c_t^a) - u(c_t^a - y)]. \quad (9.10)$$

Hence, the preference for a consumption adjustment at two dates depends only on the time difference t that separates them. It is completely independent of the shape of the remaining consumption profile; this is the implication of *stationarity*. Violation of this effect is also known as the *common difference effect*.⁶

9.3.2 Time consistency of consumption plans under EDU

A consumption profile constructed at time t is said to be *time consistent* if it continues to be optimal when reconsidered at any later time $t+k$, $k > 0$. We now derive the sufficient conditions for a time consistent plan under DU, due to Strotz (1956). Suppose that the interest rate is zero, and an individual has the endowment, y_0 . The individual maximizes (9.2) at time $t = 0$,

$$\int_{t=0}^T D(t) u(c_t) dt, \quad (9.11)$$

where $D(t)$ is a, as yet unspecified, discount function that only depends on time and u is a time-invariant, increasing, concave, utility function. The intertemporal budget constraint of the individual specifies that the present discounted value of lifetime consumption is no greater than the initial endowment, y_0 . Recalling that the interest rate is zero, we can write this constraint as

$$\int_{t=0}^T c_t dt \leq y_0. \quad (9.12)$$

The relevant Lagrangian is given by

$$\int_{t=0}^T D(t) u(c_t) dt + \mu \left(y_0 - \int_{t=0}^T c_t dt \right), \quad (9.13)$$

where μ is the Lagrangian multiplier. Maximizing the Lagrangian, pointwise, with respect to consumption at time t , c_t , the first order condition is

⁶ There is some confusion in the terminology in the literature. Some authors refer to stationarity and the common difference effect as the same things. However, it is the lack of stationarity that we will, in keeping with the standard usage, refer to as the common difference effect.

$$D(t)u_c(c_t) - \mu = 0; \quad 0 \leq t \leq T, \quad (9.14)$$

where u_c is the marginal utility of consumption. Differentiating both sides of (9.14) with respect to time, t , we get

$$D_t(t)u_c(c_t) + D(t)u_{ct}(c_t) = 0, \quad (9.15)$$

where D_t is the time derivative of the discount function, and u_{ct} is the time derivative of marginal utility. From (9.15), we get that

$$\frac{D_t(t)}{D(t)} = -\frac{u_{ct}(c_t)}{u_c(c_t)}, \quad 0 \leq t \leq T. \quad (9.16)$$

Using (9.12), (9.16), one can generate the optimal consumption profile, conditional on any initial endowment.

The optimal plan, so far, was generated from the perspective of period 0. Now suppose that we allow the individual to reoptimize at some future date $\tau > 0$. For any $t > \tau$, the analogue of condition (9.16) is given by

$$\frac{D_t(t - \tau)}{D(t - \tau)} = -\frac{u_{ct}(c_t)}{u_c(c_t)}, \quad \tau \leq t \leq T. \quad (9.17)$$

If the optimal plan for time t made at two different times, 0 and $\tau > 0$, $\tau \leq t \leq T$, is identical, then such a plan is said to be *time consistent*. What restrictions on the discount function do we need to ensure time consistency of plans? The optimal plan constructed from the perspective of the initial time period is time consistent (using (9.16), (9.17)), if and only if

$$\frac{D_t(t - 0)}{D(t - 0)} = \frac{D_t(t - \tau)}{D(t - \tau)}, \quad 0 \leq \tau \leq t \leq T. \quad (9.18)$$

Since (9.18) is an “if and only if” condition, it is also the unique condition for time consistency.

$$(9.18) \Leftrightarrow \frac{d \ln D(t)}{dt} = \text{constant}. \quad (9.19)$$

The *unique* solution to (9.19) is the exponential discount function, i.e.,

$$D(t) = e^{-\theta t}; \theta \in \mathbb{R}. \quad (9.20)$$

It can be readily checked that solution (9.20) satisfies (9.19). Furthermore, it is the unique functional form that does so. No restrictions are placed on the sign of θ . Thus the decision maker may be impatient ($\theta > 0$), or patient ($\theta < 0$). Since EDU uses (9.20), optimal consumption profiles under EDU are time consistent.

9.3.3 A note on eliciting time preferences under EDU

In testing the EDU model, one often needs to elicit the time preference of individuals. The discussion in this section is motivated by Fisher (1930). Suppose that we have a dated consumption

profile $\{c_t\}_{t=0}^T$. The instantaneous utility at time t is given by $u(c_t)$. Suppose that the individual is given an extra amount of money, m , expressed in units of consumption, at time t . Then, using a first order Taylor series approximation around c_t , we get

$$u(c_t + m) \cong u(c_t) + u'(c_t)m. \quad (9.21)$$

Suppose that the individual is indifferent between receiving an additional amount of money m at time t and receiving an additional amount $z_0(m, t)$ now, i.e.

$$(c_0 + z_0(m, t), c_1, \dots, c_t, \dots, c_T) \sim (c_0, c_1, \dots, c_t + m, \dots, c_T). \quad (9.22)$$

Denote by $\tilde{x}$, the estimated value of some variable, x . In experimental practice, the estimate of the EDU discount function in (9.5), denoted by $\tilde{D}(t)$, is given by

$$\tilde{D}(t) = \tilde{\delta}^t = \frac{z_0(m, t)}{m}. \quad (9.23)$$

Under what conditions is $\tilde{\delta}^t$ a valid estimate of δ^t ? Using the stationarity property of EDU (see above), the preference relation stated in (9.22) can be written as

$$u(c_0 + z_0(m, t)) + \delta^t u(c_t) = u(c_0) + \delta^t u(c_t + m), \quad (9.24)$$

or,

$$u(c_0 + z_0(m, t)) - u(c_0) = \delta^t [u(c_t + m) - u(c_t)]. \quad (9.25)$$

Using (9.21), and the first order Taylor series approximations of the terms $u(c_0 + z_0(m, t))$ and $u(c_t + m)$, we can rewrite (9.25) as

$$u'(c_0)z_0(m, t) = \delta^t u'(c_t)m,$$

or,

$$\delta^t \frac{u'(c_t)}{u'(c_0)} = \frac{z_0(m, t)}{m} \equiv \tilde{\delta}^t. \quad (9.26)$$

From (9.26), $\tilde{\delta}^t = \delta^t$, if only if

$$u'(c_t) = u'(c_0). \quad (9.27)$$

Equation (9.27) will hold only if the base level of consumption c_0 is comparable to c_t . If an individual expects his circumstances to alter significantly over this time, then the marginal utilities at dates 0 and t are not comparable, so $\tilde{\delta}^t$ is unlikely to be a good estimate of δ^t .⁷ A second, implicit, assumption is that the individual uses the money m at date t for consumption purposes only. So issues of intertemporal substitution are ignored.

⁷ These issues are exploited by Noor (2007).

9.3.4 Axiomatic foundations of EDU

In order to place the refutations of EDU in context, we first outline the axiomatic derivation of the EDU model (Fishburn and Rubinstein, 1982). Suppose that at time zero, $Z \subset \mathbb{R}$ is the set of possible outcomes, and $\bar{\Gamma} \subset \mathbb{R}_+$ is the set of corresponding times at which these outcomes are realized.

Elements of the outcome set Z are denoted by z or z_j , $j = 1, 2, 3, \dots$, and elements of the set $\bar{\Gamma}$ are denoted by t or t_j , $j = 1, 2, 3, \dots$. Denote by (z, t) an *outcome–time pair* such that the individual is *guaranteed* to receive the outcome z at time t .⁸ Let $\succsim$ denote the binary preference relation “at least as good as,” on the Cartesian product $Z \times \bar{\Gamma}$. Also, let $\succ, \sim$ denote, respectively, the binary relations “strictly preferred to” and “indifferent to.” Consider the following five axioms, A1–A5, in Fishburn and Rubinstein (1982); we term these as the *axioms of rationality for time discounting*.

Definition 9.2 (*Axioms of rationality for time discounting*): Let $z_m, z_n \in Z$, and $t_p, t_q \in \bar{\Gamma}$. Let $\succsim$ denote a binary preference relation on $Z \times \bar{\Gamma}$.

- A1. *Ordering*: $\succsim$ is reflexive, complete, and transitive.
- A2. *Monotonicity*: If $z_m > z_n$, then $(z_m, t) \succ (z_n, t)$ for any $t \in \bar{\Gamma}$, i.e., larger rewards are preferred if the time horizon is kept fixed.
- A3. *Impatience*: Let $t_q < t_p$, and $z > 0$, $z \in Z$. Then, $(z, t_q) \succ (z, t_p)$, i.e., the same outcome is preferred at an earlier date. If $z_j = 0$, then $(z_j, t_q) \sim (z_j, t_p)$, i.e., there is no time preference for zero rewards.
- A4. *Continuity*: Fix some outcome–time pair $(z_j, t_j) \in Z \times \bar{\Gamma}$. Then, $\{(z, t) : (z, t) \succsim (z_j, t_j)\}$ is a closed set. Similarly, $\{(z, t) : (z_j, t_j) \succsim (z, t)\}$ is also a closed set.
- A5. *Stationarity*: If $(z_m, t_q) \sim (z_n, t_p)$, then $(z_m, t_q + t) \sim (z_n, t_p + t)$, for any $t \in \bar{\Gamma}$, i.e., indifference between two outcomes depends only on the length of the time difference between the two.

Axioms A1–A4 are standard in economics and are typically used to show the existence of a well-defined utility function (Varian, 1992, p. 97). Axiom A5 is non-standard and crucial to producing the additive separability under EDU. Fishburn and Rubinstein (1982, p. 681) write: “However, we know of no persuasive argument for stationarity as a psychologically viable assumption.”⁹ The EDU preference representation is given next.

Proposition 9.1 (Fishburn and Rubinstein, 1982, p. 682): If A1–A5 hold, then for any $\delta \in (0, 1)$, there exists a continuous, increasing, real-valued function, u , defined on the domain Z such that:

- (i) For all $(z_m, t_q), (z_n, t_p)$ in $Z \times \bar{\Gamma}$, $(z_m, t_q) \succsim (z_n, t_p)$ if and only if

$$\delta^{t_q} u(z_m) \geq \delta^{t_p} u(z_n).$$

- (ii) The outcome–time pair $(0, t)$ is treated as $\delta^t u(0) = 0$.
- (iii) Given $\delta, u(\cdot)$ is unique up to multiplication by a positive constant.

⁸ While Fishburn and Rubinstein (1982) provide an axiomatization in terms of outcome–time space, Koopmans (1960) and Koopmans et al. (1964) provide an axiomatization in terms of time preferences over consumption streams.

⁹ Fishburn and Rubinstein (1982) consider other axioms in response to this problem such as “Thomsen separability.” However, it is not clear if these have significantly better grounding in psychological evidence.

9.4 Anomalies of the EDU model

Strictly speaking, a *refutation* or *rejection* of the EDU model arises if the evidence contradicts any of the axioms of EDU. However, in applications, the EDU model is typically used in conjunction with auxiliary assumptions in order to derive testable predictions. It might not be clear if the evidence rejects one of the axioms of EDU, or a set of auxiliary assumptions, or both. In this case, one may refer to *anomalies* or *puzzles* of the EDU model.

Ultimately, it is a form of defensive methodology to term rejections of the EDU model as puzzles or anomalies, subject to at least two caveats. (i) The rejections are replicated sufficiently under stringently controlled conditions, and (ii) all reasonable confounding factors have been ruled out. Likewise it is an exaggeration to term anomalies and puzzles as rejections. What may seem like a rejection at one point may be revealed to be a quirk of experimental design, or there could be other factors that are at variance with the assumptions of the relevant theory. For instance, subjects in experiments who choose between outcome–time pairs may be *uncertain* that the experimenter will honor a promise of a future reward. One may then be led to a more general theory that takes such concerns into account. We do consider these extensions in Part I and Section 10.3.3 below. The literature has used the terms such as refutations, violations, anomalies, and puzzles interchangeably. In the behavioral time discounting literature, the preferred catch-all term is *anomaly*, a practice we follow below, although this is hardly satisfactory. However, we shall also flag up instances where the evidence clearly refutes EDU.

In this section, we make the standard assumption that subjects do not integrate current temporal plans (say, the temporal choices offered in experiments) with existing temporal plans. This is an auxiliary assumption that is not an element of the axiomatic structure of EDU. Hence, refutations of EDU by the evidence may arise because this auxiliary assumption does not hold. However, *non-integration of plans*, or at least partial integration, is the typical experimental finding, so it is not an unreasonable auxiliary assumption; see for instance, Andersen et al. (2012). Fafchamps et al. (2014) find no evidence that subjects in experiments integrate asset returns in the lab with their real-life income/wealth, although winnings in one round may influence subsequent risk attitudes, which is reminiscent of the house money effect (see Part 7).

We now describe the anomalies of the EDU model. We will use the following terminology. An amount is “discounted less” (or is more *salient*) if the discount factor applied to such an amount is higher. For the EDU model, this corresponds to a higher value of δ , or equivalently, a lower value of θ (see (9.7)). The converse applies if we say that an amount is “discounted more.”

9.4.1 Hyperbolic discounting and the common difference effect

Evidence indicates that individuals exhibit a strong preference for immediate gratification that cannot be sufficiently accounted for by the impatience parameter, θ , under EDU. This desire for immediate gratification is so strong that it may lead to *preference reversals*, which cannot take place under EDU (see Section 9.3.2). Suppose that $z, z' \in Z$, and $t, \tau, t' \in \bar{T}$. The evidence indicates *preference reversals* of the following form.

$$(z, t) \succ (z', t') \text{ and } (z', t' + \tau) \succ (z, t + \tau); \tau > 0, t' > t, z' > z. \quad (9.28)$$

In (9.28), the decision maker prefers z to z' for one set of dates, but flips his preference when both dates are moved further away by the same amount. This contradicts the *constant discounting* feature of the EDU model. Preference reversals are a rejection of the EDU model because axiom

A5 on *stationarity* does not hold (see Definition 9.2). Following Loewenstein and Prelec (1992), we shall refer to a violation of the finding of constant discounting as the *common difference effect*.

Stated in terms of the EDU model, the preferences in (9.28) imply that

$$\delta^t u(z) \geq \delta^{t'} u(z') \text{ and } \delta^{t+\tau} u(z) \leq \delta^{t'+\tau} u(z'),$$

which cannot be jointly true if at least one of the preferences is strict. The preference in (9.28) holds, in particular, when $t = 0$ is the initial time period. This is the case for Thaler's (1981) apples example (see Example 9.1).

The evidence for preference reversals comes from several sources. Following Ainslie (1992), we choose here the three fold classification of the evidence.

1. When human and non-human subjects are given the choice between varying rewards with time varying delays (see e.g., (9.28)), then their choices are found by several researchers to be well described by the *hyperbolic discount function*

$$D(\tau) = \frac{1}{1 + k\tau}. \quad (9.29)$$

In (9.29), τ is the time delay between the current period, $t = 0$, and some future period $\tau > 0$, and $k > 0$ is a parameter. Using (9.4), (9.29), the corresponding discount rate per period is

$$\theta(\tau) = \frac{k}{1 + k\tau}, \quad (9.30)$$

which is plotted in Figure 9.1 for the case $k = 1.5$. Unlike EDU, which posits constant discount rates, hyperbolic discounting implies a pattern of hyperbolically declining discount rates. The implication is that the longer the time horizon before a reward is received, the lower is the per-period discount rate, so the more patient is the individual. The reader should not confuse the discount rate with the discount factor. From the perspective of period 0, the discount factor under EDU for an outcome received at time t is $\delta^t = \left(\frac{1}{1+\theta}\right)^t$, but the per-period discount rate is still θ , which is independent of time, unlike the discount rate under hyperbolic discounting in (9.30).

From (9.30), impatience is worst for very short horizons. Thus, the average implied annual discount rate under hyperbolic discounting over shorter horizons turns out to be larger (greater impatience) as compared to longer horizons. The hyperbolic discount function has been used extensively, and successfully, in behavioral and experimental psychology to fit animal and human behavior (Logue, 1988). Another natural candidate is the simplest hyperbolic function $D(t) = \frac{1}{t}$, used by Chung and Herrnstein (1967), Ainslie (1975), but (9.29) gives a better fit to the data.

Thaler (1981) asked subjects in an experiment, how much money in one month, one year, ten years, respectively, would make them indifferent to receiving \$15 now? The median response was \$20, \$50, and \$100, respectively. The implied annual discount rates for these three choices turn out to be 3.45, 1.2, 0.19, respectively.¹⁰ Hence, the longer the horizon, the

¹⁰ The implied annual discount rates can be calculated as follows. Suppose that individuals discounted at a constant rate. Then the first option implies that $15 = 20e^{-\theta/12}$, so the annual discount rate is $\theta = 3.45$. Proceeding analogously, the average annual discount rate implied by the other two choices can be found by solving $15 = 50e^{-\theta}$ and $15 = 100e^{-10\theta}$, which gives, respectively, $\theta = 1.2$ and $\theta = 0.19$.




Figure 9.1 A plot of the discount rate in (9.30) for the one-parameter hyperbolic discounting function with varying time delay.

smaller the implied annual discount rate (lower impatience). For some other supporting evidence, see Benzion et al. (1989), Loewenstein and Prelec (1992), and Laibson (1997). The evidence for hyperbolic discounting is reviewed in greater detail below.

2. In line with (9.28), when subjects in experiments are offered a choice between a smaller-sooner (SS) reward, and a larger-later (LL) reward available at some date in the future, then LL is preferred (i.e., $(z', t' + \tau) \succ (z, t + \tau)$). However, as the date of the SS reward approaches the present, the individual prefers the SS to the LL reward (i.e., $(z, t) \succ (z', t')$). The leading motivating example of this phenomenon is Thaler's (1981) apples example (Example 9.1). This pattern is found in several experimental studies. See Ainslie and Herrnstein (1981), Ainslie and Haendel (1983), Green et al. (1981), and Kirby and Herrnstein (1995).
3. When individuals have a choice between SS and LL rewards, they sometimes take actions to remove the SS choice from the menu, indicating a possible awareness of potentially harmful preference reversals in the future. This is also true for non-human subjects. We consider the possibility of such commitment devices below. Unless further auxiliary assumptions are made, such phenomena cannot be explained under EDU.

The evidence for hyperbolic discounting comes mainly from lab experiments. Chabris et al. (2008) consider the ability of lab evidence to predict 15 different kinds of field behavior such as BMI (body-mass index), smoking, drinking, exercise, saving, and gambling. They first measure the discount rates of individuals in the lab, using a *generalized hyperbolic function*, giving them alternative choices of the sooner-smaller, and larger-later type. They find that the correlation between the lab measured discount rates and field behavior is low. However, discount rates are still the most important determinants of field behavior when compared with other determinants such as age, depression, sex, cognitive ability, and education.

9.4.2 The gain-loss asymmetry

Under EDU, the discount rate is identical for gains and losses. Empirical evidence suggests that when choosing over outcome-time pairs, losses are more salient than gains. Thus, losses are

discounted less than gains, i.e., the discount rate applicable to losses is lower. Thaler (1981) estimated the discount rates for gains to be about three to ten times larger than the corresponding discount rates for losses. In fact, a substantial proportion of subjects prefer to incur a fixed loss immediately than to delay it.¹¹

The gain–loss asymmetry is not qualitatively inconsistent with the EDU model, or its axioms, provided we assume that u is concave (this is not required in the representation in Proposition 9.1). However, we show that the gain–loss asymmetry is quantitatively inconsistent with the EDU model. We follow Loewenstein and Prelec (1992) in illustrating this issue. Consider an initial consumption profile that provides consumption z at all points in time, i.e., $(z, z, \dots, z, \dots, z)$. Now construct a second consumption profile that gives the individual an extra z' units at time $t \geq 0$, i.e., $(z, z, \dots, z + z', \dots, z)$. We assume that the decision maker follows EDU, and has an increasing and concave utility function.

The effect on the individual's welfare can be found by using the method of *equivalent variation*. In this method, we adjust current consumption in the first profile by some monetary amount $q(z', t)$, so that the first and second consumption profiles yield the same intertemporal utility. Thus, under EDU, we get¹²

$$\text{Equivalent variation: } u(z + q) + \delta^t u(z) = u(z) + \delta^t u(z + z'). \quad (9.31)$$

Gain–loss asymmetry arises if the value of q is higher for losses relative to gains of an equal absolute amount, i.e., losses have relatively greater salience. From (9.31), solving out for q we get

$$q = u^{-1} [\delta^t u(z + z') + (1 - \delta^t) u(z)] - z. \quad (9.32)$$

The utility function is assumed to be concave (which is absent in the representation theorem in Proposition 9.1), and $\delta \in (0, 1)$ (which is a part of Proposition 9.1), hence,

$$\delta^t u(z + z') + (1 - \delta^t) u(z) \leq u(\delta^t(z + z') + (1 - \delta^t)z). \quad (9.33)$$

Since u^{-1} is an increasing function, using (9.32), (9.33), it follows that

$$\begin{aligned} u^{-1} [\delta^t u(z + z') + (1 - \delta^t) u(z)] &\leq u^{-1} [u(\delta^t(z + z') + (1 - \delta^t)z)] \\ &= z + \delta^t z'. \end{aligned} \quad (9.34)$$

From (9.32), (9.34) we get $q \leq z + \delta^t z' - z$, or

$$q \leq \delta^t z'. \quad (9.35)$$

Equation (9.31) must hold for all $t \geq 0$. Hence, for $z' > 0$, it must be that $q > 0$, while for $z' < 0$ it must be that $q < 0$. It follows from (9.35) that for gains, i.e., $z' > 0$ and $q > 0$,

$$\frac{q}{z'} = \left| \frac{q}{z'} \right| \leq \delta^t. \quad (9.36)$$

¹¹ For related findings, see for instance, Benzon et al. (1989), Loewenstein (1987), and Mischel et al. (1969).

¹² By contrast, under the method of *compensating variation*, we alter current consumption in the second consumption profile by an amount p so that the individual's intertemporal utility under the two consumption profiles is identical, i.e., $u(z) + \delta^t u(z) = u(z - p) + \delta^t u(z + z')$.

Analogously, for losses, i.e., $z' < 0$ and $q < 0$,

$$-q > -\delta^t z' \Rightarrow \frac{-q}{-z'} = \left| \frac{q}{z'} \right| \geq \delta^t. \quad (9.37)$$

Hence, from (9.36), (9.37), the required current compensation per unit is larger for losses of an equal absolute amount. Therefore, losses have relatively greater salience and EDU makes the correct qualitative prediction for the gain–loss asymmetry, provided we impose concavity of u , which is not a condition of the representation theorem in Proposition 9.1. We now show how this prediction is quantitatively wrong.

In going from (9.32) to (9.35), we use the concavity of the utility function; however, for perturbations of the initial consumption profile, for relatively small magnitudes (small z'), the utility function is likely to be approximately linear over the interval $(z - z', z + z')$. Hence, we cannot change the equality into an inequality as we move from (9.32) to (9.35). Therefore, the analogue of (9.35) for small perturbations is

$$q = \delta^t z'.$$

Thus, the gain–loss asymmetry disappears for small z' . Experimental evidence, however, suggests the opposite, i.e., the gain–loss asymmetry becomes stronger for small outcomes (Thaler, 1981).

9.4.3 The absolute magnitude effect

Under EDU, the discount rate is independent of the magnitudes of the outcomes. However, the evidence shows that larger monetary amounts are discounted less, i.e., the implied discount rate for larger amounts is lower. In Thaler (1981), subjects reported being indifferent to \$15 received now, and \$60 in a year's time; \$250 received now, and \$350 received in a year's time; and \$3000 received now, and \$4000 in a year's time. These choices imply the respective annual discount rates 1.39, 0.34, and 0.29. Thus, larger amounts are discounted less. Calculation of these discount rates is straightforward; for instance, $3000 = 4000e^{-\theta}$, so $\theta = \ln(4/3) \cong 0.29$. In other words, larger amounts have greater salience in an intertemporal consumption profile.

Suppose that the decision maker is risk neutral, and exhibits the following outcome–time choices: $(z_1, 0) \sim (z'_1, t)$, $(z_2, 0) \sim (z'_2, t)$, where $z'_2 > z'_1$. From the discussion in Section 9.3.3, we know that if marginal utilities are approximately constant (i.e., the outcomes are small), then, under EDU, the estimated discount factor is $\delta^t = \frac{z_1}{z'_1} = \frac{z_2}{z'_2}$. But the experimental evidence on the magnitude effect shows that $\frac{z_1}{z'_1} < \frac{z_2}{z'_2}$. This suggests that larger amounts are discounted less, and that the discount rate is magnitude dependent.¹³ We shall see below that the magnitude effect is also inconsistent with the hyperbolic discounting model, and generally with the DU model, unless one makes further assumptions on the utility function.

9.4.4 The delay–speedup asymmetry

For 2012, online e-sales exceeded 1 trillion USD, which exceeds the GDP of many countries. The following scenario is characteristic of such sales, say, from Amazon, or eBay. At date 0, Sam and

¹³ For the evidence, see Ainslie and Haendel (1983), Benzion et al. (1989), Green et al. (1994), Kirby et al. (1999), and Loewenstein (1987).

Dan have separately purchased an item online, and need to decide on a shipping option. Sam's shipping is scheduled in the near future, say, at time $s > 0$, but he has the option to delay his delivery to a future date, say, date $d > s$; in return, he receives some payment $\pi_{s \rightarrow d}$ (perhaps, a reduced shipping cost). Alternatively, Dan's shipping is scheduled for the later date, d , but he has the option to expedite his delivery to an earlier date, say, s , for an additional charge $\pi_{s \leftarrow d}$, due to more costly postage.

Suppose that we make two assumptions: (1) Individuals follow EDU, and (2) do not integrate new plans with existing plans; evidence, some reviewed above, seems to support the assumption of non-integration of plans.¹⁴ Under assumptions 1 and 2, it can be shown that $\pi_{s \rightarrow d} = \pi_{s \leftarrow d}$ (see below).

Loewenstein (1988) empirically tested to check if $\pi_{s \rightarrow d} = \pi_{s \leftarrow d}$. Subjects were asked the maximum amount they would pay for a Sony VCR with remote control (list price \$300) delivered today. They were then asked what minimum amount they would be willing to accept today to defer their receipt for the item by one year ($\pi_{s \rightarrow d}$)? In another question, subjects were asked the maximum amount they would pay for a Sony VCR with remote control (list price \$300) delivered a year from now. They were then asked what minimum amount they would be willing to pay today to speed up their receipt for the item by one year ($\pi_{s \leftarrow d}$)? It was found that $\pi_{s \rightarrow d} = \$126$, and $\pi_{s \leftarrow d} = \$54$. Following Loewenstein (1988), extensive empirical evidence shows that $\pi_{s \rightarrow d} > \pi_{s \leftarrow d}$.¹⁵ This gives rise to the *delay-speedup asymmetry*.

We now give a slightly more formal treatment. Consider three dates: $0 = r < s < d$. At date $r = 0$, an individual with initial wealth w is willing to pay the *maximum price* $p_\tau \leq w$ to contractually receive c units of the good at delivery time $\tau = s, d$ (respectively, contract- s , and contract- d). Any reasonable DU model that incorporates impatience should require $p_s > p_d$. We assume *non-integration of plans*, so that this buying decision is not integrated with existing consumption plans.

Following Loewenstein (1988), we make the simplifying assumption that the object is fully consumed at the delivery date, τ .¹⁶ Let consumption at dates s and d be denoted by c_s and c_d . A generic *consumption-time profile* is given by $(c_0, 0), (c_s, s), (c_d, d)$. Under contract- s , the consumption profile is $(w - p_s, 0), (c, s), (0, d)$, while under contract- d it is $(w - p_d, 0), (0, s), (c, d)$.

Let u be a time-invariant utility function, and we normalize $u(0) = 0$. Under EDU, we have

$$U((c_0, 0), (c_s, s), (c_d, d)) = u(c_0) + u(c_s)\delta^s + u(c_d)\delta^d.$$

The utility of the individual from the three actions, (i) not buying, (ii) signing contract- s , and (iii) signing contract- d , is respectively denoted by U_0, U_0^s , and U_0^d , where

$$\begin{cases} U_0 = U_0((w, 0), (0, s), (0, d)) = u(w). \\ U_0^s = U_0((w - p_s, 0), (c, s), (0, d)) = u(w - p_s) + u(c)\delta^s. \\ U_0^d = U_0((w - p_d, 0), (0, s), (c, d)) = u(w - p_d) + u(c)\delta^d. \end{cases} \quad (9.38)$$

Since p_τ is the *maximum price* that the individual is willing to pay to delay delivery to date $\tau = s, d$, it follows that $U_0 = U_0^s = U_0^d$, or

¹⁴ Consider, for instance, the following quote from Loewenstein (1988, p. 200): "Rather than integrating delayed consumption plans it appears that people often represent future consumption options as gains and losses or as a deviations from some standard, such as past levels of consumption."

¹⁵ For other evidence in support of the delay-speedup asymmetry, see Benzion et al. (1989), Shelley (1993), Malkoc and Zauberman (2006), and McAlvanah (2010).

¹⁶ The results also go through if the object depreciates over time, leading to a diminishing stream of consumption.

$$u(w) = u(w - p_s) + u(c)\delta^s = u(w - p_d) + u(c)\delta^d. \quad (9.39)$$

Consider an individual who has signed contract- s , and is asked to state the *minimum delay premium*, $\pi_{s \rightarrow d}$, received at time $r = 0$, to have receipt of the object delayed to time $t = d$.

Denoting the resulting lifetime utility by $U_0^{s \rightarrow d}$,

$$U_0^{s \rightarrow d} = U_0((w - p_s + \pi_{s \rightarrow d}, 0), (0, s), (c, d)) = u(w - p_s + \pi_{s \rightarrow d}) + u(c)\delta^d. \quad (9.40)$$

From the definition of $\pi_{s \rightarrow d}$, we get $U_0^s = U_0^{s \rightarrow d}$, or

$$u(w - p_s) + u(c)\delta^s = u(w - p_s + \pi_{s \rightarrow d}) + u(c)\delta^d. \quad (9.41)$$

From (9.39) and (9.41), we have $u(w - p_s + \pi_{s \rightarrow d}) = u(w - p_d)$, hence,

$$\pi_{s \rightarrow d} = p_s - p_d > 0. \quad (9.42)$$

Now consider someone who has chosen contract- d , and is asked to state the *minimum speedup premium*, $\pi_{s \leftarrow d}$, paid at time $r = 0$, to speed up the receipt of the object to time $t = s$. Denoting the resulting lifetime utility by $U_0^{s \leftarrow d}$,

$$U_0^{s \leftarrow d} = U_0((w - p_d - \pi_{s \leftarrow d}, 0), (c, s), (0, d)) = u(w - p_d - \pi_{s \leftarrow d}) + u(c)\delta^s. \quad (9.43)$$

Given the definition of $\pi_{s \leftarrow d}$ we have $U_0^d = U_0^{s \leftarrow d}$, so

$$u(w - p_d) + u(c)\delta^d = u(w - p_d - \pi_{s \leftarrow d}) + u(c)\delta^s. \quad (9.44)$$

From (9.39) and (9.44), we get $u(w - p_d - \pi_{s \leftarrow d}) = u(w - p_s)$, hence,

$$\pi_{s \leftarrow d} = p_s - p_d. \quad (9.45)$$

From (9.42), (9.45), under EDU, and non-integration of plans,

$$\pi_{s \rightarrow d} = \pi_{s \leftarrow d} = p_s - p_d > 0. \quad (9.46)$$

Any departure from (9.46) may be termed as the *delay-speedup asymmetry*. As noted above, the empirical finding of $\pi_{s \rightarrow d} > \pi_{s \leftarrow d}$, illustrates the delay-speedup asymmetry, which is not consistent with EDU.

9.4.5 Front-end delay dependence

Suppose that at time 0, an individual expresses the preference $(z, \tau) \succsim (z', \tau + d)$, where $\tau \geq 0, d > 0$. How is discounting affected across the interval $[\tau, \tau + d]$, if the front-end delay, τ , is altered? Fixing d , Read and Roelofsma (2003) find that the discount factor increases as one pushes the interval into the future by increasing τ . In particular, they looked at successive 13 week intervals $[0, 13]$, $[13, 26]$, $[26, 39]$, $[39, 52]$, and asked subjects in experiments what they would be

willing to receive at time τ in lieu of 200 Dutch Guilders at time $\tau + 13$. The elicited discount factors over the four intervals of equal length were respectively, 0.33, 0.4, 0.5, 0.52. Front-end delay dependence violates axiom A5 on stationarity (see Definition 9.2).

9.4.6 Subadditive discounting

Consider a time interval $[0, \tau_k]$ that is split up into k subintervals $[0, \tau_1], [\tau_1, \tau_2], [\tau_2, \tau_3], \dots, [\tau_{k-1}, \tau_k]$, where $0 < \tau_1 < \tau_2 < \dots < \tau_k$. Denote the discount factor over the interval $[\tau_j, \tau_{j+1}]$ by δ_j^{j+1} . Then, *subadditive discounting* implies that the total discounting is greater, or the applied discount rate is higher, when outcomes are discounted over subintervals, i.e.,¹⁷

$$\delta_0^k > \delta_0^1 \delta_1^2 \dots \delta_{k-1}^k. \quad (9.47)$$

Subadditive discounting implies that a future event is more salient when discounted over the whole interval, rather than its constituent subintervals. Under EDU, by contrast, we have $\delta_0^k = \delta_0^1 \delta_1^2 \dots \delta_{k-1}^k$. Subadditive discounting is a rejection of the stationarity axiom of EDU. Read (2001) finds evidence for subadditive discounting.¹⁸ Subadditive discounting is consistent with cyclical, i.e., non-transitive choice. An exercise asks you to show this.

9.4.7 Cyclical or non-transitive choices

The finding of *cyclical* or *non-transitive* choices is not new. Tversky et al. (1990) find evidence for cyclical choices for about 15% of the subjects in experiments after controlling for framing effects. Roelofsma and Read (2000) find cyclical choices to be pervasive in their data. Other evidence can be found in Read (2001) and Rubinstein (2003).¹⁹ Cyclical choices violate ordering, the first axiom in EDU preferences (see Definition 9.2). Such choices are also inconsistent with the hyperbolic discounting model that we consider in detail below. But other theories that we consider below are compatible with cyclical choices. These include Ok and Masatlioglu's (2007) *relative discounting model*, Manzini and Mariotti's (2006) *vague time preferences theory*, Read's (2001) *subadditive discounting*, and Rubinstein's (2003) *similarity relationship* for decision making.²⁰

The main objection to cyclical choices has been a practical one. The presence of non-transitivities seriously compromises the ability of most economic models to make predictions. Ok and Masatlioglu (2007) argue, on the other hand, that non-cyclical choices are perfectly reasonable. They give the following example. Suppose that a decision maker is willing to delay an outcome for one period, if compensated by \$50. But the same decision maker is unwilling to accept a delay of two weeks in return for an extra compensation of \$150. Then (assuming that time is measured in units of weeks, and starts at $t = 0$) the following preferences are perfectly reasonable

$$(100, 0) \succ (250, 2) \succ (200, 1) \succ (100, 0). \quad (9.48)$$

¹⁷ Subadditive discounting is similar to the idea of Tversky and Koehler (1994) who showed that the compound probability of an event is higher if it is partitioned into a greater number of subevents.

¹⁸ For other sources of evidence on subadditive discounting, see Read and Roelofsma (2003), Scholten and Read (2006a), Kinari et al. (2009), Zauberman et al. (2009), and McAlvanah (2010).

¹⁹ For a survey of non-transitive preferences in decision theory, see Fishburn (1991).

²⁰ Cyclical choices can also be explained under EDU when stationarity is replaced by Thomsen separability; see Fishburn and Rubinstein (1982).




Figure 9.2 Sequential choices for the Ok and Masatlioglu (2007) example.

Source: Reprinted from *Journal of Economic Theory*, 137(1): 214–45. Ok, E.A. and Masatlioglu, Y., “A theory of (relative) discounting.” Copyright © 2007, with permission from Elsevier.

The first strict preference follows because \$150 does not compensate for a delay of two weeks, the second because \$50 compensates for a delay of one week, and the third follows because if \$50 compensates for a delay of one week, then \$100 should too. Ok and Masatlioglu (2007) point out that non-transitive choices are problematic because the decision maker is asked to commit to making a choice *today*. What if the decision maker does not have to commit to a choice today, and realizes that choices will be made sequentially in the future? Figure 9.2 depicts these sequential choices.

The decision maker can predict that if he waited for one week, then he would end up waiting another week and choosing \$250. Faced with a choice of \$250 from waiting and \$100 from not waiting, the decision maker chooses “not waiting” because an extra \$150 is not worth the delay of two weeks. Hence, there is a unique predictable time consistent choice. Ok and Masatlioglu (2007) argue that it might be improper to disallow non-transitive choices at the outset through the imposition of the *order* axiom, and their own axiomatic framework does not impose the order axiom.

It might well be that if one accounts for time, context, and reference dependent temporal preferences, intransitivities may disappear, or weaken substantially; see Gintis (2009). For a formalization of this argument, based on reference dependent time preferences, see al-Nowaihi and Dhami (2008a, 2013). Their solution allows one to preserve the axiom of order, yet account for non-transitive choice. These issues are further considered in Sections 10.5 and 10.5.6.

9.4.8 Preferences for increasing sequences

The EDU model ascribes no preference for particular shapes of consumption profiles (see the property of *utility independence* above) provided that the exponentially discounted values of alternative consumption streams are identical. However, empirical evidence generally shows that, holding fixed the discounted values of consumption streams, people prefer increasing over decreasing sequences. So for instance, in most human cultures, the dessert follows the main meal. Furthermore, subjects in experiments indicate a preference for increasing wage profiles to flat, or declining profiles with the same discounted value (Loewenstein and Sicherman, 1991). Indeed, even when the economic benefits of decreasing wage profiles are explained to the subjects

(e.g., opportunity cost of interest forgone) there was only a very marginal decrease in the number who preferred an increasing wage profile from 76 to 69.

In another study by Hsee et al. (1991), an increasing wage profile was preferred to a decreasing one, albeit one with a higher present value. In a related phenomenon that bears a similar interpretation, Varey and Kahneman (1992) found that individuals preferred a sequence characterized by decreasing discomfort rather than increasing discomfort.²¹ Evidence indicates that monkeys also prefer increasing sequences, and in particular, sequences that have a large reward at the end (the *end bias*); see Blanchard et al. (2014).

9.4.9 *Violations of consumption independence*

Consumption independence is one of the properties of the EDU model (see Section 9.3.1, above). Loewenstein and Prelec (1993) report the following experiment. Subjects were asked where they would like to eat (French restaurant, Home, or at a Lobster restaurant) over the next five weekends from among the following two pairs of menus, A, B , where each menu has two choices, $A = \{A1, A2\}$ and $B = \{B1, B2\}$; these are shown in Table 9.1.

The fifth week choices are the same in each pair of menus. Using the consumption independence feature of EDU, the fifth week choices should, therefore, have no influence in either set of menus. Thus, leaving out week 5 in each menu, if someone chooses $A2$ in menu A , it implies that they must choose $B2$ in menu B . A majority of the respondents (89%) choose $A2$, perhaps because the consumption profile is better spread temporally. However, respondents are split almost equally between $B2$ and $B1$. Hence, almost half the respondents behave in a manner that violates EDU.

Using field data from teachers in California, Mayer and Russell (2005) show that teachers prefer their salaries to be spread out equally over a 12 month period, as compared to receiving 10 equal installments of the annual salary over a period of 10 months. Under EDU, one expects teachers to prefer the latter option because they could earn more interest. One interpretation of these results is that the 12 equal, monthly spaced, salary receipts allow one to exert better control over spending, an explanation that is consistent with hyperbolic preferences.

9.4.10 *Effects of individual-specific traits on the discount rate*

Variations in discount rates are associated with a range of individual traits as well as situations, contexts, and emotions; see Chabris et al. (2007) for a more detailed survey. The elicited

Table 9.1 Choices among temporal sequences of meals.

Options/Weekends	1	2	3	4	5
A1 (11%)	French	Home	Home	Home	Home
A2 (89%)	Home	Home	French	Home	Home
B1 (49%)	French	Home	Home	Home	Lobster
B2 (51%)	Home	Home	French	Home	Lobster

Source: Loewenstein and Prelec (1993).

²¹ See also Chapman (2000) for a similar result. For contrary results of a preference for constant or decreasing profiles, see Chapman (1996), Gigliotti and Sopher (1997), Guyse et al. (2002).

discount rates of individuals are negatively correlated with age, i.e., the older are more patient.²² Individuals with higher cognitive ability tend to be more patient.²³ The following categories of individuals are more impatient (i.e., have higher discount rates): gamblers,²⁴ smokers,²⁵ heavy alcohol drinkers,²⁶ and drug users.²⁷

²² See Green et al. (1994), Green et al. (1996), Green et al. (1999), Read and Read (2004).

²³ See for instance, Silva and Gross (2004), Frederick (2005), Kirby et al. (2005), and Benjamin et al. (2006).

²⁴ See Petry and Casarella (1999), Petry (2001), Dixon et al. (2003), Alessi and Petry (2003).

²⁵ See Bickel et al. (1999), Mitchell (1999), Baker et al. (2003), Kirby and Petry (2004), Reynolds et al. (2004).

²⁶ See Vuchinich and Simpson (1998), Bjork et al. (2004).

²⁷ See Petry (2003), Coffey et al. (2003), Bretteville-Jensen (1999), and Kirby and Petry (2004).

CHAPTER 10

Behavioral Models of Time Discounting

10.1 Introduction

In this chapter, we turn to an explanation of the anomalies of the EDU model that we considered in Chapter 9. This naturally leads us to consider developments in behavioral models of time discounting.

Section 10.2 considers explanations of the *magnitude effect* and the *sign effect* (or *gain-loss asymmetry*). The magnitude effect arises because larger magnitudes are discounted less. The sign effect arises because losses are discounted less than gains. One may wish to explain these anomalies of EDU by putting the burden of adjustment on the discount function, or the utility function (or both). The explanation in Section 10.2.1 puts all the burden of adjustment on the utility function (Loewenstein and Prelec, 1992). Loewenstein–Prelec provided the first coherent explanation of these two effects. The magnitude effect is explained by assuming that the utility function is subproportional, and less elastic for outcomes of larger absolute magnitude. The sign effect is explained by adopting a utility function with greater elasticity for losses as compared to gains.

Loewenstein–Prelec’s explanation of the magnitude effect relies on an increasing elasticity of the utility function. However, several popular classes of utility functions violate this condition. These include CARA (constant absolute risk aversion), CRRA (constant relative risk aversion), HARA (hyperbolic absolute risk aversion), logarithmic, and quadratic. In Section 10.2.2, we give a class of utility functions that satisfy Loewenstein–Prelec’s conditions (al-Nowaihi and Dhami, 2008a, 2009). This class of utility functions is known as the *simple increasing elasticity*, SIE, class of utility functions.

A great deal of effort in behavioral economics has gone into explaining the *common difference effect*. The stationarity axiom in EDU implies that the discount rate over a time interval of fixed length is always constant, irrespective of the location of the interval. The common difference effect is a violation of the stationarity axiom. Section 10.3 considers the main explanations of the common difference effect.

Hyperbolic discounting, which we consider in Section 10.3.1, is one of the leading candidates for explaining the common difference effect. Hyperbolic discounting is the finding of hyperbolically declining discount rates. Since the discount rates are higher over shorter horizons, individuals appear to be particularly impatient over choices made over short horizons as compared to longer horizons. This gives rise to the well-known phenomenon of *present-biased preferences*; Chapter 11 is devoted to applications of this finding.

We consider two variants of hyperbolic discounting. The first is the *general hyperbolic discounting function* (Loewenstein and Prelec, 1992); our exposition of the general form of hyperbolic discounting relies on al-Nowaihi and Dhami (2006). The second variant is the *quasi-hyperbolic* or (β, δ) form (Phelps and Pollak, 1968; Laibson, 1997).

In the standard models of time discounting, the discount function is additive in the sense that splitting an interval into subintervals does not alter the extent of discounting. A violation of this feature can either lead to *subadditivity* (discounting over the bigger interval is lower), or its reverse, *superadditivity*. In Section 10.3.2, we consider an explanation for the common difference effect in terms of *subadditive discounting* (Read, 2001; Scholten and Read, 2006a). We also consider the *Read–Scholten discount function*, which is one of the most flexible discount functions available, and likely to be of great use in models of time discounting.

For whatever reasons, evolutionary, or otherwise, subjects in experiments on time discounting may view the future receipt of rewards as uncertain. One method of incorporating such uncertainty into the model is to combine expected utility theory with the EDU model (Halevy, 2008). The resulting model predicts that no common difference effect should arise. Thus, it would appear that the common difference effect is a joint rejection of EDU and expected utility theory. We consider these issues in Section 10.3.3. We also show that combining EDU with a non-expected utility model, prospect theory, predicts the common difference effect when there is “near certainty” (al-Nowaihi and Dhami, 2008a). However, the common difference effect reappears when we move away from near certainty situations. Hence, it is the *certainty effect* that seems to be causing the common difference effect, not non-expected utility per se. We also consider the work of Saito (2011) who generalizes the results in Halevy (2008) and Epper et al. (2011).¹

The outcome–time pairs which are compared in models of time discounting, consist of two natural *attributes*: time of delivery (time attribute), and the monetary value (outcome attribute). *Attribute models* attempt to explain temporal choices based on a trade-off between these two attributes. Section 10.4 considers three attribute-based models. In Section 10.4.1, we consider the *model of vague time preferences* (Manzini and Mariotti, 2006) that can explain the common difference effect. In this model, individuals use a hierarchical set of rules to make choices between outcome–time pairs. If the present discounted value of one pair is higher than a threshold, it is chosen, otherwise, a secondary criterion that compares the size of outcomes, or an earlier date is used to make choices. However, the choices are sensitive to the order in which the secondary criterion is exercised, i.e., time first, or outcome first.

Section 10.4.2 considers *similarity based relations* in choosing among outcome–time pairs (Rubinstein, 2003). When presented with a choice among alternatives, decision makers look for similarity along the various attribute dimensions. For instance, the outcome attribute, or the time attribute, in the two choices might be similar. Attributes where similarity is spotted are ignored. The choice is then made on the basis of the remaining attribute. Section 10.4.3 considers the *trade-off model of intertemporal choice* (Scholten and Read, 2006b). It can be shown that this attribute-based model belongs to the class of discounted utility models. i.e., the class of time preference models which are separable in time and outcomes (al-Nowaihi and Dhami, 2008a).

Section 10.5 considers a general theory of time discounting, the *reference time theory*, RT (al-Nowaihi and Dhami, 2008a, 2013, 2014). RT can explain all the main anomalies of the EDU model, namely, gain–loss asymmetry, magnitude effect, common difference effect, delay–speedup

¹ Part 1 of the book provides a separate and non-overlapping treatment of the interaction of risk preferences and time preferences.

asymmetry, apparent intransitivity of time preferences, and non-additivity of time discounting. It can also provide a framework for incorporating recent developments in time discounting such as uncertainty and attribute models. Essentially, RT combines generalized hyperbolic discounting with subadditive discounting. Furthermore, borrowing insights from prospect theory, it incorporates a reference point for time as well as a reference point for wealth. Preferences in RT are complete and transitive, yet it is able to explain apparent intransitivity of preferences. This is accomplished by a shift in the reference point for time. In the absence of such shifts in the reference time, what might just be an apparent intransitivity in preferences could be misconstrued as real intransitivity, as we show by an example.

10.2 Explaining the sign and magnitude effects

Loewenstein and Prelec (1992) (LP) provide the leading explanation of two anomalies of the EDU model: the *magnitude effect* and *gain-loss asymmetry*. However, none of the standard utility functions satisfy their conditions. Following al-Nowaihi and Dhami (2006, 2008b), we outline their work, and then provide a scheme for constructing the appropriate class of utility functions that satisfy the relevant conditions.

10.2.1 The Loewenstein-Prelec explanation

Consider a decision maker who, at time t_0 , formulates a plan to choose c_i at time $t_i \in [0, \infty]$. Let $v(c_t)$ be the instantaneous utility from consumption at time t . It is identical to the utility function under prospect theory as in Kahneman and Tversky (1979); see Chapter 2 for the details. Thus, among other properties such as loss aversion, v satisfies the following properties.

P1. $v : (-\infty, \infty) \rightarrow (-\infty, \infty)$ is continuous, and strictly increasing.

P2. $v(0) = 0$, and is differentiable except at 0, which is the reference point.

The *elasticity* of v is given by

$$\epsilon_v(c) = \frac{c}{v} \frac{dv}{dc}, c \neq 0. \quad (10.1)$$

Consider the following four assumptions, all with good experimental basis that are followed by an explanation.²

A0: Impatience: $D : [0, \infty) \rightarrow (0, 1]$ is a continuous, strictly decreasing function with $D(0) = 1$ and $D(t) \rightarrow 0$ as $t \rightarrow \infty$.³

A1: Gain-loss asymmetry: If $0 < z < z'$ and $v(z) = D(t)v(z')$, then $v(-z) > D(t)v(-z')$.

A2: Magnitude effect: If $0 < z < z'$, $v(z) = D(t)v(z')$ and $a > 1$, then $v(az) < D(t)v(az')$.

A3: Common difference effect: Let $0 < z < z'$, $v(z) = D(t)v(z')$. The common difference effect arises if there exists some $t' > 0$ such that $D(t')v(z) < D(t' + t)v(z')$.

Starting from the condition $v(z) = D(t)v(z')$, gain-loss asymmetry requires that losses are relatively more salient; thus, when gains turn into losses, the future loss bites relatively more.

² We make no use of the fifth assumption in LP: *delay-speedup asymmetry* (LP, p. 578).

³ We shall be mainly interested in discount functions that exhibit decreasing impatience. Bleichrodt et al. (2009) propose two classes of discount functions that allow for decreasing as well as increasing impatience.

Since, $v(-z) < 0$ and $v(-z') < 0$ (recall that v is a prospect theory utility function), A1 follows immediately. Similarly, the magnitude effect (A2) requires that starting from $v(z) = D(t) v(z')$, as z and z' are scaled up simultaneously, the larger future magnitude becomes more salient, hence, $v(az) < D(t) v(az')$. The common difference effect (A3) is a violation of the stationarity assumption of EDU. The decision maker flips his preferences as the outcomes are pushed further out into the future, while maintaining the time difference between the outcomes.

We now examine the implications of A0–A3 for the shape of the utility function. LP explain the sign and magnitude effects by putting the burden of adjustment entirely on the utility function, v (rather than on the discount function, D , or on both); the exact modeling strategy is an empirical question. Thus, we would like to find a utility function that satisfies assumptions A0–A3.⁴

Proposition 10.2 shows that the gain–loss asymmetry can be explained by using a utility function that is more elastic for gains as compared to losses.

Proposition 10.2 *Assumption A1 implies that the utility function is more elastic for losses than for gains: $z > 0 \Rightarrow \epsilon_v(-z) > \epsilon_v(z)$.*

Proof: By assumption, $0 < z < z'$ and $v(z) = D(t) v(z')$, hence,

$$\frac{v(z)}{v(z')} = D(t). \quad (10.2)$$

Using assumption A1, $v(-z) > D(t) v(-z')$. Since $-z' < 0$, and v is a prospect theory utility function, it follows that $v(-z') < 0$ and $v(-z) < 0$, hence,

$$\frac{v(-z)}{v(-z')} < D(t). \quad (10.3)$$

From (10.2), (10.3) we get that $\frac{v(-z)}{v(-z')} < \frac{v(z)}{v(z')}$. Since $v(-z) < 0$ and $v(-z') < 0$, we can rewrite the previous inequality as $\frac{-v(-z)}{-v(-z')} < \frac{v(z)}{v(z')}$. Taking logs on both sides we get that

$$\frac{\ln v(z') - \ln v(z)}{\ln z' - \ln z} < \frac{\ln(-v(-z')) - \ln(-v(-z))}{\ln z' - \ln z}. \quad (10.4)$$

Define $\tilde{z} = \ln z$ and $\tilde{z}' = \ln z'$. Letting v_+, v_- be, respectively, the utility function from gains and losses, define $\tilde{v}_+ = \ln v_+$ and $\tilde{v}_- = \ln v_-$. From (10.4), $\ln v$ is a concave function of $\ln z$, thus,

$$\frac{\tilde{v}_+(\tilde{z}') - \tilde{v}_+(\tilde{z})}{\tilde{z}' - \tilde{z}} < \frac{\tilde{v}_-(\tilde{z}') - \tilde{v}_-(\tilde{z})}{\tilde{z}' - \tilde{z}}.$$

Take limits as $\tilde{z}' \rightarrow \tilde{z}$, to get $\frac{d\tilde{v}_+}{d\tilde{z}} \leq \frac{d\tilde{v}_-}{d\tilde{z}}$, from which it follows that $\epsilon_v(z) \leq \epsilon_v(-z)$. ■

We now show that a necessary condition to explain the magnitude effect is that the elasticity of the utility function must increase as the absolute magnitude of outcomes increases.

⁴ We shall note the problems that such an explanation leads to, below, and the alternative explanation proposed by al-Nowaihi and Dhami (2008a).

Proposition 10.3 (Loewenstein and Prelec, 1992; al-Nowaihi and Dhami, 2008a): *Assumption A2 implies that the utility function is more elastic for outcomes of larger absolute magnitude, i.e.,*

$$(0 < z < z' \text{ or } z' < z < 0) \rightarrow \epsilon_v(z) < \epsilon_v(z').$$

10.2.2 The class of simple increasing elasticity utility functions

Loewenstein and Prelec (1992) do not give any example of a utility function that has the properties required in Propositions 10.2 and 10.3 to explain the sign and magnitude effects. al-Nowaihi and Dhami (2009) show that none of the standard utility functions, CARA, CRRA, HARA, logarithmic, and quadratic satisfy the properties in Propositions 10.2 and 10.3 (an exercise asks you to show this).

In order to resolve this problem, al-Nowaihi and Dhami (2009) provide a simple, tractable class of utility functions, the *simple increasing elasticity* (SIE) utility functions, that satisfies the properties required in Propositions 10.2 and 10.3. Furthermore, they provide a scheme for generating such utility functions. The following method can be used to generate candidates for the, as yet undetermined, SIE utility functions.

For some set of outcomes X , choose a function, h , satisfying,

$$0 < h(x) < 1, h'(x) > 0, x \in X. \quad (10.5)$$

Then solve the following differential equation for $u(x)$:

$$\frac{x}{u} \frac{du}{dx} = h(x). \quad (10.6)$$

This method only yields candidate utility functions, which then have to be verified. For example, choose

$$h(x) = \frac{ax}{b+x} + c; x \geq 0, a, b, c > 0, a + c \leq 1. \quad (10.7)$$

Substituting from (10.7) into (10.6), we get

$$\frac{x}{u} \frac{du}{dx} = \frac{ax}{b+x} + c.$$

Separating variables, then integrating, gives

$$\int \frac{du}{u} = a \int \frac{dx}{b+x} + c \int \frac{dx}{x}.$$

The solution is

$$\begin{aligned} \ln u &= a \ln(b+x) + c \ln x + \ln K. \\ \Rightarrow u &= e^{\ln(b+x)^a} e^{\ln x^c} e^{\ln K}. \end{aligned}$$

Since the logarithm and the exponential functions are inverses of each other, we can write this expression alternatively as,

$$u(x) = K(b+x)^a x^c. \quad (10.8)$$

Choosing $a = 1 - \gamma$, $b = \frac{\gamma\mu}{\theta}$, $c = \sigma$ and $K = \frac{\gamma}{1-\gamma} \left(\frac{\theta}{\gamma}\right)^{1-\gamma}$ produces a utility function:

$$u(x) = \frac{\gamma x^\sigma}{1-\gamma} \left(\mu + \frac{\theta}{\gamma} x\right)^{1-\gamma}, \quad x \geq 0. \quad (10.9)$$

The restrictions $a, b, c > 0$, $a + c \leq 1$ give: $0 < \sigma \leq \gamma < 1$ and $\mu/\theta > 0$. To ensure that $u' > 0$, take $\theta > 0$. Hence, $\mu > 0$. For $x < 0$, define $u(x)$ by $u(x) = -\lambda u(-x)$, where $\lambda > 1$.⁵ Putting all these together gives the candidate SIE utility function

$$u(x) = \begin{cases} \frac{\gamma x^\sigma}{1-\gamma} \left(\mu + \frac{\theta_+}{\gamma} x\right)^{1-\gamma} & \text{if } x \geq 0 \\ -\lambda \frac{\gamma (-x)^\sigma}{1-\gamma} \left(\mu - \frac{\theta_-}{\gamma} x\right)^{1-\gamma} & \text{if } x < 0 \end{cases}, \quad (10.10)$$

where $\mu > 0$, $\theta_\pm > 0$, $\lambda > 1$, $0 < \sigma \leq \gamma < 1$, and λ is the coefficient of loss aversion.

It may be interesting to note that (10.10) is a product of a CRRA function, x^σ , and a HARA function, $\frac{\gamma}{1-\gamma} \left(\mu + \frac{\theta_\pm}{\gamma} x\right)^{1-\gamma}$. From (10.10), it is straightforward to prove the next proposition

Proposition 10.4 (*al-Nowaihi and Dhami, 2009*). *For the utility function u defined in (10.10):*

- (a) $u : (-\infty, \infty) \rightarrow (-\infty, \infty)$, $u(0) = 0$, u is continuous, u is C^∞ except at $x = 0$.
- (b) The elasticity of the utility function, $\epsilon_u(x) = \frac{x}{u} \frac{du}{dx}$, is given by

$$\epsilon_u(x) = \begin{cases} \sigma + \frac{1-\gamma}{1 + \frac{\gamma\mu}{x\theta_+}} > 0 & \text{if } x > 0 \\ \sigma + \frac{1-\gamma}{1 - \frac{\gamma\mu}{x\theta_-}} > 0 & \text{if } x < 0 \end{cases}. \quad (10.11)$$

- (c) $\epsilon'_u(x) > 0$ for $x > 0$ and $\epsilon'_u(x) < 0$ for $x < 0$.
- (d) $x > 0 \Rightarrow \epsilon_u(x) < \epsilon_u(-x)$.

Comparing Proposition 10.2 with Proposition 10.4(d), the SIE utility function explains the sign effect, and comparing Proposition 10.3 with Proposition 10.4(c), the SIE utility function can also explain the magnitude effect. The following proposition gives some useful properties of $\epsilon_u(x)$. Some other properties of the SIE utility function are contained in an exercise at the end of this part of the book.

Proposition 10.5 (*al-Nowaihi and Dhami, 2009*) *From (10.11), we get that $\epsilon_u(x) \rightarrow \sigma$ as $x \downarrow 0$ and as $x \uparrow 0$. Hence, $\epsilon_u(x)$ is defined for all $x \in (-\infty, \infty)$, $\epsilon_u(x) = \sigma + \frac{1-\gamma}{1 + \frac{\gamma\mu}{x\theta_+}} > 0$, $\epsilon_u(x)$ is increasing in $|x|$ and $\epsilon_u(x) \rightarrow \sigma + 1 - \gamma \leq 1$, as $|x| \rightarrow \infty$.*

The SIE class of utility functions is compatible with any theory of time preferences where preferences are separable in time and outcomes. Among others, this includes the following models that we consider below: the general hyperbolic discounting model, the (β, δ) class of quasi-hyperbolic discounting models, the subadditive discounting models, models where uncertainty is

⁵ This imposes loss aversion in the domain of losses in the same manner as under prospect theory; see Part 1 of the book.

incorporated, attribute-based models, and the reference time theory. Hence, the SIE class should be very useful.

Like LP, we have explained the sign and magnitude effects through the utility function. We now examine the resulting implications for the parameter of loss aversion, through an example. To give a feel for the issues, use the following parameter values for the SIE class of utility functions: $\mu = 1, \theta = \frac{1}{2}, \lambda = 2, 0 < \sigma = \gamma = \frac{1}{2}$. These satisfy the restrictions on the parameters in (10.10). With these parameter values, the SIE utility function (10.10) takes the form:

$$u(x) = \begin{cases} x^{\frac{1}{2}} (1+x)^{\frac{1}{2}} & \text{if } x \geq 0 \\ -2u(-x) = -2(-x)^{\frac{1}{2}} (1-x)^{\frac{1}{2}} & \text{if } x < 0 \end{cases}. \quad (10.12)$$

The *generalized hyperbolic discount function* (see (10.16) below), is

$$D(t) = (1 + \alpha t)^{-\sigma/\alpha}, \alpha, \sigma > 0. \quad (10.13)$$

Let $\alpha = 1, \sigma = 2$, then (10.13) takes the form

$$D(t) = (1 + t)^{-2}. \quad (10.14)$$

Let $V_0(x, t)$ be the value of x received at time t , discounted back to time 0.

$$V_0(x, t) = D(t) u(x).$$

Then (10.12), (10.14) give:

$$V_0(x, t) = \begin{cases} (1+t)^{-2} \left(x^{\frac{1}{2}} (1+x)^{\frac{1}{2}} \right) & \text{if } x \geq 0 \\ (1+t)^{-2} \left(-2(1+t) (-x)^{\frac{1}{2}} (1-x)^{\frac{1}{2}} \right) & \text{if } x < 0 \end{cases}. \quad (10.15)$$

Comparing the two rows in (10.15), we see that gains and losses now have the same discount function, $D(0, t) = (1 + t)^{-2}$. But the parameter of loss aversion, $\lambda = 2(1 + t)$, depends on the time, t , when x is received. In particular, $\lambda \rightarrow \infty$ as $t \rightarrow \infty$. We know of no robust empirical evidence that supports this. al-Nowaihi and Dhami (2008a) propose an alternative explanation of the sign effect. Unlike LP, they do not put the burden of adjustment on the utility function. Rather, they posit different discount functions for gains and losses, which enables them to explain the sign effect. In particular, they keep the parameter of loss aversion fixed, which is consistent with the experimental evidence.

10.3 Explaining the common difference effect

The *common difference effect*, which is a violation of *stationarity* has already been described above. Stationarity requires that preferences over outcome–time pairs do not alter as we adjust all time intervals equally. Example 9.1 nicely illustrates a violation of stationarity, which gives rise to the common difference effect. This section examines alternative explanations for the common difference effect.

10.3.1 Hyperbolic discounting

Hyperbolic discounting is one of the leading candidates for explaining the common difference effect. In this section, we consider the *general form of hyperbolic discounting*, due to Loewenstein and Prelec (1992), and the *quasi-hyperbolic form*, also known as the (β, δ) form, due to Phelps and Pollak (1968) and later popularized by Laibson (1997) who also drew out a much richer set of implications.

THE GENERAL MODEL OF HYPERBOLIC DISCOUNTING

The discount function in the general hyperbolic discounting function is given by⁶

$$D(t) = (1 + \alpha t)^{-\sigma/\alpha}; \alpha, \sigma > 0. \quad (10.16)$$

It follows that

$$\begin{aligned} \lim_{\alpha \rightarrow 0} D(t) &= \lim_{\alpha \rightarrow 0} (1 + \alpha t)^{-\sigma/\alpha} \\ &= \lim_{m \rightarrow \infty} \left(1 + \frac{t}{m}\right)^{-\sigma m} \quad (\text{by defining } m = \alpha^{-1}) \\ &= e^{-\sigma t} \quad (\text{by using the definition of } e). \end{aligned}$$

Thus, when $\alpha \rightarrow 0$, the generalized hyperbolic function collapses to the exponential discounting function. For very large values of α , this discount function approximates a step function.

We now give the axiomatic derivation in al-Nowaihi and Dhami (2006) for the generalized hyperbolic discounting function. Consider the following axioms, the first of which, impatience, has already been introduced above as A0:

B0: Impatience: $D : [0, \infty) \rightarrow (0, 1]$ is a continuous and strictly decreasing function with $D(0) = 1$ and $D(t) \rightarrow 0$ as $t \rightarrow \infty$.

B1: Common difference effect with quadratic time delay: If $0 < z < z'$, $v(z) = D(t) v(z')$ and $t' > 0$, then $D(t') v(z) = D(t' + t + \alpha t' t) v(z')$, $\alpha > 0$.

Notice that in EDU, $\alpha = 0$, hence, EDU violates B1. Under quadratic delay, relative to EDU, the value of the current reward increases disproportionately more as we approach the present, starting from some point in the future.

Proposition 10.6 (al-Nowaihi and Dhami, 2006): *B0 and B1 imply that the discount factor is a generalized hyperbola: $D(t) = (1 + \alpha t)^{-\frac{\sigma}{\alpha}}$, $\sigma > 0$, $t \geq 0$, where α is given in B1.*

Proof: Let

$$0 < z < z'. \quad (10.17)$$

Assume that an individual is indifferent between receiving the amount z now and receiving the amount z' , t periods from now. Then, letting v be the utility function and D the discount factor, we get

⁶ See also Harvey (1986) for an earlier development.

$$v(z) = D(t) v(z'). \quad (10.18)$$

Multiply (10.18) by $D(t')$, where $t' > 0$, to get

$$D(t') v(z) = D(t') D(t) v(z'). \quad (10.19)$$

Axiom B1, (10.17) and (10.18) give

$$D(t') v(z) = D(t' + t + \alpha t' t) v(z'), \alpha > 0. \quad (10.20)$$

Equations (10.19) and (10.20) give

$$D(t' + t + \alpha t' t) = D(t') D(t). \quad (10.21)$$

Let

$$X = 1 + \alpha t', Y = 1 + \alpha t, \quad (10.22)$$

hence,

$$t' = \frac{X-1}{\alpha}, \quad t = \frac{Y-1}{\alpha}, \quad t' + t + \alpha t' t = \frac{XY-1}{\alpha}. \quad (10.23)$$

Define the function $G : [1, \infty) \rightarrow (0, \infty)$ by

$$G(X) = D\left(\frac{X-1}{\alpha}\right). \quad (10.24)$$

Hence,

$$G(Y) = D\left(\frac{Y-1}{\alpha}\right), \quad G(XY) = D\left(\frac{XY-1}{\alpha}\right). \quad (10.25)$$

From (10.21), (10.23), (10.24), and (10.25)

$$G(XY) = D\left(\frac{XY-1}{\alpha}\right) = D(t' + t + \alpha t' t) = D(t') D(t) \quad (10.26)$$

$$= D\left(\frac{X-1}{\alpha}\right) D\left(\frac{Y-1}{\alpha}\right) = G(X) G(Y). \quad (10.27)$$

Define the function $h : [0, \infty) \rightarrow (0, \infty)$ by

$$h(z) = G(e^z), \quad z \geq 0. \quad (10.28)$$

Hence, and in the light of axiom B0, and (10.24), (10.27), $h : [0, \infty) \rightarrow (0, \infty)$ is strictly decreasing (in an arbitrarily small interval), and

$$h(z + z') = h(z) h(z'). \quad (10.29)$$

It is well known, and readily checked,⁷ that the solution of the functional equation (10.29) is the exponential function.

$$h(z) = e^{cz}, \quad z \geq 0, \quad c < 0. \quad (10.30)$$

From (10.28), (10.30) we get that

$$G(e^z) = e^{cz} = (e^z)^c \Leftrightarrow G(Y) = Y^c. \quad (10.31)$$

From (10.23), (10.25) we get

$$G(Y) = D\left(\frac{Y-1}{\alpha}\right) = D(t). \quad (10.32)$$

From (10.22), $Y = 1 + \alpha t$. Using this, in conjunction with (10.31), (10.32) we get that

$$D(t) = (1 + \alpha t)^c. \quad (10.33)$$

Let

$$\sigma = -\alpha c. \quad (10.34)$$

Equations (10.33), (10.34) give

$$D(t) = (1 + \alpha t)^{-\frac{\sigma}{\alpha}}, \quad \alpha, \sigma > 0, \quad t \geq 0, \quad (10.35)$$

where $\sigma > 0$ because $\alpha > 0$ and $c < 0$. Thus, under the axioms B0, B1, the discount function takes the general hyperbolic form, as claimed. ■

Recall the definition of the discount rate $\theta(t)$ in (9.4). Corollary 10.1 below derives the implications of hyperbolic discounting for the discount rate.

Corollary 10.1 *B0 and B1 imply that $\theta(t) = -\frac{\dot{D}}{D} = \frac{\sigma}{1+\alpha t}$. Hence, the discount rate, $\theta(t)$, is positive and declining in t .*

Corollary 10.1 establishes the result on declining discount rates, which is consistent with the evidence in Section 9.4.1 above.

THE QUASI-HYPERBOLIC MODEL

A simpler and more tractable method for modeling declining discount rates is the *quasi-hyperbolic form*, which is also popularly known as the (β, δ) form of hyperbolic discounting. Under quasi-hyperbolic discounting, a dated stream of consumption $(c_0, c_1, \dots, c_T)$ is evaluated in the following manner.

⁷ See, for example, Corollary 1.4.11 in Eichhorn (1978).

$$U(c_0, c_1, \dots, c_T) = u(c_0) + \beta \sum_{t=1}^T \delta^t u(c_t), 0 < \beta < 1. \quad (10.36)$$

The functional form in (10.36) was originally introduced by Phelps and Pollak (1968). Laibson (1994, 1997) is to be largely credited for reviving interest in it and developing several new implications. In this formulation, δ plays the same role that it plays in the EDU model. The innovation is the presence of the term $\beta \in (0, 1)$ that shrinks the value of future lifetime utility relative to current utility, thereby creating an additional *present bias for current consumption*. This feature leads to the time inconsistency problem. When making a choice between two dates $t > 0$ and $t' > t$, the parameter β washes out, and choices are made as under EDU. However, when these choices are brought forward by t periods, i.e., between two dates 0 and $t' - t$, then the present bias for consumption kicks in through the parameter β , potentially reversing the earlier choice. Hence, this gives rise to a common difference effect.

The discount factor between time periods 0 and 1 is $\beta\delta = \frac{1}{1+\theta_0}$, where θ_0 , the effective one-period discount rate between times 0 and 1, is

$$\theta_0 = \frac{1 - \beta\delta}{\beta\delta}. \quad (10.37)$$

By contrast, the discount factor between time periods $\tau > 0$ and $\tau + 1$ is $\delta = \frac{1}{1+\theta_\tau}$, so the effective discount rate, θ_τ , between these two periods is

$$\theta_\tau = \frac{1 - \delta}{\delta}, 0 < \tau \leq T. \quad (10.38)$$

From (10.37), (10.38), and using $0 < \beta < 1$, it follows that $\theta_\tau < \theta_0$. It also follows that the discount rate is a step function. The discount rate that applies to current (time 0) choices, θ_0 , is high, but for subsequent choices, it permanently steps down to a lower value, θ_τ , $\tau = 1, 2, 3, \dots$

The main difference between the general form of hyperbolic discounting and the (β, δ) form is the following. As we approach the present, having started off at some time in the future, the general hyperbolic form exhibits continuously increasing impatience, while in the (β, δ) form, there is a sudden and discontinuous increase in impatience at the initial time. We plot in Figure 10.1, the sequence of discount factors for the exponential discount function $(\delta, \delta^2, \dots)$, the quasi-hyperbolic discount function $(1, \beta\delta, \beta\delta^2, \dots)$, and the generalized hyperbolic discount function $(1, (1 + \alpha)^{-\sigma/\alpha}, (1 + 2\alpha)^{-\sigma/\alpha}, \dots)$.⁸ These discount factors are shown, respectively, by the continuous bold curve, the lighter step function, and the dashed curve. We use the following calibration values for the discount factors suggested by Laibson (1994): $\delta = 0.95$ for the exponential, $(\beta, \delta) = (0.5, 0.975)$ for the quasi-hyperbolic, and $(\alpha, \sigma) = (3, 1)$ for the generalized hyperbolic. Of these, only the quasi-hyperbolic discount factor is discontinuous, and plots as a step function. All three discount factors corresponding to the three different discount functions are decreasing functions. However, the two discount factors corresponding to the hyperbolic functions show relatively much steeper initial drops, and then flatten out. Comparing the exponential and the general hyperbolic functions, the latter exhibits a much steeper increase in impatience as one approaches the present time period.

⁸ Figure 10.1 plots the discount factors. It is not to be confused with Figure 9.1, which plots the discount rate for different time delays in rewards.




Figure 10.1 The discount factors for various discount functions: Exponential (continuous bold curve), quasi-hyperbolic (lighter step function), and generalized hyperbolic (dashed curve).

Example 10.2 (*Explanation of Thaler's example*): Let us use quasi-hyperbolic discounting to explain Thaler's (1981) apples example (Example 9.1, above). Let $u(x)$ be the utility from outcome x . Under quasi-hyperbolic discounting, the preference for one apple today to two apples tomorrow implies that

$$u(1) > \beta \delta u(2). \quad (10.39)$$

The preference for two apples in 51 days to one apple in 50 days implies that

$$u(1) < \delta u(2). \quad (10.40)$$

(10.39) and (10.40) imply that

$$\beta \delta u(2) < u(1) < \delta u(2). \quad (10.41)$$

(10.41) can be consistent with hyperbolic discounting. For instance, the values for the parameters used in Figure 10.1 above, $(\beta, \delta) = (0.5, 0.975)$, and an identity utility function, $u(x) = x$, satisfy (10.41).

EMPIRICAL EVIDENCE ON HYPERBOLIC DISCOUNTING

We shall only briefly survey the extensive evidence on hyperbolic discounting.⁹ Cropper et al. (1992) considered the trade-off between lives saved today and in the future, from implementing environmental cleanup programs. The cleanup of hazardous waste is costly, and while it may save some lives today, there is a greater likelihood of saving a larger number of lives in the future. They

⁹ The interested reader can consult the extensive survey in Frederick et al. (2002). A very large number of studies support hyperbolic discounting. See, for instance, Ainslie (1975), Thaler (1981), Loewenstein and Prelec (1992), Green and Myerson (1996), Kirby (1997), and Kim and Zauberman (2009).

put subjects in experiments in the role of social planners, and elicited their discount rates from the choices among hypothetical alternatives.

A total of 3,200 subjects were given choices of the following form. Without new environmental programs, 100 people will die this year from pollution, and 200 people will die $\tau > 0$ years from now. Subjects have to choose between two programs that have identical costs. Program A will save 100 lives now. Program B will save y lives τ years from now. Time τ was varied in units of years, $\tau = 5, 10, 25, 50, 100$. The number of lives saved in the future, y , was also varied, but typically $y = 200$. The main finding is that people exhibit higher discount rates for short horizons, and lower discount rates for longer horizons. There is also heterogeneity in discount rates among individuals, but no relation was found between elicited discount rates and individual characteristics, such as age, sex, marital status, education, or income.

The authors then fitted two piecewise linear discount rate curves. The first was for data pooled over $\tau = 5, 10$, and the second was for data pooled over $\tau = 25, 50$. The latter curve was significantly flatter. Henderson and Bateman (1995) fitted the data in Cropper et al. (1992) to the more general hyperbolic curve in (9.29). By normalizing $D(0) = 1$, they find that a good fit to the data is given by the hyperbolic discount function $D(\tau) = \frac{1}{1+(0.21)\tau}$.

Viscusi and Huber (2006) use survey evidence and find a declining trend in the discount rate for public goods as the time delay increases, which is consistent with hyperbolic discounting. Ahumada and Gregnani (2007) use quarterly consumer expenditure data from Argentina and test for exponential as well as hyperbolic discounting.¹⁰ They find evidence that consumers procrastinate, as predicted by the hyperbolic discounting model; see Section 11.5 below for the effects of hyperbolic discounting on procrastination. For the time period 2002–2005, a hyperbolic discounting function fits the data quite well. For an earlier period, however, exponential discounting appears more suitable.

Viscusi et al. (2008) give hypothetical choices to a large number of respondents. The respondents are asked to choose among alternative programs to improve the water quality in their regions. In each program, by paying a current cost, respondents can choose a level of improvement in water quality with a time delay. Programs differ in terms of the current cost, level of improvement, and the time delay before the improvement sets in. When the data is fitted to the quasi-hyperbolic model, the estimates of the parameter β range from 0.48 to 0.61.

Using a sample of 2987 Japanese adults, Ikeda et al. (2010) first measure the body mass index (BMI) of the respondents. Then they elicit their discount rates over different horizons by giving subjects a choice between smaller–sooner, and larger–later rewards. There are two main findings. A higher BMI leads to (1) greater impatience that takes the hyperbolic form, and (2) a lower magnitude of the gain–loss asymmetry or sign effect (subjects who exhibit the sign effect have a lower probability of being obese).

Fang and Silverman (2009) consider the 1979 youth cohort of the National Longitudinal Survey of Youth (NLSY). They were interested in the labor force participation and welfare program participation decision of never-married women. Using semi-parametric estimation methods for dynamic discrete choice models, they find evidence of quasi-hyperbolic discounting. In particular, they find that the fitted parameters of the quasi-hyperbolic function are $\beta = 0.338$ and $\delta = 0.88$.

Fang and Wang (2010) allow for the possibility that hyperbolic decision makers may be *naive*. A formal definition is given below in Section 11.5. A hyperbolic decision maker is *naive* if he

¹⁰ They use the generalized method of moments applied to a Euler equation derived in Harris and Laibson (2001a). We derive, below, the Euler equations in macroeconomic models with hyperbolic discounting.

believes that in the future, he has EDU preferences. Such a decision maker underestimates future self-control problems. They find evidence of naivety, and the parameter β in the quasi-hyperbolic model equals 0.72.

Salois and Moss (2011) use annual data on farmland valuations and returns in the US from a 43 year panel. Their estimation equation derives from the change in the value of farmlands between two successive periods. They find that exponential discounting is rejected for six out of their eight panel regions, while quasi-hyperbolic discounting provides a good fit to their data. The parameter β ranges from 0.304 to 0.637 over the eight regions.

Using a variety of time-delay questions, typically of the smaller-sooner versus larger-later choice variety, Benhabib et al. (2010) propose a fixed cost of time delay. Suppose that (z, t) denotes a reward z to be received at time t . Using an incentive compatible elicitation mechanism, they estimate the following discount function, D , that nests several possibilities.

$$D(z, t; \nu, r, \alpha, b) = \begin{cases} 1 & \text{if } t = 0 \\ \alpha [1 - (1 - \nu) rt]^{\frac{1}{1-\nu}} - \frac{b}{z} & \text{if } t > 0 \end{cases}.$$

A simple calculation shows that

$$D(z, t; 1, r, 1, 0) = e^{-rt}; D(z, t; 2, r, 1, 0) = \frac{1}{1 + rt}; D(z, t; 1, r, \alpha, 0) = \alpha e^{-rt}. \quad (10.42)$$

From (10.42), we get the following cases: (i) For $\nu = 1, \alpha = 1, b = 0$, we get exponential discounting. (ii) For $\nu = 2, \alpha = 1, b = 0$, we get hyperbolic discounting. (iii) For $\nu = 1, \alpha < 1, b = 0$, we get quasi-hyperbolic discounting. (iv) For the most general case with $b > 0$, we have a fixed cost of receiving a future reward. The larger the reward, z , the smaller is this cost. Not surprisingly, this formulation can account for the magnitude effect described above. Let us call this case *discounting with fixed cost*. But there is no presumption that this discount function can explain several other anomalies such as subadditivity, although the sign effect can be explained if one uses the SIE utility function in conjunction with this discount function. Obviously, it can also explain the common difference effect.

The main findings of Benhabib et al. (2010) are as follows. (1) Exponential discounting is rejected. (2) The discount rate is declining over time, which indicates the presence of a *present bias* in time preferences. (3) The most general formulation, discounting with fixed costs ($b > 0$), is preferred to hyperbolic discounting. The authors are careful to stress that (p. 223) “Experiments with relatively large rewards are needed to confirm the fixed cost representation of present bias that we identify in the experimental data.” By contrast, Andersen et al. (2011) show that their experimental results with 413 adult Danes are consistent with exponential discounting, and the elicited discount rates are quite low. They find no evidence of quasi-hyperbolic discounting, or of discounting with fixed costs.

Tanaka et al. (2010) estimate the Benhabib et al. (2010) discount function using high stakes experiments in Vietnam in a tightly controlled set of experiments. They find evidence for quasi-hyperbolic discounting. The estimated value of β is 0.644, which is consistent with the typical estimates of this parameter. Interestingly, the present bias is found irrespective of the level of incomes, or the economic circumstances of the subjects.

Kosse and Pfeiffer (2013) focus on the quasi-hyperbolic (β, δ) model. They examined the possibility of intergenerational transmission of patience between mother and child. They first elicit the β, δ parameters for 213 mothers, using choices between smaller-sooner and larger-later rewards

in an incentive compatible manner. Then they elicit the same discounting parameters for the children. Children were given a choice of receiving a pack of gummy bears at the start of the mother's interview or two packs at the end of the interview. A probit regression was then run in which the child's β parameter was the dependent parameter, and the mother's elicited β, δ parameters were the independent variables. A higher value of the mother's β parameter (recall that $\beta = 1$ gives the EDU model) was found to significantly, and positively, influence the probability that the child would choose the delayed reward (i.e., exhibit greater patience).

An emerging literature explores the relation between risk preferences and time preferences. We consider this issue in more detail in Section 10.3.3 below; see also Part 1 of the book for a consideration of these issues. The basic idea is that subjects in experiments may consider future rewards to be uncertain, even when no such uncertainty can be inferred from the instructions in the experiment. Hence, in making choices between, say, smaller-sooner and larger-later rewards, perceived uncertainty of the larger-later reward may distort the elicited time preferences. Thus, the claim is that time preferences can potentially be reversed when risk preferences are added to the problem and vice-versa (Keren and Roelofsma, 1995; Weber and Chapman, 2005; Baucells and Heukampf, 2010). Greater risk aversion may also be associated with greater impatience (Anderhub et al., 2001).

A range of studies allow for the simultaneous presence of uncertainty and time delay, and find significant interaction effects between the two.¹¹ Epper et al. (2011) allow for non-linear probability weighting under uncertainty, and consider the interaction effects between uncertainty and time delay. An incentive compatible mechanism is used to elicit choices. They find that their data is best described by non-linear probability weighting and hyperbolic time delay. Subjects who exhibit more pronounced departures from linear weighting of probabilities also display a greater departure from exponential discounting.

In summary, the weight of the evidence supports declining discount rates. There may be less agreement about whether the pattern of declining rates is hyperbolic, or not. The estimated β parameter in (β, δ) models is typically below 1 ($\beta = 1$ corresponds to EDU), suggesting that present bias is present in the data. More realistic models may need to interact hyperbolic discounting with risk preferences; this is already an active area of research that we report on extensively in this book.

10.3.2 Subadditive discounting

The explanation of further developments in time discounting requires a slightly more formal definition of a discount function. Suppose that we are interested in the discount function between time periods $r \geq 0$ and $t \geq r$, then the following definition of an *interval discount function* is suitable.

Definition 10.3 (*Interval discount functions*): Let

$$\Delta = \{(r, t) \in \mathbb{R} \times \mathbb{R} : 0 \leq r \leq t\}. \quad (10.43)$$

A discount function is a mapping, $D : \Delta \rightarrow (0, 1]$, satisfying:

- (i) For each $r \in [0, \infty)$, $D(r, t)$ is a strictly decreasing function of $t \in [r, \infty)$ into $(0, 1]$ with $D(r, r) = 1$.

¹¹ See, for instance, Abdellaoui et al. (2011), Coble and Lusk (2010), Baucells and Heukampf (2010), Anderson and Stafford (2009), Weber and Chapman (2005), and Keren and Roelofsma (1995).

- (ii) For each $t \in [0, \infty)$, $D(r, t)$ is a strictly increasing function of $r \in [0, t]$ into $(0, 1]$. Furthermore, if D satisfies (i) with “into” replaced with “onto,” then we call D a continuous discount function.

As one expands the interval $[r, t]$, the discount factor applied to the interval falls. This can happen in one of two ways: an increase in t (Definition 10.3(i)) and a decrease in r (Definition 10.3(ii)). A continuous discount function $D(r, t)$ is continuous in t ; see al-Nowaihi and Dhami (2008a) for a formal proof. We leave the proof of the next proposition as an exercise.

Proposition 10.7 (al-Nowaihi and Dhami, 2008a): Let D be a continuous discount function and let $v(x)$ be the utility from outcome x . Suppose $r \geq 0$. If $0 < x \leq y$, or if $y \leq x < 0$, then $v(x) = v(y) D(r, t)$ for some $t \in [r, \infty)$.

Read (2001) and Scholten and Read (2006a) found experimental evidence for *subadditivity*, which we now formally define in terms of the interval discount function.

Definition 10.4 (Additivity): A discount function, $D(r, t)$, is:

$$\begin{cases} \text{Additive if} & D(r, s) D(s, t) = D(r, t), & \text{for } r \leq s \leq t, \\ \text{Subadditive if} & D(r, s) D(s, t) < D(r, t), & \text{for } r < s < t, \\ \text{Superadditive if} & D(r, s) D(s, t) > D(r, t), & \text{for } r < s < t. \end{cases}$$

In Definition 10.4, additivity implies that discounting a quantity from time t back to time s , and then further back to time r , is the same as discounting that quantity from time t back to time r in one step. In other words, breaking an interval into subintervals has no effect on discounting. However, in the other two cases, it does have an effect. Under subadditive discounting, there is more discounting over the subdivided intervals (future utilities are shrunk more), while the converse is true under superadditive discounting.

Example 10.3 Consider the following example that illustrates the importance of subadditive discounting. Consider three dates, $0 \leq r < s < t$. A risk neutral investor faces an investment opportunity to invest \$ a at date r that materializes at date t and yields the amount \$ c . At date r , the investor is assumed to have no opportunity cost of his funds.

Suppose that there are two kinds of investors. Investor 1 discounts over the whole interval $[r, t]$, so he invests if $a < D(r, t)c$. By contrast, investor 2 uses subadditive discounting and first discounts the final amount from time t to s and then back from time s to r . Thus, investor 2 invests if $a < D(r, s) D(s, t) c$. If the following condition holds

$$D(r, s) D(s, t) c < a < D(r, t) c,$$

then investor 1 invests, but investor 2 does not invest.

Under exponential discounting, if we move a fixed interval of time, say, $[r, s]$, $r < s$, to another date, $[r + t, s + t]$, where $t \geq 0$, then discounting over that interval stays the same (constant impatience). However, two other possibilities may also arise. In making decisions about time intervals of a fixed size, that recede into the future, one may become more patient (*declining impatience*), as in models of hyperbolic discounting. The converse of this is *increasing impatience*. We formalize these concepts in the next definition.

Definition 10.5 A discount function, $D(r, s)$, exhibits:¹²

¹² Some authors use “present bias” for what we call “declining impatience.” But other authors use “present bias” to mean that the discount function, $D(s, t)$ is declining in t . So we prefer “declining impatience” to avoid confusion.

$$\begin{cases} \text{declining impatience if} & D(r, s) < D(r + t, s + t), & \text{for } t > 0 \text{ and } r < s, \\ \text{constant impatience if} & D(r, s) = D(r + t, s + t), & \text{for } t \geq 0 \text{ and } r \leq s, \\ \text{increasing impatience if} & D(r, s) > D(r + t, s + t), & \text{for } t > 0 \text{ and } r < s. \end{cases}$$

Based on the empirical evidence, Scholten and Read (2006a) developed a specific interval discount function, which we shall call the *RS-discount function*; this is stated below.

$$D(r, t) = [1 + \alpha(t^\tau - r^\tau)^\rho]^{-\frac{\beta}{\alpha}}, 0 \leq r \leq t, \alpha > 0, \beta > 0, \rho > 0, \tau > 0. \quad (10.44)$$

The RS discount function is extremely flexible, and depending on various parameter values, it can incorporate additivity, subadditivity, superadditivity, declining impatience, constant impatience, and increasing impatience; see al-Nowaihi and Dhami (2008a) for the relevant proofs. On account of its generality relative to the other discounting functions, we state some of its properties in the following proposition (one of the exercises asks you to prove it).

Proposition 10.8 (*al-Nowaihi and Dhami, 2008a*) Let $D(r, t)$ be the RS-discount function in (10.44), then:

- (a) If $0 < \rho \leq 1$, then D is subadditive.
- (b) If $\rho > 1$, then D is neither subadditive, additive, nor superadditive.
- (ci) If $0 < \tau < 1$, then D exhibits declining impatience.
- (cii) If $\tau = 1$, then D exhibits constant impatience.
- (ciii) If $\tau > 1$, then D exhibits increasing impatience.

We can now see the interpretation of the parameters τ and ρ in the RS-discount function (10.44).¹³ The parameter τ controls impatience, independently of the values of the other parameters α, β, ρ . In particular, $0 < \tau < 1$ gives declining impatience, $\tau = 1$ gives constant impatience, and $\tau > 1$ gives increasing impatience. If $0 < \rho \leq 1$, then we get subadditivity, irrespective of the values of the other parameters α, β, τ . However, if $\rho > 1$, then (10.44) can be neither subadditive, additive, nor superadditive. Depending on the particular values, we may have declining impatience, constant impatience, or increasing impatience.

Example 10.4 Consider the Read–Scholten discount function that we define in (10.44), with $\tau = 1$, $\rho = 0.6$, $\alpha = \beta = 0.5$. Let the utility function be $u(x) = x$ for $x \geq 0$. We can then compute

$$D(0, 4) = (1 + 0.5(4)^{0.6})^{-1}, D(4, 10) = (1 + 0.5(6)^{0.6})^{-1}, D(0, 10) = (1 + 0.5(10)^{0.6})^{-1},$$

which respectively, equal 0.46540, 0.40567, 0.33439. As expected, $D(0, t)$ is declining in t , and $D(r, 10)$ is increasing in r (see Definition 10.3). Let the unit of time be months. Then a decision maker who discounts over the entire interval $[0, 10]$ is indifferent between receiving \$15 in 10 months time and \$5.0159 at date $r = 0$ because

$$5.0159 = D(0, 10) 15.$$

It is common to use “stationarity” for what we call “constant impatience.” We prefer the latter, to be in conformity with “declining impatience” and “increasing impatience.”

¹³ Another reason for stating Proposition 10.8 formally is that there are several errors in the original exposition. Scholten and Read (2006a), bottom of p. 1425, state erroneously that $\alpha > 0$ implies subadditivity and $\rho > 1$ implies superadditivity. Their statement that $0 < \tau < 1$ implies declining impatience is correct but incomplete.

However, unlike the additivity assumption made in the exponential discount function, we have subadditive discounting,

$$D(0, 4) D(4, 10) = 0.1888 < D(0, 10) = 0.33439.$$

Thus, \$15 discounted back to time $r = 0$ using the two subintervals $(0, 4)$ and $(4, 10)$ leads to a lower present value $D(0, 4) D(4, 10) 15 = 2.832 < 5.0159$. It follows that under subadditive discounting, two different decision makers, one who uses a longer interval to discount and another who uses the two smaller subintervals, may make very different decisions. Since $\tau = 1$, we have from Proposition 10.8(cii) that the discount function exhibits constant impatience. For instance, $D(0, 4) = D(0 + 4, 4 + 4) = 0.46540$.

Example 10.5 (Explanation of Thaler's apples example using subadditivity; al-Nowaihi and Dhami, 2008a): Consider the explanation of the common difference effect, using subadditivity of the discount function, and Thaler's apples example. Denote by $u(x)$, the utility from outcome x . A decision maker prefers one apple today to two apples tomorrow, so that

$$u(1) > D(0, 1)u(2). \quad (10.45)$$

However, the decision maker, today, prefers to receive two apples in 51 days time to receiving one apple in 50 days time, so that

$$D(0, 50)u(1) < D(0, 51)u(2). \quad (10.46)$$

From (10.45) and (10.46) we get

$$D(0, 1)D(0, 50) < D(0, 51). \quad (10.47)$$

If we assume additivity (as in hyperbolic discounting and in the EDU model), then $D(0, 51) = D(0, 50)D(50, 51)$. Hence, under additivity, (10.47) implies that

$$D(0, 1) < D(50, 51),$$

so, under additivity, the discount function of the decision maker must exhibit declining impatience (see Definition 10.5) in order to explain the common difference effect.

An alternative explanation of the common difference effect is to relax additivity, and suppose instead that the discount function exhibits subadditivity. To see this, assume constant impatience (see Definition 10.5) of the discount function, so that $D(0, 1) = D(50, 51)$. Then (10.47) gives

$$D(0, 50)D(50, 51) < D(0, 51),$$

which is exactly the definition of subadditivity given in Definition 10.4, above. Thus, an alternative explanation of the common difference effect is to assume constant impatience (as in the EDU model), but assume subadditivity of the discount function (unlike the EDU model). It is possible that the common difference effect could be due to both declining impatience and subadditivity. Read (2001) conducted a series of experiments that tested for the common difference effect, and the experimental design could also discriminate between subadditivity and declining impatience. He found support for the common difference effect and for subadditivity, but rejected declining impatience in favor of constant impatience.

10.3.3 Uncertainty

A test of a theory (T) is always a test of T plus auxiliary assumptions (A). Thus, a refutation of T and A may be a refutation of A, rather than T. However, since A is often left implicit, a refutation of T and A may be misconstrued as a refutation of T, rather than A. A case in point may be T = “exponential discounting” and A = “uncertainty is not relevant.” Halevy (2008) argues that the common difference effect may, in fact, be a refutation of expected utility (EU) rather than exponential discounting.

Despite the experimenters’ best efforts to eliminate uncertainty, there may be a residual risk in the minds of subjects that they will not receive their promised payoffs. This could be a purely evolutionary instinct in which subjects associate the future with a risk that the outcomes might not materialize (supply side uncertainty). Or there might be a probability that the decision maker himself is not alive to receive the outcome in the future (demand side uncertainty). Tanaka et al. (2010) use an experimental design that minimizes such uncertainty by giving the money to be distributed to subjects in experiments, up-front, to a trusted third party. Further, they require the trusted third party to sign documents that the money will indeed be given to the subjects in experiments. However, this level of experimental control is not a feature of most time discounting experiments.

If subjects are EU maximizers, then risk does not matter (Example 10.6, below). However, if subjects overweight low probabilities and underweight high probabilities (as in many non-EU theories), then risk matters (Example 10.7, below). Moreover, surprisingly, the *lower* the residual risk, the *greater* will be its effect (Example 10.8, below).¹⁴

Consider the following three examples. They all involve a choice between receiving \$1000 now, or \$1100 next year, and a second choice between receiving these two sums 10 and 11 years from now, respectively. There is some probability that a future payoff is not received. We use the SIE utility function (10.12), so that

$$u(1000) = 1000.5 \text{ and } u(1100) = 1100.5. \quad (10.48)$$

Let the discount function be $D(s, t)$. We shall invoke EU and some non-EU theories of decision making in these examples; the details of the relevant decision theories can be found in Part I of the book. Let the probability weighting function be $w(p)$, where p is the probability that the outcome will actually materialize one year from now. In the special case of EU, $w(p) = p$. We assume independence of probabilities across time, so that the probability of receiving the payoff t years from now is p^t . We take the current level of wealth, w_0 , as reference wealth, and the present time to be $r = 0$.

Thus, if there is no common difference effect, then

$$(1100, 1) \prec (1000, 0) \Rightarrow (1100, 11) \prec (1000, 10). \quad (10.49)$$

Consider first the choice between periods 0 and 1. There is a probability $p \geq 0$ that the amount in period 1 will not be received. Using Definition 10.3, under *rank dependent utility* (RDU) and

¹⁴ So, experimenters, by doing their best to reduce residual risk with the aim of getting a sharper refutation of exponential discounting may, actually, be achieving the opposite.

prospect theory (PT), the present discounted utility of (1100, 1) is $D(0, 1) w(p) u(1100)$, thus¹⁵

$$(1100, 1) \prec (1000, 0) \Leftrightarrow D(0, 1) w(p) u(1100) < u(1000). \quad (10.50)$$

Similarly, the probability of receiving the amounts in periods 10 and 11 is, respectively, p^{10} and p^{11} , hence,

$$(1100, 11) \prec (1000, 10) \Leftrightarrow D(0, 11) w(p^{11}) u(1100) < D(0, 10) w(p^{10}) u(1000). \quad (10.51)$$

Example 10.6 (EDU with EU; al-Nowaihi and Dhami 2008a): Assume that the decision makers use (i) EDU, so $D(s, t) = e^{-\theta(t-s)}$, $\theta > 0$, and (ii) EU, so $w(p) = p$. Substitute in (10.50), (10.51) to get:

$$(1100, 1) \prec (1000, 0) \Leftrightarrow e^{-\theta} p u(1100) < u(1000), \quad (10.52)$$

$$(1100, 11) \prec (1000, 10) \Leftrightarrow e^{-11\theta} p^{11} u(1100) < e^{-10\theta} p^{10} u(1000), \quad (10.53)$$

(10.53) is equivalent to:

$$(1100, 11) \prec (1000, 10) \Leftrightarrow e^{-\theta} p u(1100) < u(1000). \quad (10.54)$$

From (10.52) and (10.54), we see that $(1100, 1) \prec (1000, 0) \Leftrightarrow (1100, 11) \prec (1000, 10)$. Thus, (1) payoff risk about the future, (2) EDU, and (3) EU imply no common difference effect. Hence, the observation of a common difference effect is a rejection of the joint hypothesis of payoff risk, EDU, and EU. Thus, it would imply the rejection of one or the other (or all three), but not, necessarily, EDU.

Example 10.7 (EDU with non-EU; al-Nowaihi and Dhami 2008a): We take cumulative prospect theory (Tversky and Kahneman, 1992) as our model of non-EU, while continuing to assume EDU and payoff risk. Take $D(r, t) = e^{-\theta(t-r)}$, $\theta = 0.04$, and $w(p) = e^{-(\ln p)^{0.5}}$ (the Prelec (1998) probability weighting function with $\beta = 1$, $\alpha = 0.5$; see Chapter 2 for the details). Let $p = 0.98$, so payoff risk is quite low. Substitute in (10.50), (10.51), using the utility values in (10.48), to get:

$$(1100, 1) \prec (1000, 0) \Leftrightarrow e^{-0.04} e^{-(\ln 0.98)^{0.5}} 1100.5 < 1000.5, \quad (10.55)$$

$$(1100, 11) \prec (1000, 10) \Leftrightarrow e^{-0.44} e^{-(\ln(0.98)^{11})^{0.5}} 1100.5 < e^{-0.4} e^{-(\ln(0.98)^{10})^{0.5}} 1000.5. \quad (10.56)$$

A simple calculation shows that the inequality in (10.55) holds, while the corresponding inequality in (10.56) does not hold. Hence, $(1100, 1) \prec (1000, 0)$, but $(1100, 11) \succ (1000, 10)$, so the common difference effect is found. Thus, it is quite possible that the finding of a common difference effect is not a refutation of EDU, but of payoff uncertainty and EU.

¹⁵ As in EU, one can define outcomes relative to a reference point under RDU. However, EU and RDU make no distinction in human behavior in the domain of gains and losses. By contrast, PT fundamentally distinguishes between behavior in gains and losses. The arguments below do not hinge on this difference. Hence, the reader may either take u to be the standard utility function under EU and RDU with the caveat that all outcomes are expressed relative to the reference wealth, w_0 . Or, the reader may imagine that the decision maker uses PT to deal with risk, in which case, u is the usual utility function under PT.

Example 10.8 (It's the certainty effect that's doing the work; al-Nowaihi and Dhami 2008a): Recall that in Example 10.7 there was near certainty because $p = 0.98$. We call this the certainty effect. Consider the same structure as for Example 10.7, except that now we remove the certainty effect by reducing the probability of non-receipt of the reward to $p = 0.4$, rather than $p = 0.98$. Substitute in (10.50), (10.51), and using the utility values in (10.48), to get:

$$(1100, 1) < (1000, 0) \Leftrightarrow e^{-0.04} e^{-(\ln 0.4)^{0.5}} 1100.5 < 1000.5. \quad (10.57)$$

$$(1100, 11) < (1000, 10) \Leftrightarrow e^{-0.44} e^{-(\ln(0.4)^{11})^{0.5}} 1100.5 < e^{-0.4} e^{-(\ln(0.4)^{10})^{0.5}} 1000.5. \quad (10.58)$$

A simple calculation shows that both inequalities, (10.57) and (10.58), are satisfied. Hence, $(1100, 1) < (1000, 0)$ and $(1100, 11) < (1000, 10)$. Thus, the common difference effect is due to the certainty effect in particular (as in Example 10.7), rather than uncertainty as such (as in the current example).

The typical empirical exercises use EU in conjunction with EDU. Tanaka et al. (2010) relax both these assumptions. They estimate a prospect theory utility function with the Benhabib et al. (2010) discount function that can nest the EDU and the quasi-hyperbolic functions. Conducting high stakes experiments in Vietnamese villages they find that there is no correlation in the degree of concavity of the utility function (which measures risk under EU) with income. However, mean income in the villages is correlated with loss aversion.

Examples 10.7 and 10.8 suggest that if the common difference effect is due to the certainty effect alone, rather than a combination of the certainty effect and non-exponential discounting, then the phenomenon should disappear for probabilities around 0.4. A second issue is the relevance of risk for experimental settings when outcome–time pairs are offered. Do subjects mistrust that the experimenter (typically a university researcher) will pay?¹⁶ Do subjects invoke misplaced heuristics and mistakenly believe that future payments may not be made? These questions are deserving of further exploration.

Saito (2011) generalizes the results in Halevy (2008). Saito assumes that as the delay increases, the perceived probability by the decision maker of receiving the outcome, declines. Recall, from Part I of the book, the definitions of the *common ratio effect* and the *certainty effect*. Here, we state these definitions in terms of preferences.

Denote by $X \subset \mathbb{R}$, the set of monetary outcomes. Consider binary atemporal lotteries of the form $(x, p; 0, 1 - p)$, where $x \in X$, and $p \in [0, 1]$ is the perceived probability of receiving x . We are first interested in *risk preferences*, $\succeq^r$, over atemporal lotteries. Denote the preference relation “at least as good as” over the set of such binary lotteries by $\succeq^r$, and let $\sim^r$ and $>^r$ denote, respectively, the corresponding indifference and strict preference relations.

Definition 10.6 (*Common ratio effect*): For any $x, y \in X$, $x < y$, $\bar{p} \in [0, 1]$ and $\mu \in [0, 1]$, let $(x, \bar{p}) \sim^r (y, \mu\bar{p})$. Preferences, $\succeq^r$, exhibit the common ratio effect if

$$\begin{cases} (x, p) <^r (y, \mu p) & \text{for all } 0 < p < \bar{p} \\ (x, p) >^r (y, \mu p) & \text{for all } \bar{p} \leq p \leq 1 \end{cases}.$$

In Definition 10.6, the ratio of probabilities in the two lotteries is always identical. Yet, starting from indifference at $\bar{p}$, a fall in $\bar{p}$ switches preferences towards the lottery with the higher outcome,

¹⁶ Most experiments make clear that there is no risk of not getting the reward; see for instance Thaler (1981, p. 203). See also the remark on Tanaka et al. (2010), above.

but an increase switches preferences to the lottery with the lower outcome. For the special case of $\bar{p} = 1$, the common ratio gives rise to the certainty effect.

Definition 10.7 (*Certainty effect*): For any $x, y \in X$, $x < y$, and $\mu \in [0, 1]$, let $(x, 1) \sim^r (y, \mu)$. Preferences, $\succeq^r$, exhibit the certainty effect, if for $\eta \in (0, 1)$ we have

$$(x, \eta) \prec^r (y, \mu\eta).$$

In Definition 10.7, keeping the ratio of probabilities fixed, the *certainty effect* arises when individuals switch preferences towards the smaller outcome if it is received with certainty. When even a small degree of uncertainty is introduced, the decision maker shifts to the higher outcome.

From a consideration of risk preferences, let us now move on to time preferences. Consider the sequence of dates shown in Figure 10.2. Suppose that a decision maker has time preferences $\succeq_0$ between outcome–time pairs at *time* 0, such that $0 \leq d < t < s$. The outcome–time pairs at time 0 are summarized in the set

$$T_0(X) = \{(x, t) \mid x \in X, t \in \mathbb{R}_+ \text{ and } t \geq 0\}. \quad (10.59)$$

Let $p(t)$ be the probability of receiving an outcome at time t ; this arises on account of the uncertainty about receiving an outcome in the future. The assumption that the probability of receiving a future reward diminishes as one moves further out into the future is formalized as follows:

A1: There is certainty of receiving a current outcome, i.e., $p(0) = 1$. (ii) $p(t)$ is a decreasing and continuous function, and $p(\infty) = 0$.

Let $p(t \mid d)$ be the probability of receiving the outcome at time $t \geq d$, conditional on the decision maker being alive at time d . Thus, $p(t \mid d) = \frac{p(t)}{p(d)}$. Conditional on the individual having survived until time d , define the analogue of (10.59) by

$$T_d(X) = \{(x, t) \mid x \in X, t \in \mathbb{R}_+ \text{ and } t \geq d\}.$$

The corresponding time preferences at time d are denoted by $\succeq_d$. We now focus on the time preferences, $\succeq_d$, of the decision maker over $T_d(X)$, taking into account the future uncertainty of receiving the prize. This also clarifies the relation between $\succeq_d$ (time preferences over outcome–time pairs at time d) and $\succeq^r$ (risk preferences over atemporal lotteries).

Definition 10.8 (*Time preference*): For all times $d \in \mathbb{R}_+$ and $(x, t), (y, s) \in T_d(X)$,

$$(x, t) \succeq_d (y, s) \Leftrightarrow (x, p(t \mid d)) \succeq^r (y, p(s \mid d)).$$

We are now in a position to define quasi-hyperbolic and hyperbolic discounting.

Definition 10.9 (*Quasi-hyperbolic discounting*): Suppose that $x, y \in X$, $t, s \in \mathbb{R}_+$ such that $x < y$ and $t < s$. if $(x, t) \sim_t (y, s)$, then preferences, $\succeq_d$, exhibit quasi-hyperbolic discounting if



Figure 10.2 Sequence of outcome–time pairs.

Source: Saito (2011).

$$(x, t) \prec_d (y, s) \text{ for all } d < t.$$

In Definition 10.9, there are two decision times, t and d such that $d < t$. At time d , the decision maker reports that at the future time period t , he will choose the larger–later option (a patient choice). However, once time t arrives, he switches his preference towards the smaller–sooner option, (x, t) (an impatient choice).

We now define hyperbolic discounting.

Definition 10.10 (*Hyperbolic discounting*): Let $x, y \in X$, $\bar{d}, t, s \in \mathbb{R}_+$ such that $\bar{d} < t < s$, and $x \prec y$. If $(x, t) \sim_{\bar{d}} (y, s)$, then preferences, $\succeq_d$, exhibit hyperbolic discounting if¹⁷

$$\begin{aligned} (x, t) \prec_d (y, s) & \text{ for all } d < \bar{d} \\ (x, t) \succ_d (y, s) & \text{ for all } \bar{d} < d \leq t \end{aligned}$$

The definition of hyperbolic discounting has the following features. (i) A decision maker may at some time $d < \bar{d}$ claim that he will prefer a larger–later reward $((x, t) \prec_d (y, s))$. (ii) However, when time $\bar{d}$ arrives, he professes indifference between (x, t) and (y, s) $((x, t) \sim_{\bar{d}} (y, s))$. (iii) At any time $\bar{d} < d \leq t$, the decision maker becomes sufficiently impatient that he prefers the smaller–sooner reward.

Lemma 10.1 (*Saito, 2011*): If $\succeq_d$ exhibits hyperbolic discounting, it also exhibits quasi-hyperbolic discounting.

We now state the main result that relates risk attitudes to time preferences. This is a generalization of the result in Halevy (2008) and Epper et al. (2011). The reader is asked to prove the result in the problem set.

Proposition 10.9 (*Saito, 2011*): Suppose that Assumption A1 holds.

- (a) $\succeq^r$ exhibits the common ratio effect if and only if $\succeq_d$ exhibits hyperbolic discounting.
- (b) $\succeq^r$ exhibits the certainty effect if and only if $\succeq_d$ exhibits quasi-hyperbolic discounting.

The importance of Proposition 10.9 is that it relates what was traditionally considered the subject matter of risk (common ratio effect) with the subject matter of time discounting (hyperbolic discounting). Since the Allais paradox can be stated as a common ratio violation (see Part I of the book), one intriguing interpretation of Proposition 10.9 is that the Allais paradox causes hyperbolic discounting, a possibility that is already noted in Halevy (2008) and Epper et al. (2011). Yet, and as the author notes, this is not the only possibility. Assumption A1 captures *regular uncertainty*; however, *non-regular uncertainty* (e.g., probability of receiving a future outcome may be drawn from a Poisson distribution) as in Dasgupta and Maskin (2005), may also explain hyperbolic discounting. It is also possible that the Allais paradox and hyperbolic discounting are both caused by a third factor; see for instance, Fudenberg and Levine (2009).

10.4 Attribute-based models

The outcome–time pairs that are to be compared by decision makers may be thought of as consisting of two natural attributes: time of delivery and monetary value. One can then consider the trade-off between the two attributes to determine which time–money pair will be chosen.

¹⁷ This definition also goes through with $\bar{d} \leq t \leq s$ but we wish to make the comparison particularly stark; a similar remark also applies to Definition 10.9, which also goes through with $t \leq s$.

10.4.1 Vague time preferences model

Manzini and Mariotti (2006) develop a theory of *vague time preferences* and discuss the psychological foundations for such an approach. They argue that the choice between, say, receiving \$1000 now and \$1100 next year is *clearer* than the choice between these two sums received 10 and 11 years from now, respectively. Three criteria are proposed to choose between two future outcome–time pairs (x, t) and (y, s) .¹⁸

The *primary criterion* requires that if the difference between the present discounted utility of (x, t) and (y, s) is *not similar* or it is *significantly different*, then choose the alternative that has the higher value. If the two present discounted values are *similar*, then the subject invokes the *secondary criterion* to choose the alternative that has the highest monetary value. If they have the same monetary values, so that the secondary criterion fails, then the subject behaves according to the *tertiary criterion*: “choose the outcome that is delivered sooner.” If all three criteria fail, then the subject is indifferent among the two alternatives.

This model achieves a complete, but intransitive, preference ordering. In particular, indifference here is not an equivalence relationship.¹⁹ In order to make the model operational, we require a definition of the similarity of present values of two outcome–time pairs (see the primary criterion, above). Suppose that two present values are *significantly different* if their difference is greater than $\sigma > 0$. We can state the criteria in this model, formally, as follows by using the interval discount function $D(r, t)$ (Definition 10.3).

Let $<_V$ be the binary relation induced by the vague time preferences model. Let outcomes $x, y \in \mathbb{R}$ and let r, s, t be three time periods such that $0 \leq r < t, r < s$. Then $(x, t) <_V (y, s)$ at time r , if and only if, one of the following holds:

1. (Primary criterion) $D(r, s)u(y) - D(r, t)u(x) > \sigma$, or
2. (Secondary criterion) $D(r, s)u(y) - D(r, t)u(x) \leq \sigma$, and $x < y$, or
3. (Tertiary criterion) $D(r, s)u(y) - D(r, t)u(x) \leq \sigma$, $x = y$, and $s < t$.

One can similarly state the conditions for $(y, s) <_V (x, t)$. More generally, σ can be a “vague-ness function” rather than a scalar; see the original paper for the details. Obviously, if x and y are vectors, then extra criteria can be added. In this model, present utility values whose difference is less than σ are regarded as not significantly different. This could be because, for example, the decision maker is vague about the appropriate objective function, or discount function to use. Therefore, the decision maker does not want the decision to depend too critically on the choice of these functions. On the other hand, the decision maker might be absolutely sure that more is better than less and sooner is better than later. Example 10.9, below, shows how this theory can explain the common difference effect.

Example 10.9 (*al-Nowaihi and Dhami, 2008a*): Consider the choice between receiving \$1000 now and \$1100 next year, and the choice between these two sums received 10 and 11 years from now, respectively. We use the SIE utility function (10.12), so that the utilities of the outcomes are given in (10.48): $u(1000) = 1000.5$ and $u(1100) = 1100.5$. We use the exponential discount function $D(r, t) = e^{-\theta(t-r)}$ with $\theta = 0.1$ and $r = 0$, so $D(0, t) = e^{-0.1t}$.

¹⁸ For an attribute-based model that applies to the problem of cyclic temporal choices, see Manzini and Mariotti (2007).

¹⁹ An equivalence relation on a set is a binary relation that satisfies reflexivity, symmetry, and transitivity. For instance, consider the set of students in a large university and define the equivalence relation: “having the same year of birth as.”

We take $\sigma = 3$. Using these values, we get $u(1000) - u(1100)e^{-0.1} = 1000.5 - 1100.5e^{-0.1} = 4.7264 > 3$. Hence, the primary criterion holds and the decision maker prefers \$1000 now to \$1100 next year.

Let us now consider the choice between the two sums received 10 and 11 years from now,

$$|e^{-1.1}u(1100) - e^{-1}u(1000)| = |e^{-1.1}1100.5 - e^{-1}1000.5| = 1.7388 < 3.$$

Hence, the primary criterion fails, and the decision maker considers the second criterion. Since $1000 < 1100$, the second criterion requires that the decision maker prefer \$1100 received 11 years from now to \$1000 received 10 years from now. We then have an illustration of the common difference effect using the model of vague time preferences.

The experimental results of Roelofsma and Read (2000) supported “sooner is better than later” against “larger is better than sooner.” However, if the order of the secondary criterion in the vague time preference model is reversed, so that sooner is better than larger, then, in Example 10.9, \$1000 received 10 years from now would be better than \$1100 received 11 years from now, and we would not get a common difference effect.

The validity of this explanation of the common difference effect is ultimately an empirical question. However, the main contribution of the model lies in directing the attention of theorists towards the use of primary and secondary criteria. Casual empiricism suggests that this may not be an unreasonable proposal.

10.4.2 The similarity relation and time preference

Rubinstein (2003) argues that hyperbolic discounting is a conservative extension of EDU. He proposes an even more radical departure from EDU that attempts to open the black box of human decision making by relying on the *procedural aspects of decision making*. Based on his earlier work in decision theory (see Rubinstein, 1988), he proposes the use of a *similarity relationship* in evaluating outcome–time pairs. Suppose that an economic agent must compare two outcome–time alternatives: (z_1, t_1) and (z_2, t_2) . There are two *attributes* in such decisions. *Size* of the prizes, z_1, z_2 , and the *time delays*, t_1, t_2 . A decision maker who uses the similarity relation, applies the following three stage procedure.

1. If $z_1 > z_2$ and $t_2 > t_1$, then the larger prize is available at an earlier date, hence, (z_1, t_1) is the preferred alternative.
2. If there is *similarity* in one attribute dimension, then the choice is made by comparing the two alternatives in the remaining attribute dimension. So, for instance, if z_1 and z_2 are similar, then the decision maker chooses the alternative that is available earlier ($\min\{t_1, t_2\}$). Similarly, if t_1 and t_2 are perceived to be similar, then the chosen alternative is the one that offers a larger prize ($\max\{z_1, z_2\}$).
3. If there is no similarity in either dimension, then no decision criterion is provided by the theory. Hence, the preference relationship in this model is incomplete.

This model is a promising step, insofar as it attempts to open the black box of human decision making. However, it needs further development, as is illustrated in the first of three examples in Rubinstein (2003) that is based on hypothetical payoffs.²⁰

²⁰ A small sample of the participants were selected and paid a fee of \$20 that was independent of their choices in the experiment.

Example 10.10 *The experiment was conducted in Fall 2002. The students in the experiment were randomly allocated to one of two questions.*

Q1: Choose among one of the two options:

A. Receive \$467.00 on June 17 2004.

B. Receive \$607.07 on June 17 2005.

Q2: Choose among one of the two options:

A. Receive \$467.00 on June 16 2005.

B. Receive \$467.39 on June 17 2005.

In Q1, 55% chose the delay option (B) while in Q2, 46% chose the delay option (B).

Rubinstein (2003) argued that the choices made in the experiment are not consistent with hyperbolic discounting, but are consistent with the use of the similarity procedure. To see this, assume that the utility function is $u(x) = \ln x$, and the discount function is of the hyperbolic form, given in (9.29) with $k = 1.5$.

Questions Q1 and Q2 are not explicit about the date at which students should consider their choices. Assume that the date of the experiment is the relevant date. However, the exact date is not given, so we take September 1, 2002, the start of Fall in 2002, as the decision date. There are 655 days till June 17, 2004, 1020 days until June 17, 2005, and so on. Using this information, and measuring time in days, we get that the decision maker chooses the no delay option (A) in both cases because²¹

$$\frac{1}{1 + 1.5(655)} \ln(467) > \frac{1}{1 + 1.5(1020)} \ln(607.07),$$

$$\frac{1}{1 + 1.5(1019)} \ln(467) > \frac{1}{1 + 1.5(1020)} \ln(467.39).$$

No standard error figures are given, hence, it is not possible to know if the proportions of the students choosing the no-delay option are significantly different between the two questions. In Q2, hyperbolic discounting explains more than half the choices and in Q1, it explains less than half the choices.

The similarity relation can be used to explain the no-delay choice (A) in Q2 in the following manner. The two amounts differ only by 39 cents, so they could be perceived to be *similar*. The decision maker then chooses the earlier option, which is A. This raises the question of what exactly is “similar”? From the two questions, it appears that the smallest unit of time is a day. Hence, one may argue that there is similarity in the time dimension in Q2, in which case the decision maker should choose the larger amount associated with option B. Thus, depending on what one means by similarity, one may justify any of the available choices, which is problematic for theory.

There are also other issues with making the similarity relation operational. First, the similarity relationship does not specify how close two entities must be in order to qualify to be *similar*. Second, the procedure does not apply to choice situations when there is no similarity in either dimension (see stage 3 of the similarity procedure above). Third, the procedure becomes less suitable as the number of choice pairs (x_i, t_i) increase, because the order in which choices are compared may become important in this case. Fourth, it is not clear which attribute to choose first in order to use the similarity relation. Hence, although the similarity relationship seems to be a

²¹ This conclusion is not altered if we choose the simplest discount function, $D(t) = \frac{1}{t}$.

promising choice heuristic, its domain of application would appear to be limited. As such, barring further development, it does not seem to be in a position to replace hyperbolic discounting in terms of the analytical convenience to model problems of present-biased preferences; a point also recognized by Rubinstein (2003).

10.4.3 Trade-off attribute model of intertemporal choice

We now consider the multiple attribute model of Scholten and Read (2006b), the *trade-off model of intertemporal choice* that focuses on a “trade-off” between the time dimension and the outcome dimension.

Suppose that a decision maker has a choice between a smaller-sooner (SS) reward, (z_S, t_S) , and a larger-later (LL) reward, (z_L, t_L) , such that $0 < z_S < z_L$ and $t_S < t_L$. The advantage of the SS reward, denoted by $a_S(t_S, t_L)$, is that it arrives earlier. The disadvantage of the SS reward, denoted by $d_S(z_S, z_L)$, is that one forgoes the larger outcome in the future. Now consider the case of losses $z_L < z_S < 0$. Here, $a_L(t_S, t_L)$ is the advantage of the LL loss because it occurs at a later time, so discounted losses feel smaller. Similarly, $d_L(z_S, z_L)$ is the disadvantage of the LL loss because the absolute magnitude of the loss is larger.

In the simplest form, suppose that the functional forms for a, d are given by:²²

$$\begin{cases} a_S(t_S, t_L) = \kappa(t_L - t_S); d_S(z_S, z_L) = z_L - z_S & \text{if } 0 < z_S < z_L \\ a_L(t_S, t_L) = \kappa(t_L - t_S); d_L(z_S, z_L) = \lambda |z_L - z_S| & \text{if } z_L < z_S < 0 \end{cases}, \quad (10.60)$$

where $\kappa > 0$ ensures compatibility between the units in which time and outcomes are measured. For instance, moving from time units of weeks to months increases κ by a factor of four. $\lambda > 1$ is akin to the parameter of loss aversion in prospect theory.

Let $\sim$ denote the indifference relation, and $>$ the strict preference relation. In the trade-off model of intertemporal choice, the decision maker uses the following decision rule for $0 < z_S < z_L$:

$$\begin{aligned} a_S(t_S, t_L) = d_S(z_S, z_L) & \text{ then } SS \sim LL \\ a_S(t_S, t_L) > d_S(z_S, z_L) & \text{ then } SS > LL \\ a_S(t_S, t_L) < d_S(z_S, z_L) & \text{ then } LL > SS \end{aligned} \quad (10.61)$$

z_S, z_L are measured relative to some reference point. If the reference point is current wealth, then the amounts offered in experiments may be thought to constitute gains and losses relative to current wealth.

This model can explain the magnitude effect, present bias, and the sign effect. Suppose that $a_S(t_S, t_L), d_S(z_S, z_L)$ are given by (10.60), $a_S(t_S, t_L) = d_S(z_S, z_L)$, and $z_S > 0, z_L > 0$. Then, the trade-off model predicts $SS \sim LL$. Starting with $a_S(t_S, t_L) = d_S(z_S, z_L)$, consider the following three possible perturbations.

1. Suppose that both outcomes $z_S > 0, z_L > 0$ are multiplied by a factor, $\gamma > 1$. Then, we compare the two outcome-time pairs $(\gamma z_S, t_S), (\gamma z_L, t_L)$. Since $a_S(t_S, t_L) = d_S(z_S, z_L)$, we get from (10.60) that $a_S(t_S, t_L) < d_S(\gamma z_S, \gamma z_L)$. Using (10.61), we get $LL > SS$, as predicted by the magnitude effect.

²² We consider here only the linear version of the model that is not able to take account of declining sensitivity of gains and losses. A non-linear extension is also available.

2. Suppose that the decision maker reveals at some time $t < t_S$ that $a_S(t_S, t_L) = d_S(z_S, z_L)$. Now suppose that we bring time t_S closer to t by multiplying t_S by $\rho < 1$ to generate a new pair of choices $(z_S, \rho t_S), (z_L, t_L)$. From (10.60), we now get that $a_S(\rho t_S, t_L) > d_S(z_S, z_L)$, so (10.61) implies that $SS > LL$. This prediction is consistent with models of present bias such as hyperbolic discounting; as rewards are brought closer to the present, the decision maker switches preferences to the more immediate reward.
3. Initially we have $0 < z_S < z_L$ and $a_S(t_S, t_L) = d_S(z_S, z_L)$, so (10.61) implies that $SS \sim LL$. Now suppose that the signs of z_S, z_L are reversed and we have the situation $z_L < z_S < 0$. Since $\lambda > 1$ we now get that $a_L(t_S, t_L) < d_L(z_S, z_L)$, therefore, we have $SS > LL$. Thus, losses are more salient, which is consistent with the sign effect.

10.5 The reference time theory (RT)

Within the class of discounted utility models, what should a general theory of time preference aspire to? Two desirable elements seem to present themselves. First, it should be able to explain the main anomalies of the EDU model, namely, gain-loss asymmetry, magnitude effect, common difference effect, delay-speedup asymmetry, apparent intransitivity of time preferences, and non-additivity of time discounting. Second, it should provide a framework that can incorporate some of the recent developments in time discounting such as uncertainty and attribute models.

al-Nowaihi and Dhami (2008a, 2013, 2014) incorporate both these desirable elements in their *reference time theory* (RT). RT is a synthesis of three earlier lines of work, namely, LP (Loewenstein and Prelec, 1992), PPL (Phelps and Pollak, 1968; Laibson, 1994, 1997), and RRS (Roelofsma and Read, 2000; Read, 2001; Scholten and Read, 2006a). In a nutshell, RT is basically LP extended to allow for non-additive time discounting by incorporating a reference point for time and reference point for outcomes.

LP introduced reference dependence for outcomes in the DU model. Suppose that a decision maker, at time $t = 0$, takes an action that gives rise to the sequence of outcome-time pairs: $(w_1, t_1), (w_2, t_2), \dots, (w_n, t_n)$, where w_i is the outcome at time t_i , $i = 1, 2, \dots, n$. LP proposed the following intertemporal utility function for the decision maker:

$$U = \sum_{i=1}^n v(w_i - w_0) D(0, t_i). \quad (10.62)$$

In (10.62), v is the utility function under prospect theory, due to Kahneman and Tversky (1979). v is defined over outcomes relative to a reference point, $x_i = w_i - w_0$, where w_0 is the reference outcome level. When $x_i \geq 0$, the individual is said to be in the *domain of gains*, and if $x_i < 0$, the individual is in the *domain of losses*. In LP, $D(0, t)$ is the general hyperbolic discounting function that discounts the utility at time $t \geq 0$ back to the present, time $t = 0$.

In RT, a *reference time* is introduced, in addition to the reference outcome level. In LP, there is a reference outcome level, but, by default, the time period 0 is the reference time to which all outcomes are discounted back. In many cases, it may be desirable to have a reference time, $r \geq 0$. Consider the following two examples.

Example 10.11 At time 0, a business may be contemplating an investment in a new plant that will become operational at time $t_1 > 0$. In calculating the viability of the investment, the business may wish to compare the future profits and costs from the investment discounted back to time t_1 , which then becomes the reference time.

Example 10.12 At time 0, a student starting a three year degree may wish to use a government loan to finance the degree. The government may require loan repayments to be made only upon completion of the degree (as in the UK). The student may wish to calculate the future benefits and costs of this action at time 0, discounted back to the time at which the degree is completed, t_1 , which then becomes the reference time.

RT alters the LP framework in two ways. First, it augments it to allow for a reference time as well as a reference outcome level. Second, a general class of discount functions, D , is considered that allows for declining impatience, subadditivity, or a combination of both. We show that if the discount function is *additive* (Definition 10.4), then preferences do not depend on the reference time. Then we show that some (not all) of the observed intransitivities in outcome–time pairs may be due to reference time shifts. One’s reference time may depend on the different binary choices under consideration.

The trade-off model of Scholten and Read (2006b) can be accommodated within RT. Using examples we also show that the treatment of uncertainty in Halevy (2007) and the vague time preference model of Manzini and Mariotti (2006) can also be incorporated within RT. Finally we compare our model with the model of intransitive preferences proposed in Ok and Masatlioglu (2007).

10.5.1 A basic description of RT

Consider a decision maker who, at time t_0 , takes an action that results in the outcome w_i at time t_i , $i = 1, 2, \dots, n$, where

$$t_0 \leq r \leq t_1 < \dots < t_n.$$

Time r is the *reference time*. It is the time back to which all values are to be discounted, using a discount function, $D(r, t)$ (see Definition 10.3). The reference time r need not be the same as t_0 (recall Examples 10.11, 10.12). Without loss of generality, we normalize the time at which the decision is made to be $t_0 = 0$.

As in LP, the decision maker has a *reference outcome level*, w_0 , relative to which all outcomes are to be evaluated. Letting $x_i = w_i - w_0$, the utility of outcome x_i is $v(x_i)$, where v is the prospect theory utility function that has four main properties: *reference dependence*, *monotonicity*, *declining sensitivity*, and *loss aversion*.

Definition 10.11 The properties of v are:

- (i) $v : (-\infty, \infty) \rightarrow (-\infty, \infty)$ is increasing (monotonicity).
- (ii) $v(0) = 0$ (reference dependence).
- (iii) For $x > 0$: $-v(-x) > v(x)$ (loss aversion).
- (iv) v is concave for gains, $x \geq 0$, but convex for losses, $x \leq 0$ (declining sensitivity).
- (v) v is continuous, and twice differentiable except at 0.

For each *reference outcome* and *reference time* pair $(w_0, r) \in (-\infty, \infty) \times [0, \infty)$, the decision maker has a complete and transitive preference relation, $\preceq_{w_0, r}$ on $(-\infty, \infty) \times [r, \infty)$ given by

$$(w_1, t_1) \preceq_{w_0, r} (w_2, t_2) \Leftrightarrow v(w_1 - w_0) D(r, t_1) \leq v(w_2 - w_0) D(r, t_2). \quad (10.63)$$

Let S be a non-empty set of outcome–time sequences from $(-\infty, \infty) \times [0, \infty)$ of the form $(x_1, t_1), (x_2, t_2), \dots, (x_i, t_i), \dots$, where outcomes are expressed relative to the reference outcome.

Using (10.63), we extend $\preceq_{w_0, r}$ to a complete transitive preference relation on sequences in S , as follows (this also holds for infinite sequences, provided the sums in (10.64) converge):

$$\begin{aligned} ((x_1, s_1), (x_2, s_2), \dots, (x_m, s_m)) &\preceq_{w_0, r} ((y_1, t_1), (y_2, t_2), \dots, (y_n, t_n)) \\ \Leftrightarrow \sum_{i=1}^m v(x_i) D(r, t_i) &\leq \sum_{i=1}^n v(y_i) D(r, t_i). \end{aligned} \quad (10.64)$$

Thus, the decision maker's intertemporal utility function is given by:

$$V_r((w_1, t_1), (w_2, t_2), \dots, (w_n, t_n) \mid w_0) = \sum_{i=1}^n v(x_i) D(r, t_i) \quad (10.65)$$

Comparing (10.65) to the utility function (10.62), proposed by LP, and successfully employed to explain several anomalies in intertemporal choice, we see that LP implicitly set $r = 0$. Also, LP implicitly assume that the discount function is additive (Definition 10.4), in which case, as we show below, there is no loss in assuming $r = 0$. To accommodate the empirical evidence, however, we allow the discount function to be non-additive, in which case the choice of reference time does matter.

We now consider the determination of the reference point for time, r . Let

$$T = \{t \in [0, \infty) : t = t_i \text{ for some sequence } \{(x_1, t_1), (x_2, t_2), \dots, (x_i, t_i), \dots\} \text{ in } S\}. \quad (10.66)$$

Since T is bounded below by 0 and non-empty, it follows that T has a greatest lower bound, r . We make the following assumption for reference time that is motivated by the status-quo justification of reference points, as the time closest to, or identical to the current time that an individual contemplates a decision (Kahneman and Tversky, 2000).

A0: Given S , T , r , as described above, the decision maker takes r (the greatest lowest bound of T) as the reference point for time.

A0 should be regarded as a tentative assumption, whose implications are to be more fully explored. However, the theory is flexible enough to allow for other reference points for time.

Example 10.13 *If a decision maker wishes to compare (x, s) with (y, t) , $s \leq t$, then $S = \{(x, s), (y, t)\}$ and $T = \{s, t\}$. Thus, A0 implies that $r = s$. If $v(x) < v(y) D(s, t)$, then the decision maker chooses (y, t) over (x, s) ; note from Definition 12, $D(s, s) = 1$.*

10.5.2 The implications of additivity of discount functions for RT

We first show that if the discount function is additive, then the choice of the reference point for time is irrelevant.

Proposition 10.10 *(al-Nowaihi and Dhami, 2008a, 2013): Suppose that the discount function is additive (Definition 10.4), then $(x_1, t_1) \preceq_{w_0, r} (x_2, t_2)$ if, and only if, $(x_1, t_1) \preceq_{w_0, 0} (x_2, t_2)$.²³*

Proof: First suppose that $(x_1, t_1) \preceq_{w_0, r} (x_2, t_2)$ and recall that $x_i = w_i - w_0$, $i = 1, 2$. Then we get

²³ This proposition can be easily extended to sequences of time outcome pairs. For instance, al-Nowaihi and Dhami (2013) prove the following result. If $D(r, t)$ is additive, then

$$\begin{aligned} ((x_1, s_1), (x_2, s_2), \dots, (x_m, s_m)) &\preceq_{w_0, r} ((y_1, t_1), (y_2, t_2), \dots, (y_n, t_n)) \\ \Leftrightarrow ((x_1, s_1), (x_2, s_2), \dots, (x_m, s_m)) &\preceq_{w_0, 0} ((y_1, t_1), (y_2, t_2), \dots, (y_n, t_n)). \end{aligned}$$

$$D(r, t_1) v(x_1) \leq D(r, t_2) v(x_2). \quad (10.67)$$

From Definition 10.3, $D(0, t) > 0$ for any finite t , hence, multiplying both sides of (10.67) by $D(0, r) > 0$ we get

$$D(0, r)D(r, t_1) v(x_1) \leq D(0, r)D(r, t_2) v(x_2). \quad (10.68)$$

Additivity (Definition 10.4) implies that

$$D(0, r)D(r, t) = D(0, t). \quad (10.69)$$

Using (10.68), (10.69) we get that

$$D(0, t_1)v(x_1) \leq D(0, t_2)v(x_2) \Rightarrow (x_1, t_1) \preceq_{w_0, 0} (x_2, t_2).$$

Now assume that $(x_1, t_1) \preceq_{w_0, 0} (x_2, t_2)$, then (10.67) implies that

$$D(0, t_1) v(x_1) \leq D(0, t_2) v(x_2). \quad (10.70)$$

Using (10.69), (10.70) we get

$$D(r, t_1) v(x_1) \leq D(r, t_2) v(x_2) \Rightarrow (x_1, t_1) \preceq_{w_0, r} (x_2, t_2). \blacksquare$$

In contrast to the result in Proposition 10.10, if the discount function is non-additive, then the choice of the reference point for time matters. RT uses this feature to explain (apparent) intransitivity in preferences as a framing effect due to a change in the reference point for time.

10.5.3 Explanation of intransitive preferences under RT

We now illustrate one possible resolution of apparent intransitivity of preferences using the RT theory (al-Nowaihi and Dhami, 2008a). This argument is similar in spirit to Gintis (2009) who argues that apparent intransitivities may be explained by taking account of the context, frame, and the state of the world. Let the decision time be the current period, $t = 0$, and consider the following preferences over outcome–time pairs.

$$(1, 0) \prec (3, 2); (3, 2) \prec (2, 1); (2, 1) \prec (1, 0). \quad (10.71)$$

Take the reference point for wealth to be the current level of wealth, so each outcome in (10.71) is regarded as a gain relative to current wealth. Evaluated from the perspective of the current time, $t = 0$, the stated preferences in (10.71) satisfy $V_0(1, 0) < V_0(3, 2) < V_0(2, 1) < V_0(1, 0)$, where V_0 (for $r = 0$) is defined in (10.65). Thus, preferences are intransitive. Ok and Masatlioglu (2007, p.215) use a similar example to motivate their intransitive theory of relative discounting.

Alternatively, we may view the preference pattern in (10.71) as due to a framing effect that results in a shift in the reference point for time. Assume that, in conformity with Assumption A0, the choice of reference time in each pairwise comparison is the sooner of the two dates (the status-quo argument for a reference point). Thus, (10.71) implies the following reference points

in pairwise comparisons reading from left to right: $r = 0$, $r = 1$, and $r = 0$. Then, recalling the definition of V_r in (10.65), (10.71) can be formalized as follows.

$$\begin{aligned} V_0(1, 0) &< V_0(3, 2), \\ V_1(3, 2) &< V_1(2, 1), \\ V_0(2, 1) &< V_0(1, 0). \end{aligned} \tag{10.72}$$

There is no intransitivity in (10.72). Thus, the apparent intransitivity in (10.71) arises from conflating $V_0(3, 2)$ with $V_1(3, 2)$ and $V_1(2, 1)$ with $V_0(2, 1)$. The following example shows that (10.72) is consistent with RT.

Example 10.14 (*al-Nowaihi and Dhami, 2008a*): Take the utility function to be a SIE utility function as in (10.12):

$$v(x) = x^{0.5} (1 + x)^{0.5}, \quad x \geq 0. \tag{10.73}$$

Thus (working to five significant figures),

$$v(1) = 1.414\,2, \quad v(2) = 2.449\,5 \text{ and } v(3) = 3.464\,1. \tag{10.74}$$

As our discount function, we take the Read-Scholten discount function (10.44) above, with $\alpha = \beta = 1$ and $\rho = \tau = 0.5$:

$$D(r, t) = \left(1 + (t^{0.5} - r^{0.5})^{0.5}\right)^{-1}. \tag{10.75}$$

Thus,

$$D(0, 1) = 0.5, \quad D(1, 2) = 0.608\,42 \text{ and } D(0, 2) = 0.456\,79. \tag{10.76}$$

From (10.74) and (10.76) we get

$$V_0(1, 0) = v(1) D(0, 0) = 1.414\,2, \tag{10.77}$$

$$V_0(3, 2) = v(3) D(0, 2) = 1.582\,4, \tag{10.78}$$

$$V_1(3, 2) = v(3) D(1, 2) = 2.107\,6, \tag{10.79}$$

$$V_1(2, 1) = v(2) D(1, 1) = 2.449\,5, \tag{10.80}$$

$$V_0(2, 1) = v(2) D(0, 1) = 1.224\,8. \tag{10.81}$$

From (10.77) to (10.81), we get

$$V_0(1, 0) < V_0(3, 2); \quad V_1(3, 2) < V_1(2, 1); \quad V_0(2, 1) < V_0(1, 0), \tag{10.82}$$

confirming (10.72).

As indicated by Proposition 10.10, no additive discount function (e.g., exponential, (β, δ) form of hyperbolic discounting, or LP's generalized hyperbolic discounting) can explain apparently intransitive choices as exhibited in (10.71). The reason is that with an additive discount function, all utilities are necessarily discounting back to the same base period (typically time $t = 0$), hence, showing up apparent intransitivities.

Table 10.1 Examples of discount functions when there is a reference point for time.

Type	$D(r, t)$
Exponential	$e^{-\beta(t-r)}, \beta > 0.$
PPL	$\begin{cases} 1 & \text{if } r = t = 0 \\ e^{-(\delta+\beta t)}, \beta > 0, \delta > 0, & \text{if } 0 = r < t \\ e^{-\beta(t-r)} & \text{if } 0 < r \leq t \end{cases}$
LP	$\left(\frac{1+\alpha t}{1+\alpha r}\right)^{-\frac{\beta}{\alpha}}, t \geq 0, r \geq 0, \alpha > 0, \beta > 0.$
RS	$[1 + \alpha(t^\tau - r^\tau)^\rho]^{-\frac{\beta}{\alpha}}, 0 \leq r \leq t, \alpha > 0, \beta > 0, \rho > 0, \tau > 0.$
GRS	$e^{-Q[w(t)-w(r)]}, 0 \leq r \leq t, \text{ where}$ $Q : [0, \infty) \xrightarrow{\text{onto}} [0, \infty) \text{ is strictly increasing,}$ $w : [0, \infty) \xrightarrow{\text{onto}} [0, \infty) \text{ is strictly increasing.}$

10.5.4 Recasting common discount functions in terms of RT

In theories of time discounting that do not have a reference time, the discount functions are stated under the implicit assumption that $r = 0$. In the presence of reference time, we first need to restate the main discount functions for the case $r > 0$. We extend five common discount functions for the case $r > 0$. The standard versions of these functions can simply be obtained by setting $r = 0$.

The first three discount functions in Table 10.1 have already been discussed above. These are, respectively, the *exponential discount function* proposed by Samuelson (1937), the $\beta - \delta$ or *quasi-hyperbolic discount function* proposed by Phelps and Pollak (1968) and Laibson (1997), and the *generalized hyperbolic discount function* proposed by Loewenstein and Prelec (1992). The labels PPL and LP reflect the names of the authors who proposed these functions. These three discount functions are *additive* (Definition 10.4). They can account for the common difference effect through declining impatience (Definition 10.5), but they cannot account for either non-additivity, or intransitivity.

The RS interval discount function in Table 10.1 was introduced by Read (2001) and Scholten and Read (2006a). It can account for both non-additivity and intransitivity. It can account for the common difference effect through declining impatience, subadditivity, or a combination of both.

Note that the LP approaches the exponential as $\alpha \rightarrow 0$. In general, neither of LP, or RS is a special case of the other. However, for $r = 0$ (and only for $r = 0$), RS reduces to LP when $\rho = \tau = 1$. While ρ, τ are parameters, r is a variable. Hence, neither discount function is a special case of the other.²⁴

It is straightforward to check that each of exponential, LP, RS, and generalized RS is a continuous discount function in the sense of Definition 10.3. It is also straightforward to check that PPL is a discount function. The reason the latter is not a *continuous* discount function is that $\lim_{t \rightarrow 0+} D(0, t) = e^{-\delta} < 1 = D(0, 0)$.

²⁴ Scholten and Read (2006a) report incorrectly that the LP discount function is a special case of the RS discount function.

10.5.5 Explanation of the EDU anomalies in RT

The RT theory is able to explain the six main anomalies of the EDU model in one unified framework. We only explain here the thrust of these explanations. These anomalies are: gain–loss asymmetry, magnitude effect, common difference effect, delay–speedup asymmetry, apparent intransitivity of time preferences, and non-additivity of time discounting. We now outline the intuition behind RT’s explanation of the anomalies; for the details, see al-Nowaihi and Dhami (2008a, 2013, 2014).

RT explains the magnitude and sign effects as in LP (see Section 10.2). However, it employs the SIE class of value functions, which are due to al-Nowaihi and Dhami (2009), that satisfy the requirements in LP. RT can also explain the sign effect by using different discount functions for gains and losses.

RT can explain the common difference effect as due to either declining impatience (Definition 10.5), subadditivity (Definition 10.4), or a combination of both. However, the approach of RT is more general, as RT can also explain the common difference effect as due to the presence of a small amount of irremovable uncertainty, as in Halevy (2008), or as a consequence of multiple decision criteria, as in Manzini and Mariotti (2006). Several of the attribute-based models considered in Section 10.4 can be incorporated within RT. We shall not consider these issues here but the interested reader can consult al-Nowaihi and Dhami (2008a). Thus, RT can accommodate all the known explanations of the common difference effect. In the spirit of scientific inquiry, RT leaves it to empirical evidence to select the correct explanation.

Here, we give an updated exposition of RT theory from al-Nowaihi and Dhami (2014).²⁵ We specify our axioms in terms of three functions: The *delay function*, Ψ , the *speedup function*, f , and the *generating function*, φ . Each of these functions can be uniquely elicited from behavior. Ψ determines stationarity and the common difference effect (Proposition 10.13, below). f determines whether impatience is declining, constant, or increasing (Proposition 10.14, below). Ψ and f jointly determine whether the preferences are subadditive, additive, or superadditive (Proposition 10.15, below). The f and φ functions jointly determine a unique discount function. Conversely, a continuous discount function determines unique f and φ .²⁶

DELAY FUNCTION, Ψ

Let the reference time be $r \geq 0$. Suppose that a decision maker is indifferent between x received at time r and y received at time $t \geq r$, thus,

$$v(x) = v(y) D(r, t). \quad (10.83)$$

Now suppose that the receipt of x is *delayed* to time s . We ask, at what time, T , will y received at time T be equivalent to x received at time s , i.e., for what T does the following hold?

$$v(x) D(r, s) = v(y) D(r, T). \quad (10.84)$$

²⁵ For an earlier version of the theory, see al-Nowaihi and Dhami (2008a). The main difference between the two papers is that in the later version, al-Nowaihi and Dhami (2014) use only three underlying functions, each of which can be elicited from observed behavior, but they are able to reproduce all the results from the earlier paper. By contrast, al-Nowaihi and Dhami (2008a) employ four underlying functions, one of which may be difficult to elicit from behavior.

²⁶ We do not formally state this result below but the interested reader may wish to construct the proof; the proofs are available in al-Nowaihi and Dhami (2014).




Figure 10.3 The delay and the speedup functions.

Let us conjecture that T depends on r, s, t through a functional relation, say, $T = \Psi(r, s, t)$; see Figure 10.3. From (10.83), (10.84) we get that $\Psi(r, s, t)$ must satisfy

$$D(r, s) D(r, t) = D(r, \Psi(r, s, t)). \quad (10.85)$$

We shall call the function $\Psi(r, s, t)$, if it exists, a delay function (see Definition 10.12). For the exponential discount function (see Table 10.1), $\Psi(r, s, t) = s + t$.

Definition 10.12 (Delay functions): Let D be a discount function. Suppose that the function, Ψ , has the property $D(r, s) D(r, t) = D(r, \Psi(r, s, t))$, $s \geq r, t \geq r$. Then we call Ψ a delay function corresponding to the discount function, D .

Next, we show that a delay function exists, is unique, and depends on r, s, t .

Proposition 10.11 (Existence and properties of a delay function; al-Nowaihi and Dhami, 2014):

- (i) Let D be a discount function and Ψ a corresponding delay function. Then, Ψ has the following properties:
 - (a) Ψ is unique,
 - (b) $\Psi(r, s, t)$ is strictly increasing in each of s and t ,
 - (c) $\Psi(r, s, t) = \Psi(r, t, s)$,
 - (d) $\Psi(r, r, t) = \Psi(r, t, r) = t$,
 - (e) $v(x) = v(y) D(r, t)$ if and only if, $v(x) D(r, s) = v(y) D(r, \Psi(r, s, t))$.
- (ii) A continuous discount function has a unique delay function, $\Psi(r, s, t)$. $\Psi(r, s, t)$ is continuous in s, t .

SPEEDUP FUNCTION, f

Suppose that the decision maker is indifferent between x received at time r and y received at time t , $0 \leq r \leq t$, where time r is the reference time. Then, recalling that $D(r, r) = 1$ we have

$$v(x) = v(y) D(r, t). \quad (10.86)$$

Suppose that the receipt of x is brought forward from time r to time 0, where time 0 is the new reference time. We ask, at what time, T , will y received at time T be equivalent to x received at time 0? Or, for what time, T , will the following hold?

Table 10.2 The delay function, speedup function, and generating function corresponding to common discount functions.

	$\varphi(t)$	$f(r, t)$	$\Psi(r, s, t)$
Exponential	$e^{-\beta t}$	$t - r$	$s + t$
PPL ($r = t = 0$)	1	0	s
PPL ($0 = r < t$)	$e^{-(\delta + \beta t)}$	t	$\frac{\delta}{\beta} + s + t$
PPL ($0 < r \leq t$)	$e^{-\beta t}$	$t - r$	$s + t$
LP	$(1 + \alpha t)^{-\frac{\beta}{\alpha}}$	$\frac{t-r}{1+\alpha r}$	$s + t + \alpha st$
RS	$[1 + \alpha t^{\tau\rho}]^{-\frac{\beta}{\alpha}}$	$(t^{\tau} - r^{\tau})^{\frac{1}{\tau}}$	$[s^{\tau\rho} + t^{\tau\rho} + \alpha(st)^{\tau\rho}]^{\frac{1}{\tau\rho}}$
Generalized RS	$e^{-Q[w(t)]}$	$w^{-1}[w(t) - w(r)]$	$w^{-1}Q^{-1}[Q(w(s)) + Q(w(t))]$

$$v(x) = v(y) D(0, T). \quad (10.87)$$

For the exponential discount function, we have $T = t - r$. More generally, let us conjecture that T depends on r, t , so that we can write $T = f(r, t)$, where $f: \Delta \rightarrow [0, \infty)$ will be called the *speedup function* (Δ is given in Definition 10.3); see Figure 10.3.

Definition 10.13 (*Speedup function*): Let $f: \Delta \rightarrow [0, \infty)$ satisfy:

- (a) For each $r \in [0, \infty)$, $f(r, t)$ is a strictly increasing function of $t \in [r, \infty)$ into $[0, \infty)$, with $f(r, r) = 0$.
- (b) For each $t \in [0, \infty)$, $f(r, t)$ is a strictly decreasing function of $r \in [0, t]$ into $[0, t]$, with $f(0, t) = t$.

Then we call f a *speedup function*. If, in (a), “into” is replaced with “onto,” then we call f a *continuous speedup function*.

A “continuous speedup function,” $f(r, t)$, is continuous in t .

THE GENERATING FUNCTION

Our third useful function is the generating function, φ . Like the delay function, Ψ , and the speedup function, f , the generating function, φ , can be uniquely elicited from behavior.

Definition 10.14 (*The generating function*): Let $\varphi: [0, \infty) \rightarrow (0, 1]$ be a strictly decreasing function with $\varphi(0) = 1$. Then we call φ a *generating function*. If, in addition, φ is onto, we call φ a *continuous generating function*.

On account of reference dependence of the prospect theory utility function, $v(0) = 0$. Such a function is measurable on the ratio scale, i.e., u and v represent the same prospect theory preferences if, and only if, $u = \alpha v$ for some positive constant, α . The following proposition is an immediate consequence of this.

Proposition 10.12 (*Uniqueness of the generating function; al-Nowaihi and Dhami, 2014*): The generating function, φ , is uniquely determined by behavior.

Example 10.15 Table 10.2 gives the generating function, φ , and the speedup function, f , of each of the discount functions $D(r, t)$, given in Table 10.1.

CHARACTERIZING OBSERVED BEHAVIOR FROM THE UNDERLYING FUNCTIONS

We now show how restrictions on the three functions, delay, speedup, and generating functions, can be used to characterize many of the important behaviors in Part 3. Among others, these include stationarity, common difference effect, declining impatience, and subadditivity.

Proposition 10.13 (*al-Nowaihi and Dhami, 2014*): A continuous discount function with the delay function, Ψ , exhibits:

- (a) Stationarity if, and only if, $\Psi(r, s, t) = s + t$, for all $r \geq 0, s \geq r, t \geq r$.
- (b) The common difference effect if, and only if, $\Psi(r, s, t) > s + t$, for all $r \geq 0, s > r, t > r$.

Proposition 10.14 (*al-Nowaihi and Dhami, 2014*): A continuous discount function with the speedup function, f , exhibits:

- (a) Declining impatience if, and only if, $f(r, s) > f(r + t, s + t)$, for all $t > 0, s > r$.
- (b) Constant impatience if, and only if, $f(r, s) = f(r + t, s + t)$, for all $t \geq 0, s \geq r$.
- (c) Increasing impatience if, and only if, $f(r, s) < f(r + t, s + t)$, for all $t > 0, s > r$.

Proposition 10.15 (*al-Nowaihi and Dhami, 2014*): A continuous discount function with delay function, Ψ , and speedup function, f , exhibits:

- (a) Subadditivity if, and only if, $f(r, t) < \Psi(0, f(r, s), f(s, t))$, for all $0 \leq r < s < t$.
- (b) Additivity if, and only if, $f(r, t) = \Psi(0, f(r, s), f(s, t))$, for all $0 \leq r \leq s \leq t$.
- (c) Superadditivity if, and only if, $f(r, t) > \Psi(0, f(r, s), f(s, t))$, for all $0 \leq r < s < t$.

Let us provide a brief summary of these results. Each of the four discount functions in Table 10.1, PPL, LP, RS (with $\rho \leq 1$ and $\tau\rho \leq 1$), and generalized RS, when combined with the SIE utility function (10.10), satisfies the following anomalies of the EDU model: The magnitude effect, the gain–loss asymmetry, common difference effect, delay–speedup asymmetry, and apparent intransitivity of preferences. If one is interested in incorporating departures from additivity, then Proposition 10.15 gives the relevant conditions for the delay and speedup functions. As an example, recall Proposition 10.8. Whether preferences exhibit stationarity or the common difference effect is determined by Ψ . Whether impatience is declining, constant, or increasing is determined by f . Ψ and f jointly determine whether preferences are subadditive, additive, or superadditive.

10.5.6 Intransitive preferences and relative discounting

Ok and Masatlioglu (2007) develop a theory of intransitive time preferences. At time 0, a decision maker has a binary relationship, $\preceq$, on the set $\chi = X \times [0, \infty)$, where X is a non-empty set. Let $x, y \in X$ and $s, t \in [0, \infty)$, then $(x, s) \preceq (y, t)$ is to be interpreted in the following manner. If the individual is asked at time 0 to commit to a choice between (x, s) and (y, t) he states that “ y received at time t is (weakly) preferred to x received at time s .” For this reason, $\preceq$ may be described as representing the *commitment preferences* of the individual.

Let $\sim$ and $<$ be the symmetric and asymmetric parts of $\preceq$, respectively. For each $t \in [0, \infty)$, $\preceq_t$ is the t -th time projection of $\preceq$ onto X , i.e., $x \preceq_t y$, if, and only if, $(x, t) \preceq (y, t)$. In particular, $\preceq_0$ is the projection of $\preceq$ onto X at time 0 (and, similarly, for $\sim_t$ and $<_t$). Thus, $\preceq_0$ describes the undated preferences of the individual in the present, i.e., time 0.

If X is a metric space, then further structure can be imposed on $\preceq$.

Definition 10.15 (*Time preferences; Ok and Masatlioglu, 2007, p. 218*): Let X be a metric space, then $\preceq$ is a time preference on χ if

- (i) $\preceq$ is complete and continuous,
- (ii) $\preceq_0$ is complete and transitive,
- (iii) $\preceq_t = \preceq_0$ for each $t \in [0, \infty)$.

In Definition 10.15, note that transitivity is imposed on the undated present preferences $\preceq_0$ (and, hence, also on $\preceq_t$), but transitivity is not imposed on the commitment preferences $\preceq$. In other words, static preferences $(\preceq_t, \preceq_0)$ are always transitive, but intransitivity may arise with the passage of time. Hence, neither $\prec$ nor $\sim$ are, necessarily, transitive. In particular, $\sim$ is *not*, in general, an equivalence relationship.

Example 10.16 Suppose that $x, y, z \in X$. Then Definition 10.15(ii) implies that if $(x, 0) \preceq_0 (y, 0)$ and $(y, 0) \preceq_0 (z, 0)$, then, by transitivity, we get $(x, 0) \preceq_0 (z, 0)$. Further, from Definition 10.15(iii) we get that for any $t \in [0, \infty)$ the t -th time projection of $\preceq$ onto X satisfies: if $(x, t) \preceq_t (y, t)$ and $(y, t) \preceq_t (z, t)$, then, $(x, t) \preceq_t (z, t)$. However, from Definition 10.15(i) we may have for $s, t, u \in [0, \infty)$: $(x, s) \preceq (y, t)$ and $(y, t) \preceq (z, u)$, but $(z, u) \not\preceq (x, s)$, which violates transitivity.

OM introduce a set of six properties, P1–P6. We only give a heuristic discussion of these properties because our focus is on the representation theorem that follows from them. *Time sensitivity* (P1) requires that any outcome can be made less desirable than another outcome by postponing it for a sufficiently long period. *Outcome sensitivity* (P2) requires that any delay can be compensated by a high enough prize. *Monotonicity* (P3) requires that $\preceq$ be increasing in outcomes and decreasing in time. Suppose that $(x, 0) \sim (y, 1)$ and $(z, 0) \sim (w, 1)$, then define the *delay premium* in these respective cases by $y - x$ and $w - z$. *Separability* (P4) requires that if the delay premium for delaying a reward x at any time $t \in [0, \infty)$ for an time interval of length, l , is $y - x$, then over the same interval, $w - z$ is the delay premium for delaying a reward of z at time t . *Path independence* (P5) is less persuasive as compared to the other properties. It requires that for any outcomes $x, y, w, z \in X$ and times $t_1, t_2, t_3 \in [0, \infty)$ if $(x, t_1) \sim (y, t_2) \sim (w, t_3)$, then $(x, t_2) \sim (z, t_3)$ where $(z, t_1) \sim (w, t_2)$. Finally *monotonicity in outcomes* (P6) requires that someone who uses the preference $\preceq_0$ prefers a higher to a lower amount.

The next proposition studies the implications of these properties. Let $\mathbb{R}$ be the set of real numbers, $\mathbb{R}_+$ the set of non-negative reals and $\mathbb{R}_{++}$ the set of positive reals. Recall that a homeomorphism is a mapping that is 1-1, onto, continuous, and its inverse is also continuous.

Proposition 10.16 (Ok and Masatlioglu (2007), Theorem 1): Let X be an open interval in $\mathbb{R}$ and $\preceq$ a binary relation on X . $\preceq$ is a time preference (Definition 10.15) on X that satisfies properties P1–P6 if, and only if, there exists an increasing homeomorphism $U : X \rightarrow \mathbb{R}_{++}$ and a continuous map $D : \mathbb{R}_+^2 \rightarrow \mathbb{R}_{++}$ such that, for all $x, y \in X$ and $s, t \in [0, \infty)$,

$$(x, s) \preceq (y, t) \text{ iff } U(x) \leq U(y) D(s, t), \quad (10.88)$$

where (i) for given s , $D(s, t)$ is decreasing in t with $D(s, \infty) = 0$, and (ii) $D(t, s) = 1/D(s, t)$.

Suppose $s \leq t$. Then (10.88) says that y received at time t is (weakly) preferred to x received at time s if, and only if, the (undated) utility of x is less or equal to the (undated) utility of y discounted from time t back to time s by the discount factor $D(s, t)$. In this case, part (i) of Proposition 10.16 implies the following. Fix the time, s , at which x is received. Let the time, t , at which y is received, recede into the future. Then the value of the utility of y , discounted back to time s , decreases. In the limit, as the receipt of y is indefinitely postponed, the value of its utility,

discounted back to time s , approaches zero. Part (ii) of Proposition 10.16 says that compounding forward, from time s to time t , is the inverse of discounting backwards from time t to time s .

For each $r \in [0, \infty)$, let $\preceq^r$ be the restriction of $\preceq$ to $X \times [r, \infty)$, i.e., to times $t \geq r$. Thus, for $r \leq s$ and $r \leq t$, $(x, s) \preceq^r (y, t)$ if, and only if, $(x, s) \preceq (y, t)$. We can now point to the main differences between RT and the theory proposed by Ok and Masatlioglu.

First, note that U in Proposition 10.16 can take only positive values while v in the RT theory, which uses prospect theory, takes both positive and negative values.²⁷ To bypass this problem, we consider only the domain of strictly positive gains. Let w_0 be the reference point for wealth. Take $X = \{w - w_0 : w > w_0\} = (0, \infty)$ and let $\preceq$ satisfy the conditions of Proposition 10.16. Let (U, D) be the representation of $\preceq$ guaranteed by Proposition 10.16, where U is the utility function and D is the interval discount function.

From Section 10.5.1 recall that, under RT, for each $r \in [0, \infty)$, $\preceq_{w_0, r}$ is a complete transitive order on $(-\infty, \infty) \times [r, \infty)$ and, hence, also on $X \times [r, \infty)$. Thus, in general, $\preceq_{w_0, r}$, unlike $\preceq^r$, is *not* the restriction to $X \times [r, \infty)$ of some complete binary relationship on $X \times [0, \infty)$. Thus the proposal by Ok and Masatlioglu and RT are not compatible and neither is a special case of the other.

Third, $\preceq_{w_0, r}$ is *transitive* while, in general, $\preceq^r$ is *not transitive*. To elaborate this point, consider (x, r) , (y, s) , and (z, t) , where $x, y, z \in X$ and $s, t \in [r, \infty)$. Suppose $(x, r) \preceq_{w_0, r} (y, s)$ and $(y, s) \preceq_{w_0, r} (z, t)$. Since $\preceq_{w_0, r}$ is transitive, we can conclude that $(x, r) \preceq_{w_0, r} (z, t)$. Now, suppose that $(x, r) \preceq^r (y, s)$ and $(y, s) \preceq^r (z, t)$. Since $\preceq^r$ is not, in general, transitive, we *cannot* conclude that $(x, r) \preceq^r (z, t)$.²⁸ More generally, given a compact subset $C \subset X \times [r, \infty)$, there is no guarantee in Ok and Masatlioglu that it has a maximum under $\preceq^r$ (i.e., an $m \in C$ such that $c \preceq^r m$ for all $c \in C$). This, obviously, will cause great difficulty for any economic theory formulated in the OM framework. On the other hand, in RT, if D is continuous, C will always have a maximum under $\preceq_{w_0, r}$.

Fourth, and finally, these problems with OM can all be resolved in the special case where $\preceq$ is transitive. But then $\preceq$ would also be additive. In this case, OM would reduce to the standard discounting model.

²⁷ Hence, in its present formulation, the proposal by Ok and Masatlioglu cannot explain *gain-loss asymmetry*. However, it can explain the *magnitude effect* using the SIE value function (see Proposition 10.4).

²⁸ As Ok and Masatlioglu clearly explain, it is for this reason that we should think of $D(s, t)$ in their theory as the *relative* discount function between times s and t .

CHAPTER 11

Applications of Present-Biased Preferences

11.1 Introduction

In the two previous chapters we saw that humans have a much greater bias towards the present relative to traditional models of discounting, such as EDU. In particular, most evidence indicates that the discount rate increases (equivalently impatience increases) as the delay with which an outcome is received, is shortened. Models where the degree of impatience increases, hyperbolically, or otherwise, when the delay of an outcome falls, are termed as models of *present-biased preferences*. This chapter considers applications of such models with particular emphasis on models of quasi-hyperbolic discounting.¹ The applications that we consider range from issues of optimal lifetime consumption, optimal retirement decisions, procrastination, choice between alternative gym memberships, and optimal taxes on activities that cause harm to the individual (sin taxes).

In applications of present-biased preferences, it is typical to consider a decision maker as a succession of *multiple selves*, one for each time period. Suppose that all selves exhibit hyperbolic discounting. This leads to a *taste for immediate gratification*, so it is not surprising that it may cause *self-control problems*. The first feature of models of multiple selves is that current selves have imperfect control over future selves. Hence, future selves may take actions that are contrary to the current self's own interest.

As an example, a future self may splurge and undersave because of present-biased preferences. This leads to a smaller pension pot, which causes concern to the current self. However, the current self may take actions that indirectly influence the future selves. For instance, the current self may invest in illiquid assets, which constrains the ability of future selves to splurge on consumption. A second feature of such models is that there might be *limited awareness* on the part of the current selves about the self-control problems of their future selves.

In Section 11.2, we consider two classes of models. In the *planner–doer model*, a long-lived current planner-self of an individual takes actions to influence the actions of future, short-lived, myopic-selves (Thaler and Shefrin, 1981). We show that this model can predict inadequacy of savings, as well as variation in the marginal propensity to consume out of permanent income. The second class of models considers the basic machinery of the models of multiple selves, by

¹ We do not delve into the reasons why humans exhibit hyperbolic discounting. For evolutionary arguments that justify hyperbolic discounting, see Dasgupta and Maskin (2005).

distinguishing among various assumptions about the degree of self-awareness of future self-control problems (O'Donoghue and Rabin, 2001).

Section 11.3 applies hyperbolic discounting and alternative assumptions about the degree of awareness of future self-control problems to a standard life-cycle model (Laibson, 1997, 1998). The present bias of individuals leads to greater current consumption and lower savings, relative to the benchmark case of time consistent behavior under the EDU model. Such a model can explain a range of puzzles that are difficult to explain in the standard model. These include the following. Why does consumption track income so closely? Why do individuals undersave for retirement? Why is there a sharp drop in consumption at retirement? Why do individuals hold illiquid assets and credit card debt simultaneously? Why have national savings rates been declining? Why is marginal propensity to consume asset specific? We also derive the Euler equation under the behavioral life-cycle model (*behavioral Euler equation*), show why high consumption variability is not associated with high consumption growth, and clarify the relation between intertemporal substitution and risk.

Section 11.4 considers the retirement decision under present-biased preferences (Diamond and Köszegi, 2003). The present bias of future selves might induce them to retire too early. The current self might disagree with this decision because longer horizon discount rates under hyperbolic discounting are lower. However, the current self cannot *directly* influence the future selves. One *indirect* method of influencing the future selves may be to *strategically undersave* now and, so, force the future selves to work rather than retire. If such an approach does not work, then, the current self may have to engage in *resigned oversavings*, i.e., save “too much” for the subsequent selves. We also examine the effect of employment legislation, such as mandatory, or flexible retirement policies, on the savings and retirement decisions.

Why do we procrastinate on important tasks, or sometimes give in to temptation (O'Donoghue and Rabin, 1999a, 1999b, 2001)? We consider these issues in the framework of quasi-hyperbolic discounting and different degrees of awareness of future self-control problems.² Section 11.5 considers two generic cases. First, the case of *delayed rewards*, such as the decision to quit smoking, in which costs must be incurred now, but rewards accrue in the future. Second, the case of *delayed costs*, such as credit card debt, and vacation decisions, in which rewards are received now, but costs accrue in the future. The appropriate equilibrium concept here is one of a *perception-perfect strategy*. For immediate costs, lack of awareness of future self-control problems leads to greater current *procrastination*, because (i) present bias makes current costs more salient, relative to future gains, and (ii) lack of awareness of future self-control problems reduces the perceived costs of procrastination. For the case of immediate rewards, the converse result obtains.

Section 11.6 examines the effect of market competition on the propensity to procrastinate (Brocas and Carillo, 2001). Suppose, for instance, that there are two employees who have to complete a task and face a “winner takes all” incentive scheme. Procrastination by any of the employees is now doubly costly if the other employee does not procrastinate. This induces greater work effort from both employees. This is the case if costs of procrastination are low. However, for intermediate costs of procrastination, conditional on the beliefs of any employee about the other's procrastination, a multiplicity of equilibria can arise.

In standard public economics, *Pigouvian taxes* are often advocated in order to reduce the level of a harmful externality that one individual imposes on another. However, when individuals consume a *sin-good*, such as smoking, and use hyperbolic discounting, there are additional issues to

² An alternative approach to issues of self-control, the Gul-Pesendorfer approach to temptation preferences, that uses the revealed preference framework, can be found in Part 6.

take into account (O'Donoghue and Rabin, 2006). In a multiple selves framework, the current self of an individual, on account of his present bias, may consume far in excess of what a patient long-lived self would advocate. Pigouvian taxes can then be levied on the current self of the individual, in order to protect the future selves. Section 11.7 explores these issues in a formal model.

In Section 11.8, we consider the take-up of alternative pricing plans for *investment goods* and *leisure goods* (DellaVigna and Malmendier, 2004). Investment goods have current costs and future benefits (e.g., health club memberships). Leisure goods, conversely, have current benefits, but future costs (e.g., credit card borrowing). For instance, consider a potential visitor to a health club who has hyperbolic preferences and is sophisticated about his future self-control problems. He will know that the future selves, on account of their self-control problems, will not use the club membership if the monthly fee is too high. Hence, he may be willing to pay a larger up-front lump-sum payment (membership fee) in return for a smaller unit price (such as the monthly fee) for later visits. Furthermore, a unit price below marginal cost is optimal for the health club. The evidence presented by DellaVigna and Malmendier (2004) confirms these results. By using a converse set of arguments, one can show that optimal price is above marginal costs for leisure goods. In each of these cases, the variation in the price, relative to the neoclassical model, allows individuals with self-control problems to gain access to a commitment device.

A range of other applications of hyperbolic discounting is briefly described in Section 11.9. This section highlights an expanding set of problems to which hyperbolic discounting can be applied in a tractable manner to address important issues.

11.2 Two frameworks for modeling present-biased preferences

This section describes two frameworks for modeling models of *multiple selves* in the presence of hyperbolic discounting. Consider an individual who lives through several time periods. Drawing on an influential idea from psychology, suppose that a *distinct self* of the same individual lives in any particular period. Hence, over time, there are *multiple selves* of the same individual.³ The individual in period t is fully described by *self* t of the individual.

In the first of these frameworks, considered in Section 11.2.1, Thaler and Shefrin (1981) envisage the current self as a *long-sighted planner* who cares about the utilities of all future selves, and the future selves as *myopic doers* who care only about their own utilities. The planner can take actions that directly constrain the actions of the doers, or even directly influence the utility of the doers. However, in any period t , other than the initial period in which the planner takes the action, it is the doer corresponding to period t that decides on the action to be taken. For instance, the planner-self of a smoking addict might commit, in period 0, to pay a non-refundable fee in return for enrollment at a smoking clinic in period 1. This constrains the choices of the myopic doer in period 1, although it is the doer who must make the smoking decision in period 1. The idea that future selves have a potential self-control problem that a current self might rectify, sometimes known as the *two-selves models*, can also be found in Elster (1979), Margolis (1982), and Schelling (1984).

The second of the two frameworks of multiple selves that we consider in Section 11.2.2 has proved to be very useful in applications of present-biased preferences. On account of the present

³ We consider further motivation for models of multiple selves and models of *dual selves* in Part 6 of the book on emotions.

bias of each self in any period, different selves typically put different weights on future outcomes. However, unlike planner–doer models, none of the selves can *directly* control their future selves. A major question in this literature, therefore, is the following. Can the current self take any actions that, *indirectly*, influences future selves to internalize better the preferences of the current self?

An important issue, which we examine in Section 11.2.2, is the *degree of awareness* that a current self has about the present bias of one's own future selves. While experimental evidence is inconclusive, it appears that completely *sophisticated behavior*, i.e., full awareness of future self-control problems, and the polar extreme of completely *naive behavior*, are less likely relative to the intermediate case of *partial awareness* of self-control problems in the future.

11.2.1 Planner–doer models

Within a life-cycle model, suppose that the current self takes the form of a *planner*, and cares for the utility of all the subsequent selves. The subsequent selves, one for each time period, termed as the *doers*, are myopic in the sense that they care only about maximizing their own utility. Thaler and Shefrin (1981) term models with these features as *Behavioral Life-Cycle Models* (BLC) that we now discuss. We present a modified version of the model of Thaler and Shefrin (1981) in which there is strategic substitutability between the actions of the planner and the doers. There is no presumption that this modification would be agreed to by the authors of the BLC model, but it allows for an elegant and simple illustration of the model.

The planner–doer framework can be used to describe several types of behaviors in which the current planner may take potentially costly actions to bind the actions of future myopic doers. Consider the following examples.

1. In the case of alcohol addiction, the planner may take a drug (e.g., antabuse) that makes the future selves ill if they drink alcohol.
2. The planner-self of an academic might agree to present a paper at a conference at a future date, which binds the behavior of future selves.
3. The planner-self in a life-cycle model might inculcate a *debt ethic* (no borrowing under most circumstances) that gradually becomes a habit and, in models of habit formation, constrain the future selves.

Consider a three-period model; the time periods are $t = 0, 1, 2$. In the initial period, period 0, the planner chooses some action $\theta \in [0, \bar{\theta}] \subset R$ that influences the utility of the subsequent selves. The interpretation of $\theta = 0$ is that the planner takes no action, and $\bar{\theta}$ is the maximum action that can be taken. The cost of taking action θ to the planner is $\phi(\theta)$, which is increasing and convex: $\phi_1 > 0$ and $\phi_{11} \geq 0$, where subscripts denote partial derivatives.

Call the self in period $j = 1, 2$ as self j (or doer j). For pedagogical simplicity, the parameter, θ , chosen by the planner affects only self 1, but not self 2. The individual works in the second period, inelastically supplying one unit of labor, and earning labor income, y . The individual is retired at the end of period 1, and earns no income in period 2. Hence, the income profile over the three periods is $(0, y, 0)$. Consumption in period $j = 1, 2$ is denoted by c_j .

The planner's problem is to choose θ in order to maximize the sum of utilities of selves 1 and 2, giving equal weights to both. Self 2 is retired, has no economically meaningful decision to make, and is unaffected by the planner's action, θ . So, for pedagogical simplicity, we assume that period 2 utility is given by

$$U_2(c_2) = c_2.$$

The utility function of self 1 is given by

$$U_1(s, \theta) = u(c_1, \theta) + \eta U_2(c_2), \quad 0 < \eta < 1. \quad (11.1)$$

So, as compared to the planner who puts equal weights on the utilities of the two future selves, self 1 is myopic in the sense that he puts relatively greater weight on his own utility. Self 1 is also influenced by the planner's choice of θ . The own utility function of self 1, u , has continuous partial derivatives of order 2. Other properties of u are (subscripts on u denote partial derivatives):

$$u_1 > 0; u_{11} < 0; u_2 < 0; u_{12} < 0. \quad (11.2)$$

Hence, u is increasing in consumption, but at a diminishing rate. By choosing a higher action, θ , the planner can reduce the utility of self 1, i.e., $u_2 < 0$. Furthermore, there is *strategic substitutability* between the action chosen by the planner and the consumption of self 1, i.e., $u_{12} < 0$. Hence, the planner has the ability to alter the level of utility and the marginal utility of self 1.

The next set of assumptions ensure an interior solution.

$$(i) \eta < u_1(y, 0). (ii) \forall \theta > 0: (a) \eta > u_1(y, \theta), \text{ and } (b) \eta > u_1(0, \theta). \quad (11.3)$$

The budget constraint of self 1 is given by

$$c_1 + s = y, \quad (11.4)$$

where $s \in [0, y]$, is the savings made by self 1.

The solution to the problem is by backward induction. Self 2 has no economically meaningful decision to make and relies entirely on the savings of self 1, so $U_2(c_2) = s$. Substituting this and (11.4) in (11.1), the unconstrained maximization problem of self 1 is to choose s , given θ , i.e.,

$$s^* = \arg \max_{s \in [0, y]} U_1(s, \theta) = u(y - s, \theta) + \eta s.$$

It follows that

$$\frac{dU_1(s, \theta)}{ds} = -u_1(y - s, \theta) + \eta. \quad (11.5)$$

Since u is a continuous and strictly concave function over a compact set, we have a unique solution, $s^*(\theta)$, to the problem of self 1. Substituting s^* in (11.4), the optimal consumption choice is $c_1^*(\theta) = y - s^*(\theta)$.

If the planner undertakes no costly action, $\theta = 0$, then, using the first inequality in (11.3), and (11.5), it is optimal for self 1 to consume all of the current income, because $\frac{dU_1(s, \theta)}{ds} < 0$. Thus, in the absence of any intervention from the planner, $s^*(0) = 0$, and self 2 has zero income. This is the sense in which self 1 is *present-biased*.

However, for any $\theta > 0$, (11.3) and (11.5) imply that $s^* > 0$. The last inequality in (11.3), and (11.5) imply a unique interior optimal solution $0 < s^*(\theta) < y$. The optimal consumption $c_1^*(\theta) = y - s^*(\theta)$, is also interior so implicitly differentiating (11.5) at equality, we get

$$\frac{dc_1^*}{d\theta} = -\frac{ds^*}{d\theta} = \frac{u_{12}}{-u_{11}} < 0. \quad (11.6)$$

The sign in (11.6) follows because of strategic substitutability, $u_{12} < 0$. Thus, an increase in the planner's action increases savings made by self 1.

The planner's objective is to maximize the sum of utilities of selves 1 and 2, hence, anticipating correctly the subsequent choices, the unconstrained problem of the planner is,

$$\theta^* \in \arg \max U_0(\theta) = u(y - s^*(\theta), \theta) + \eta s^*(\theta) + s^*(\theta) - \phi(\theta); \theta \in [0, \bar{\theta}].$$

Using the envelope theorem, we can write

$$\frac{dU_0(\theta)}{d\theta} = U'_0 = u_2(y - s^*(\theta), \theta) + \frac{ds^*(\theta)}{d\theta} - \frac{d\phi(\theta)}{d\theta}. \quad (11.7)$$

U_0 is a continuous function on a compact set, $[0, \bar{\theta}]$, so it has a maximum. The maximum is unique if $U'_0 < 0$. Given (11.2) and (11.6), the first term on the RHS of (11.7) is negative and the second is positive. These terms capture, respectively, the marginal cost arising from the reduction in consumption of self 1, and the marginal benefit arising from an increase in the consumption of self 2. The third term is the marginal cost to the planner of using the action, θ . These terms reflect the basic trade-off faced by the planner. Let us now consider some predictions of the planner–doer models.

1. **Savings inadequacy:** Consider the case where self 1 has to subscribe to some mandatory retirement savings plan. In the life-cycle model, individuals smooth income perfectly over their lifetime, so issues of savings inadequacy arising from self-control problems are absent. However, in the planner–doer framework, it is natural to ask if the actions taken by the planner ensure that the future selves make sufficient savings. The answer depends on the nature of marginal benefits and costs shown in (11.7). Kotlikoff et al. (1982) find the answer in the negative. Diamond and Hausman (1984) also find that savings are inadequate. Hence, in this model, mandatory pension, and social security contributions, seem empirically important to supplement the effort of the planner to control the doers.
2. **Non-proportionality:** In Friedman's permanent income hypothesis, individuals consume a fixed fraction of their permanent income. This fraction does not vary with the magnitude of permanent income. Friedman called this the *proportionality principle*. Shefrin and Thaler (1988) report evidence against the proportionality principle; see also, Diamond and Hausman (1984). In the planner–doer framework, there is no presumption that the proportionality principle holds because each of the doers is myopic and does not smooth consumption as well as is required in the neoclassical model. Further, the actions of the doers are controlled imperfectly by a planner through a costly choice of θ that potentially reflects willpower (Shefrin and Thaler, 1988). Thus, individuals at different levels of permanent income (y in the model above) consume different fractions of it.

The planner–doer model has played a lesser role in the subsequent literature in economics relative to the model of multiple selves that we consider next. For this reason, we do not consider the planner–doer model any further; this should not be taken, in any way, as an adverse judgment on the model.

11.2.2 Multiple selves and the degree of awareness

We now describe the workhorse model of *multiple selves* in the presence of quasi-hyperbolic discounting. This section outlines the basics of this framework, followed by a range of applications in subsequent sections. Suppose that time $t \in \Gamma = \{0, 1, 2, \dots, T\}$. Let there be multiple selves of the same individual—one for each period. Each self who lives in any time period $t \in \Gamma$ has quasi-hyperbolic preferences given by.

$$U_t(c_t, c_{t+1}, \dots, c_T) = u(c_t) + \beta \sum_{\tau=t+1}^{\tau=T} \delta^\tau u(c_\tau), 0 < \beta < 1. \quad (11.8)$$

The parameter β in (11.8) creates a bias towards immediate gratification for each self, and $\delta > 0$ is the discount factor in the EDU model. In making economic decisions, the self at time $t \in \Gamma$ must form beliefs about the behavior of all the future selves. Such beliefs reflect the *degree of self-awareness* of the current self at any time $t \in \Gamma$. Self t has beliefs that the preferences of all the future selves at time $k > t$ are given by

$$U_k(c_k, \dots, c_T) = u(c_k) + \hat{\beta} \sum_{\tau=k+1}^{\tau=T} \delta^\tau u(c_\tau), t < k \leq T - 1. \quad (11.9)$$

The relation between the actual impatience parameter of self t , β , and his beliefs about the impatience parameter of a future self, $\hat{\beta}$, is crucial for current decisions. In a more general model, one's beliefs about the present bias of all the future selves may be heterogeneous. Thus, there could be a sequence of $\hat{\beta}_k$ values $t < k \leq T - 1$, but we simplify by setting $\hat{\beta}_k = \hat{\beta}$ for all k . O'Donoghue and Rabin (2001) propose the following classification.

1. *Time consistent*: People with standard time consistent preferences behave as if they are using the EDU model, so

$$\beta = \hat{\beta} = 1. \quad (11.10)$$

These individuals imagine that none of their future selves ever has a present bias (over and above that caused by the presence of δ).

2. *Sophisticates*: Sophisticates know that they have a bias towards immediate gratification, and believe that their future selves will also have an identical bias, i.e.,

$$\beta = \hat{\beta} < 1. \quad (11.11)$$

3. *Partial naifs*: Partial naifs know that they have a bias towards immediate gratification, but believe (mistakenly) that their future selves will have a smaller, but non zero, bias, i.e.,

$$\beta < \hat{\beta} < 1. \quad (11.12)$$

4. *Naifs*: Naifs know that their current selves are present-biased, but believe that their future selves will have no bias towards immediate gratification, i.e., they believe that future selves are time consistent, hence

$$\beta < \hat{\beta} = 1. \quad (11.13)$$

The existing evidence is not conclusive, but O'Donoghue and Rabin (2001) feel that people lie in between the two extremes of sophisticates and naifs, i.e., most people are probably partial naifs. This suggests checking for the robustness of results with respect to this proposed classification.

Example 11.17 (*The snooze button in alarm clocks*): Many of us set the alarm to wake up in the morning. Let us term the self who sets the alarm at night, the night-self, and the self that wakes up in the morning, the morning-self. We may often find it difficult to precommit to waking up on the alarm, because the night-self cannot predict fully how the morning-self will react to the alarm. It is possible for the morning-self to just switch the alarm off, with possibly negative consequences for both selves. In this case, the night-self might desire some commitment technology that disciplines the morning-self.

The makers of alarm clocks seem aware of the problem, so they supply us with alarm clocks that have a snooze button that reactivates the alarm to set-off after a short interval. This acts as a disciplining device on the morning-self, and also provides the night-self with some degree of precommitment.

An even stronger precommitment from the night-self would be to place the alarm clock away from the bed, increasing the cost of the morning-self to reverse the night-self's decision. Perhaps the night-self does not often engage in this strong precommitment because it acts as a partial naif, and underestimates the self-control problems that the morning-self will have.

11.3 Optimal consumption under hyperbolic discounting

Consider a simple three period life-cycle model of consumption with multiple selves and quasi-hyperbolic discounting. We show that in the presence of hyperbolic discounting, there are lower savings due to the present bias of individuals. The results on the relation between the degree of awareness (naif, partial naif, and sophisticates), and the shape of the consumption profile, are interesting, but are dependent on the functional forms used; the Euler equations corresponding to the more general life-cycle model are considered below in Section 11.3.4. A more direct aim of the simple model is to expose the reader to the method of solution under hyperbolic discounting and limited awareness.

These sorts of models have provided an important starting point for the emerging and exciting new discipline of *behavioral macroeconomics*. For a fuller treatment in a more general model, the reader may wish to consult Laibson et al. (1998) and Angeletos et al. (2001). A calibrated economy in which consumers use hyperbolic discounting explains much better the data on consumption, illiquid assets, and credit card debt, relative to a model where consumers use exponential discounting. In a follow-up paper, Laibson et al. (2005) estimate the quasi-hyperbolic parameters (on an annual basis) to be: $\beta \approx 0.7$ and $\delta \approx 0.96$. For a more general model with infinite horizon and uncertainty, Harris and Laibson (2001a,b) show that the derived consumption functions may exhibit discontinuities and non-monotonicity. However, when the model is calibrated at empirically plausible parameter values, such pathologies do not arise. Judd (2003) looks into further issues using asymptotic theory and perturbation methods to examine the existence and uniqueness of the solution under hyperbolic discounting.

Let the time periods be 1, 2, and 3, so $\Gamma = \{1, 2, 3\}$. The preference of self t , $t \in \Gamma$, are given in (11.8). We set $\delta = 1$; this is also known as $(\beta, 1)$ preferences. Beliefs about the future selves

are given in (11.9). Self 1 has exogenously given labor income $y > 0$, but subsequent selves have no labor income; this can be relaxed, but it is not essential for the results. Consumption in the three periods is denoted by c_1, c_2, c_3 . We assume that the individual can only save in one-period financial instruments, such as one-period bonds. For simplicity, the interest rate is normalized to zero. We relax some of these assumptions below. For instance, we will consider the possibility of saving in assets with a longer maturity, and the use of illiquid assets.

The preferences of self 1 are given by

$$U_1 = u(c_1) + \beta u(c_2) + \beta u(c_3), \quad (11.14)$$

where u is a strictly increasing and concave utility function. The problem of self 1 is to choose savings in the first period, $s_1 \in [0, y]$. For pedagogical ease, we focus only on interior solutions. The budget constraints in each time period are given by:

$$\begin{cases} c_1 = y - s_1, & s_1 \in [0, y], \\ c_2 = s_1 - s_2, & s_2 \in [0, s_1], \\ c_3 = s_2, \end{cases} \quad (11.15)$$

where s_2 is the savings made by self 2, which we restrict to be non-negative. Except for the first period, consumption in any of the two remaining periods is only made possible by savings carried over from the previous period in the form of a one-period financial instrument that bears zero rate of interest.

11.3.1 Benchmark case: time consistent preferences

In the classical case, decision makers follow the EDU model, which we know from Section 9.3.2 gives rise to time consistent choices. Denote the optimal values in this case by the superscript c (for the “consistent” solution). From (11.10), in this case, $\beta = \hat{\beta} = 1$. Substitute $\beta = 1$ and (11.15) in (11.14) to get the following unconstrained problem.

$$s_1^c, s_2^c \in \arg \max_{s_1 \in [0, y]; s_2 \in [0, s_1]} u(y - s_1) + u(s_1 - s_2) + u(s_2).$$

The first order conditions at the interior solution are

$$s_1 : u'(c_1) = u'(c_2), \quad (11.16)$$

$$s_2 : u'(c_2) = u'(c_3). \quad (11.17)$$

From (11.16), (11.17), we get that along the optimal consumption path, since the interest rate is zero and $\delta = 1$, marginal utilities are equated. Thus, $u'(c_1) = u'(c_2) = u'(c_3)$, which in conjunction with the budget constraints in (11.15), implies that the optimal solution is given by

$$c_1^c = c_2^c = c_3^c = \frac{y}{3}. \quad (11.18)$$

Hence, an individual with time consistent preferences perfectly smooths consumption across the three periods.

11.3.2 The solution under limited awareness

Consider now any of the cases of different degrees of awareness described in Section 11.2.2. When maximizing (11.14), subject to the budget constraint in (11.15), self 1 knows that self 2 will choose s_2 . Self 1 holds the belief that the impatience parameter of self 2 is $\hat{\beta}$. The relation between β and $\hat{\beta}$ is given by (11.11), (11.12), and (11.13), respectively, for *sophisticates*, *partial naifs*, and *naifs*. Let superscript $j = s, p, n$ on the optimal solution denote, respectively, sophisticates, partial-naifs, and naifs. Hence, self 1 believes that the preferences of self 2 are given by

$$U_2 = u(c_2) + \hat{\beta}u(c_3). \quad (11.19)$$

Substituting (11.15) in (11.19), self 1 believes that the unconstrained problem of self 2 is given by

$$s_2^j \in \arg \max_{0 \leq s_2 \leq s_1} u(s_1 - s_2) + \hat{\beta}u(s_2); j = s, p, n. \quad (11.20)$$

In (11.20), self 2 takes as given, the choice of the first period saving, s_1 , made by self 1. The first order condition to the problem in (11.20) is given by

$$u'(c_2) = \hat{\beta}u'(c_3). \quad (11.21)$$

The optimal solution to (11.21) is given by $s_2^j(s_1)$. The higher is $\hat{\beta}$, the greater is the weight placed on the utility of self 3, hence, the greater is s_2^j .

Self 1 now maximizes (11.14), given the budget constraints in (11.15), and the subsequent behavior of self 2, encapsulated in the reaction function $s_2^j(s_1)$. Substituting these constraints into (11.14), the unconstrained problem of self 1 is given by

$$s_1^j \in \arg \max_{s_1 \in [0, y]} u(y - s_1) + \beta u(s_1 - s_2^j(s_1)) + \beta u(s_2^j(s_1)); j = s, p, n.$$

The first order condition is given by

$$-u'(y - s_1) + \beta u'(s_1 - s_2^j(s_1)) \left(1 - \frac{ds_2^j(s_1)}{ds_1}\right) + \beta u'(s_2^j(s_1)) \frac{ds_2^j(s_1)}{ds_1} = 0. \quad (11.22)$$

In saving an extra dollar, self 1 faces a reduction in own utility (first term), but a benefit in terms of an increased utility for self 3 (third term). The marginal effect on self 2 depends on whether s_1 offsets s_2^j more than, or less than one for one.

In order to obtain analytical results, let us assume that the instantaneous utility function is of the CRRA type, i.e.,

$$u(c) = \left(\frac{1}{1-\gamma}\right) c^{1-\gamma}; 0 < \gamma \text{ and } \gamma \neq 1. \quad (11.23)$$

Pollak (1968) showed that the degree of awareness has no effect in the case of log preferences, hence, we restrict (11.23) to the case $\gamma \neq 1$. For CRRA preferences, the first order condition of self 2, (11.21), simplifies to

$$s_2^j(s_1) = \widehat{\mu}(s_1 - s_2^j(s_1)), \text{ where } \widehat{\mu} = \widehat{\beta}^{\frac{1}{\gamma}}. \quad (11.24)$$

Hence,

$$s_2^j(s_1) = \frac{\widehat{\mu}}{1 + \widehat{\mu}} s_1. \quad (11.25)$$

We can now use (11.25) in (11.22) to find the optimal savings made by self 1, s_1^j , $j = s, p, n$. These values can then be substituted in the budget constraints in (11.15) to determine the optimal values of consumption, c_1^j, c_2^j, c_3^j , $j = s, p, n$. The next proposition states the results of this simple exercise.

Proposition 11.17 *Suppose that preferences take the CARA form. Let $\mu = \beta^{1/\gamma}$ and $\widehat{\mu} = \widehat{\beta}^{\frac{1}{\gamma}}$. Then the following solution arises in the three-period life-cycle model.*

(a) *When self 1 is a naif, so $\beta < \widehat{\beta} = 1$ (see (11.13)), then*

$$c_1^n = (1 - \theta_1) y; c_2^n = c_3^n = \frac{\theta_1}{2} y, \text{ where } \theta_1 = \left(1 + \frac{1}{2\mu}\right)^{-1}. \quad (11.26)$$

(b) *When self 1 is a partial naif, so $\beta < \widehat{\beta} < 1$ (see (11.12)), then*

$$c_1^p = (1 - \theta_2) y; c_2^p = \left(\frac{\theta_2}{1 + \widehat{\mu}}\right) y; c_3^p = \left(\frac{\widehat{\mu}\theta_2}{1 + \widehat{\mu}}\right) y, \quad (11.27)$$

where

$$\theta_2 = \left(1 + \frac{1}{\mu} \left(\left(\frac{1}{1 + \widehat{\mu}}\right)^{1-\gamma} (1 + \widehat{\mu}^{1-\gamma})\right)^{-1/\gamma}\right)^{-1}.$$

(c) *When self 1 is a sophisticate, so $\beta = \widehat{\beta} < 1$ (see (11.11)), then*

$$c_1^s = (1 - \theta_3) y; c_2^s = \left(\frac{\theta_3}{1 + \mu}\right) y; c_3^s = \left(\frac{\theta_3\mu}{1 + \mu}\right) y, \quad (11.28)$$

where

$$\theta_3 = \left(1 + \frac{1}{\mu} \left(\left(\frac{1}{1 + \mu}\right)^{1-\gamma} (1 + \mu^{1-\gamma})\right)^{-1/\gamma}\right)^{-1}.$$

Comparing (11.18), (11.26), and using $\beta^{1/\gamma} < 1$, it follows that $c_1^n > c_1^c$. Hence, as one would expect, the relatively greater desire of the naif for immediate gratification as compared to a time consistent individual, leads to greater current consumption. Thus, a naif engages in a lower degree of consumption smoothing.

Table 11.1 Optimal consumption profiles for two values of the coefficient of relative risk aversion, 0.5 (columns 2–4) and 1.5 (columns 5–7).

Individual type	c_1	c_2	c_3	c_1	c_2	c_3
Time consistents	333.3	333.3	333.3	333.3	333.3	333.3
Sophisticates	731.9	197.1	71.0	499.6	292.4	208.0
Partial naifs	790.9	127.5	81.6	515.1	260.4	224.4
Naifs	581.4	307.8	110.81	412.8	343.14	244.1

Partial naifs are more aware than naifs that they will have future self-control problems. Does relatively greater self-awareness, as compared to the naifs, lead the partial naifs to consume less in the first period? Comparing (11.26), (11.27), the answer ultimately depends on the values taken by the parameters. For instance, when $\hat{\beta} = 0.8$ and $\gamma = 0.5$, a simple calculation shows that $c_1^p > c_1^n$. Hence, for these illustrative values, the greater awareness of partial naifs of the self-control problems of their future selves, reduces their current savings. This result is not robust to all utility functions and other applications below deal with this issue.

We now report some calibration results from O'Donoghue and Rabin (2003). Let us normalize $y = \$1000$. Table 11.1 reports illustrative results for first period consumption for various values of $\hat{\beta}$, and for different values of the coefficient of relative risk aversion, γ . In columns 2–4, $\gamma = 0.5$, and in columns 5–7 we have $\gamma = 1.5$. In general, the relation between the degree of self-awareness and over-consumption relative to the benchmark case of time consistent consumers, depends on the shape of the utility function. However, present bias, in general, leads to overconsumption because the first period consumption of the time consistents is lower relative to any of the other types with potential self-control problems.

11.3.3 Mitigating self-control problems in the life-cycle model

The most robust feature of Table 11.1 is the self-control problem (in terms of lower savings) faced by all three types of individuals who exhibit hyperbolic discounting, relative to individuals with time consistent preferences. Can individuals who have self-control problems take some actions to influence the choices of the future selves? One possibility, suggested by Angeletos et al. (2001), is investment in *illiquid assets*; we give below an analytical solution for a simpler model.

Suppose that, in the three period model considered above, self 1 can invest in some illiquid asset, z , that has a maturity of two periods, so it matures in period 3. The asset is illiquid in the sense that self 2 cannot sell it, nor borrow against it in the second period. Such an asset gives self 1 an opportunity to control the behavior of self 2. We continue to assume that utility takes the CRRA form in (11.23).

SOLUTION IN THE PRESENCE OF COMMITMENT

As a benchmark, we first derive the optimal consumption profile of self 1 under *commitment*. In the presence of commitment, the current self can directly specify the actions of future selves, and the specified actions are binding. Since direct control is possible, self 1 does not need to make a costly commitment decision using illiquid assets. Substituting (11.15) in (11.14), we get the following unconstrained problem of self 1.

$$s_1^*, s_2^* \in \arg \max_{s_1 \in [0, y]; s_2 \in [0, s_1]} u(y - s_1) + \beta u(s_1 - s_2) + \beta u(s_2).$$

On account of the commitment technology, self 1 can directly choose s_1, s_2 . A simple proof that is omitted gives the next proposition.

Proposition 11.18 *Let $\mu = \beta^{1/\gamma}$ and utility be of the CRRA form in (11.23). The optimal solution to the problem of self 1 under commitment is given by*

$$s_1^* = \left(1 + \frac{1}{2\mu}\right)^{-1} y. \quad (11.29)$$

$$s_2^* = \left(1 + \frac{1}{2\mu}\right)^{-1} \frac{y}{2}. \quad (11.30)$$

The corresponding optimal consumption profile is given by

$$c_1^* = y \left[1 - \left(1 + \frac{1}{2\mu}\right)^{-1}\right]. \quad (11.31)$$

$$c_2^* = c_3^* = s_2^* = \left(1 + \frac{1}{2\mu}\right)^{-1} \frac{y}{2}. \quad (11.32)$$

Despite the availability of commitment, an individual who has present-biased preferences does not perfectly smooth consumption, unlike a time consistent individual (see (11.18)).

ILLIQUID ASSETS AND ABSENCE OF COMMITMENT

Suppose that no commitment technology is available to self 1. We shall now show that self 1 can use a long-term, two-period, illiquid asset, z , to implement the consumption profile given in (11.31), (11.32). Investment in illiquid assets could take several forms, such as housing and defined pension contributions (for instance the 401k plan in the US). This alters the budget constraint of self 3, given in (11.15), to $c_3 = s_2 + z$. Since the problem of self 3 is trivial, we first solve the optimization problem of self 2 in the presence of illiquid assets.

Self 2 takes as given s_1, z , which are chosen by self 1, and solves

$$s_2^j \in \arg \max_{(0 \leq s_2 \leq s_1)} u(s_1 - s_2) + \widehat{\beta} u(s_2 + z); j = s, p, n. \quad (11.33)$$

The objective function in (11.33) is identical to (11.20), except that c_3 now includes the illiquid asset, z , that was invested in the first period, but matures in the third. Given the nature of the illiquid asset, the individual cannot borrow against z in the second period.⁴ A simple calculation shows that for the case of CRRA preferences, the solution is given by

$$s_2^j(s_1, z) = \max \left\{ 0, \frac{1}{1 + \widehat{\mu}} (s_1 - \widehat{\mu} z) \right\}, \text{ where } \widehat{\mu} = \widehat{\beta}^{\frac{1}{\gamma}}. \quad (11.34)$$

Given the subsequent behavior of self 2, as captured in (11.34), self 1 now takes two main financial decisions. How much to invest in short-term savings, s_1 , and how much to invest in the long-term

⁴ If, for instance, self 2 could cash a fraction λ of the asset, then his objective function would be $u(s_1 + \lambda z - s_2) + \widehat{\beta} u(s_2 + (1 - \lambda)z)$.

illiquid asset, z ? Lacking a commitment technology, self 1 cannot directly control the behavior of self 2, but he can do so indirectly by a choice of z . We now show how the optimal consumption profile given in (11.31), (11.32) can be implemented; the solution is derived recursively.

Suppose that the solution to (11.34) is $s_2^j = 0$. Hence, self 2 is unable to alter the consumption profile chosen by self 1. This implies that $c_3 \equiv z$. To ensure that $c_3 = c_3^*$, choose

$$z = c_3^* = \left(1 + \frac{1}{2\mu}\right)^{-1} \frac{y}{2}. \quad (11.35)$$

Since $s_2^j = 0$, we have that $c_2 \equiv s_1$. To ensure that $c_2 = c_2^*$, choose

$$s_1 = c_2^* = \left(1 + \frac{1}{2\mu}\right)^{-1} \frac{y}{2}. \quad (11.36)$$

The choice of c_2^*, c_3^* ensures that $c_1 = c_1^*$ will satisfy the intertemporal budget constraint. Hence, in the presence of illiquid assets, self 1 can implement the commitment solution. We leave it to the reader to show that this result is unchanged if the solution to (11.34) is $s_2^j > 0$. With a relatively high value of $\gamma = 5$ and $y = \$1000$, $\beta = 0.6$, $\hat{\beta} = 0.8$, it turns out that the optimal consumption levels are $c_1 = 388.8$ and $c_2 = c_3 = 305.6$, which is closer to the consumption profile of an individual with time consistent preferences (see Table 11.1).

We now use the model above to illustrate how one might explain a range of puzzles in macroeconomics.

1. **Why does consumption track income so closely?** In our stylized three period model, the individual has income only during the first period, hence, the income profile is $(y, 0, 0)$. From Table 11.1, the time consistent individual engages in perfect consumption smoothing over the three periods. However, for any of the other three types of present-biased individuals, consumption tracks income more closely relative to the time consistent individual. This feature survives modeling over a longer horizon where income is received over several periods.
 2. **Why do individuals undersave for retirement?** Time consistent individuals do not undersave for retirement. However, the retirement plans of any of the other types can be scuppered by the present bias of one of the future selves who may splurge on current consumption. For instance, under commitment, (11.30) gives desired second period savings, s_2^* , which, from (11.32), equal desired third period consumption c_3^* . However, in the absence of commitment and illiquid assets, second period saving, $s_2^j(s_1) = c_3^j$ given in (11.25), is decided by self 2. Given the present bias of self 2, c_3^j falls short of c_3^* (see Table 11.1). So, in general, individuals with a present bias will not save adequately for their retirement. These issues are considered in a richer framework in Section 11.4 below.
- Brown et al. (2009) explore two possible reasons for undersavings. (i) Bounded rationality (individuals simply lack the necessary sophistication to formulate optimal dynamic plans), and (ii) present-biased preferences. They find that while initially subjects' decisions are not close to optimal decisions due to bounded rationality, individuals learn quickly. However, when *visceral influences* are introduced (in this case, thirst), the behavior of individuals is

consistent with hyperbolic discounting.⁵ The parameter β in hyperbolic discounting models is estimated to be 0.6–0.7.

3. **Why is there a sharp drop in consumption at retirement?** In our three-period model, with the income profile $(y, 0, 0)$, the individual can be interpreted as working for the first period of his life and retired for the next two.⁶ It is apparent from Table 11.1 that for an individual with time consistent preferences, there is no drop in consumption at retirement. However, for all other types of individuals, there is a drop in consumption from the first period onwards. Evidence suggests that the drop in consumption at retirement can be fairly significant. Based on US data, Bernheim et al., (1997) report that the bottom wealth quartile suffers a 30% drop in consumption at retirement. Kotlikoff et al. (1982) find that in the absence of social security and private pensions, consumption in old age would, on average, be about 40% lower.
4. **Why do individuals hold illiquid assets and credit card debt simultaneously?** Illiquid asset can discipline one's future selves only if the future selves find it difficult to borrow against future wealth. Using field data from the Philippines, Ashraf et al. (2006) show that there is substantial demand for illiquid assets as a commitment device, and that such actions enhance savings.⁷ In our three-period model, suppose that self 1 has invested in illiquid assets, but a credit card company unexpectedly gives self 2 an opportunity to borrow in period 2. Self 2 is likely to accept this opportunity if her optimal choice is constrained by the investment in illiquid assets undertaken by self 1. In this instance, the expansion of financial markets may worsen the self-control problems for an individual.
5. **Why have national savings rates been declining?** Countries in the developed world have been experiencing a decline in the savings rate since the 1970s. This period has also coincided with an unprecedented expansion in the credit market. Easy availability of credit on reasonable terms puts a constraint on the ability of current selves to control actions such as consumption splurges of their future selves. This may lead to a reduction in savings (Laibson, 1997).
6. **Why is the marginal propensity to consume asset specific?** Thaler (1990) argues that consumers have different marginal propensities to consume (MPC) from different accounts for different assets. For instance, evidence suggests that the MPC from an unexpected job related bonus is close to unity, while the MPC from an unexpected increase in the value of an equity portfolio is close to zero. In our stylized three period model, we have not attempted to explain the mechanisms that give rise to variations in MPC. A possible explanation might run as follows. The current self tries to restrict the actions of the future selves by choosing to invest in illiquid assets that bind the choices of the future selves. Suppose that in (11.34), $s_2^j(s_1, z) = 0$, i.e., self 2 is liquidity constrained on account of the action taken by self 1. If self 2 now unexpectedly receives a bonus on the job, then she will consume it immediately because the choice of self 1 has restricted her liquidity. However, an unexpected appreciation of the illiquid asset, in the example, will benefit only self 3, without effecting the marginal propensity to consume of self 2.

⁵ For the effects of visceral influences on human choices and actions, see Part 6.

⁶ The alternative assumption of working for the first two periods, and retiring for the third period can be easily accommodated in the model.

⁷ On the demand for commitment devices for retirement savings in US, see Benartzi and Thaler (2004).

11.3.4 Euler equations in a general life-cycle problem

In this section, we derive the Euler equation in a more general life-cycle model, under exponential and hyperbolic discounting.

THE EULER EQUATION UNDER EXPONENTIAL DISCOUNTING

Consider an individual who uses the EDU model and lives for T periods. His date 0 utility function under EDU is

$$U_0(c_0, c_1, \dots, c_T) = E_0 \sum_{t=0}^T \delta^t u(c_t),$$

where E_0 is the expectation operator, conditional on the information set at $t = 0$, and $\delta \in [0, 1]$ is the discount factor. u is a strictly increasing, and concave, instantaneous utility function. $c_t \in R_+$ is consumption at time t . In any time period t , the income of the individual, y_t , is stochastic. The budget constraint in period t is

$$c_t + s_t = y_t + Rs_{t-1} \equiv W_t; t = 1, 2, \dots, T. \quad (11.37)$$

In (11.37), s_t is savings in period t and Rs_{t-1} is the financial income at time t arising on account of savings made in one-period maturity bonds at time $t - 1$. R is the time-invariant gross interest rate on savings. The total income from all sources in period t is denoted by W_t . The intertemporal budget constraint is given by

$$\sum_{t=0}^T R^{-t} c_t \leq \sum_{t=0}^T R^{-t} y_t \equiv W, \quad (11.38)$$

where W is the present discounted value of lifetime income. Since we assume local non-satiation of u , hence, (11.38) binds at the optimal solution. Letting ν be the Lagrangian multiplier on the constraint, the individual's optimization problem is given by

$$(c_0^*, c_1^*, \dots, c_T^*) \in \arg \max_{c_t \in R_+} E_0 \left[\sum_{t=0}^T \delta^t u(c_t) + \nu \sum_{t=0}^T R^{-t} (y_t - c_t) \right],$$

where c_t^* , $t = 0, 1, \dots, T$, is the optimal consumption in period t . The first order condition with respect to c_t for an interior solution is⁸

$$\delta^t E_0 u'(c_t) - \nu R^{-t} = 0; 0 \leq t \leq T. \quad (11.39)$$

Using (11.39), for any two consecutive dates, $j, j + 1$, $0 \leq j < T$ we get

$$E_0 u'(c_j) = R \delta E_0 u'(c_{j+1}). \quad (11.40)$$

⁸ The imposition of the Inada conditions and the finiteness of the budget constraint typically ensure an interior solution in these models.

In particular, for $j = 0$, the current date, we get

$$u'(c_0) = RE_0 \delta u'(c_1). \quad (11.41)$$

(11.40) gives the Euler equation in the EDU model. A slight rearrangement of (11.40) gives the familiar condition of equality between the marginal rate of substitution and the marginal rate of transformation: $\frac{E_0 u'(c_j)}{\delta E_0 u'(c_{j+1})} = R$. The Euler equation can also easily be derived using an arbitrage argument. The LHS of (11.40) is the change in current utility when current consumption is reduced by a dollar; this is the marginal cost. Over one period, the saved dollar is worth R dollars and it is weighted by the discounted marginal utility of consumption, $\delta u'(c_{j+1})$, in the next period; this is the marginal benefit.

THE EULER EQUATION UNDER HYPERBOLIC DISCOUNTING AND SOPHISTICATED MULTIPLE SELVES

We now draw on Laibson's (1998) treatment of the Euler equations in the presence of sophisticated multiple selves of an individual who uses hyperbolic discounting. Assume that the individual is a sophisticate (see (11.11)), hence, $\beta = \hat{\beta} < 1$. Consider the hyperbolic preferences stated in (11.8), and let $T \rightarrow \infty$. The problem is to maximize (11.8) by choosing an optimal consumption profile, subject to the budget constraint given in (11.38). We give a heuristic derivation of the *behavioral Euler equation* below.

Consider the consumption-savings choice of self 0. Suppose that self 0 reduces current consumption by one unit. This alters current utility by $u'(c_0)$; this is the marginal cost of the decision. However, the extra unit saved by self 0 creates a ripple effect on future instantaneous utilities that is felt by all future selves; the aggregate of this ripple effect is the marginal benefit. Self 0 reallocates income in this manner till the marginal costs and marginal benefits are equated. We first state the relevant arbitrage condition in (11.42), and then give a heuristic derivation.

$$u'(c_0) = \beta \left[\delta RE_0 u'(c_1) + \delta^2 R^2 \sum_{i=0}^{\infty} \delta^i R^i E_0 \left[\frac{\partial c_{2+i}}{\partial W_{2+i}} \prod_{j=0}^{i-1} \left(1 - \frac{\partial c_{1+j}}{\partial W_{1+j}} \right) \right] u'(c_{2+i}) \right]. \quad (11.42)$$

In any period t , $\frac{\partial c_t}{\partial W_t}$ is the marginal propensity to consume, so $1 - \frac{\partial c_t}{\partial W_t}$ is the marginal propensity to save. When self 0 saves an extra dollar, the marginal benefit, on account of the effect on self 1 is

$$\beta \delta RE_0 u'(c_1); \quad (11.43)$$

this is analogous to the RHS of (11.41) under EDU, but under hyperbolic discounting there is an extra parameter β that captures the present bias of the individual. This is the first term in the brackets on the RHS of (11.42). Self 1, however, does not consume the entire dollar received from self 0. Of that extra dollar, $1 - \frac{\partial c_1}{\partial W_1}$ is saved for self 2. The term corresponding to self 2 on the RHS of (11.42) is

$$\beta \delta^2 R^2 E_0 \frac{\partial c_2}{\partial W_2} \left(1 - \frac{\partial c_1}{\partial W_1} \right) u'(c_2). \quad (11.44)$$

To understand (11.44), consumption of self 2 increases by $\frac{\partial c_2}{\partial W_2} \left(1 - \frac{\partial c_1}{\partial W_1}\right)$, and this is weighted by her marginal utility of consumption $u'(c_2)$. Since this term arises in period 2, in order to convert it to time 0 units, we need to multiply by $\delta^2 R^2$ to account for (i) discounting over two periods (the term δ^2), and (ii) the return on the one-period financial instrument over two periods (the term R^2).

Of the original dollar saved by self 0, a fraction $\left(1 - \frac{\partial c_1}{\partial W_1}\right) \left(1 - \frac{\partial c_2}{\partial W_2}\right)$ goes to self 3. Hence, we can construct the present discounted value of the marginal benefit (to self 0), in utility terms, on account of self 3 as

$$\beta \delta^3 R^3 E_0 \frac{\partial c_3}{\partial W_3} \left(1 - \frac{\partial c_1}{\partial W_1}\right) \left(1 - \frac{\partial c_2}{\partial W_2}\right) u'(c_3). \quad (11.45)$$

The explanation of (11.45) is parallel to that of (11.44).

We can construct the marginal returns for all future selves in this manner. Adding together the resulting terms such as (11.43), (11.44), (11.45) one obtains the composite present discounted value of marginal benefit of a dollar saved by self 0; this is simply the RHS of (11.42).

Moving (11.42) forward in time by one period, while maintaining the information set at time 0, we get

$$u'(c_1) = \beta \left[\delta R E_0 u'(c_2) + \delta^2 R^2 \sum_{i=0}^{\infty} \delta^i R^i E_0 \left[\frac{\partial c_{3+i}}{\partial W_{3+i}} \prod_{j=0}^{i-1} \left(1 - \frac{\partial c_{2+j}}{\partial W_{2+j}}\right) \right] u'(c_{3+i}) \right]. \quad (11.46)$$

Multiply both sides of (11.46) by $\delta R \left(1 - \frac{\partial c_1}{\partial W_1}\right)$ and subtract from (11.42) to get

$$u'(c_0) = RE_0 \delta \left[1 - (1 - \beta) \frac{\partial c_1}{\partial W_1} \right] u'(c_1). \quad (11.47)$$

(11.47) gives the *behavioral Euler equation* under hyperbolic discounting. Notice that with $\beta = 1$, (11.47) reduces to the Euler equation under EDU given in (11.41). Defining

$$\tilde{\delta}_1 = \delta \left[1 - (1 - \beta) \frac{\partial c_1}{\partial W_1} \right], \quad (11.48)$$

as the *effective discount factor* under hyperbolic discounting, we can express the Euler equation under hyperbolic discounting in a form that is similar to (11.41)

$$u'(c_0) = RE_0 \tilde{\delta}_1 u'(c_1). \quad (11.49)$$

In general, we could have set the starting date to any arbitrary positive $t < T$. In this case, the Euler equation in (11.49) can be written as

$$u'(c_t) = RE_t \tilde{\delta}_{t+1} u'(c_{t+1}) \text{ where } \tilde{\delta}_{t+1} = \delta \left[1 - (1 - \beta) \frac{\partial c_{t+1}}{\partial W_{t+1}} \right]. \quad (11.50)$$

CRR preferences have the following property that considerably simplifies the Euler equation. The optimal consumption in any period, t , c_t , is a constant fraction of W , the lifetime present

discounted value of income defined in (11.38). This fraction is given by $\omega_t = \omega_t(\beta, \delta, R) \in (0, 1)$, so that $c_t = \omega_t W = \omega_{j=t} \sum_{j=t}^T R^{j-t} y_j$. Hence, the marginal propensity to consume out of lifetime income is constant. Furthermore, using $y_j = W_j - R s_{j-1}$ (see (11.37)), we can write

$$c_t = \omega_t \sum_{j=t}^T R^{j-t} (W_j - R s_{j-1}). \quad (11.51)$$

Using (11.51) we can compute

$$\frac{\partial c_t}{\partial W_t} = \omega_t. \quad (11.52)$$

Substituting (11.52) in (11.50), for CRRA preferences, the Euler equation under hyperbolic discounting is given by

$$u'(c_t) = R\delta E_t [1 - (1 - \beta) \omega_{t+1}] u'(c_{t+1}), \quad (11.53)$$

where ω_{t+1} is given by moving forward (11.52) by one period. We will now use (11.53) to illustrate how hyperbolic discounting can explain several anomalies of the life-cycle model under the assumption of EDU.

11.3.5 Missing precautionary savings effects

Individuals engage in *precautionary savings* if they respond to uncertainty in future consumption by saving more in the current period. Let us first derive a tractable expression for the Euler equation (11.41) under the EDU model in the following special case. Suppose that preferences are of the CRRA form and the log of consumption is normally distributed, i.e., $\ln c_t \sim N(\mu, \sigma^2)$. Rewriting (11.41) for CRRA preferences we get

$$c_0^{-\gamma} = R\delta E_0 [c_1^{-\gamma}]. \quad (11.54)$$

Since the exponential function is the inverse of the log function, we have

$$E_0 [e^{-\gamma \ln c_1}] = E_0 [c_1^{-\gamma}]. \quad (11.55)$$

Hence, using (11.55) we get⁹

$$E_0 [e^{-\gamma \ln c_1}] = E_0 [c_1^{-\gamma}] = e^{-\gamma \mu + \frac{1}{2} \gamma^2 \sigma^2}, \quad (11.56)$$

where $\mu = E_0 [\ln c_1]$ is the unconditional mean of $\ln c_t$ and σ^2 is the unconditional variance. Substituting (11.56) in (11.54) we get

$$c_0^{-\gamma} = R\delta e^{-\gamma E_0 [\ln c_1] + \frac{1}{2} \gamma^2 \sigma^2}. \quad (11.57)$$

⁹ We use a standard result in statistics that if $x \sim N(\mu, \sigma^2)$ then $E[e^{-\rho x}] = e^{-\rho(\mu - \frac{\rho}{2}\sigma^2)}$.

Taking logs on both sides of (11.57) and dividing through by $-\gamma$ we get

$$\ln c_0 = \left(\frac{1}{-\gamma} \right) [\ln R + \ln \delta] + E_0 [\ln c_1] - \frac{1}{2} \gamma \sigma^2. \quad (11.58)$$

Defining $\Delta E_0 [\ln c_1] = E_0 [\ln c_1] - \ln c_0$ and simplifying (11.58) we get

$$\begin{aligned} \Delta E_0 [\ln c_1] &= \frac{1}{\gamma} [\ln R + \ln \delta] + \frac{1}{2} \gamma \sigma^2. \\ \Rightarrow \Delta \ln c_1 &= \frac{1}{\gamma} [\ln R + \ln \delta] + \frac{1}{2} \gamma \sigma^2 + \epsilon_1, \end{aligned} \quad (11.59)$$

where ϵ_1 is a random term that cannot be predicted at time $t = 0$ when the decision is made. When equations of the form (11.59) are estimated, the estimated coefficient of σ^2 is found to be very small. So high consumption variability is not associated, in the data, with high consumption growth; for instance Browning and Lusardi (1996) and Lusardi (1998). Hence, the precautionary motive seems extremely small, or absent in the data.

These results seem consistent with the hyperbolic model. Intuitively, the current bias will prevent the individual from undertaking adequate precautionary saving. Laibson (1998) points out that a high variability in consumption growth is normally associated with a high value of the one period ahead marginal propensity to consume, i.e., when σ^2 is high, ω_1 is high. Consider the Euler equation under hyperbolic discounting in (11.53) for $t = 0$. A high value of ω_1 reduces the effective discount rate and brings c_0 and c_1 closer, reducing the slope of the optimal consumption profile. This offsets the precautionary savings effect under hyperbolic discounting relative to the EDU model, which is consistent with the evidence (Laibson, 1998).

11.3.6 Relation between intertemporal substitution and risk

Let us focus on the case of full certainty. The *elasticity of intertemporal substitution of consumption* (EIS) between two consecutive time periods is the percentage change in the ratio of consumption that results from a given percentage change in the relative price of consumption, along the optimal path. In the intertemporal context, the price ratio is simply the gross interest rate, hence, in order to compute the EIS, we need to find $\partial \ln (c_{t+1}/c_t) / \partial \ln R$. For CRRA preferences, and in the absence of uncertainty, (11.54) can be rewritten for any two consecutive time periods as

$$\begin{aligned} c_t^{-\gamma} &= R \delta c_{t+1}^{-\gamma}. \\ \Rightarrow \ln \frac{c_{t+1}}{c_t} &= \frac{1}{\gamma} [\ln R + \ln \delta], \end{aligned}$$

hence,

$$\frac{\partial \ln (c_{t+1}/c_t)}{\partial \ln R} = \frac{1}{\gamma}. \quad (11.60)$$

Under EDU, this sets a fairly strong restriction on the relation between EIS and γ ; each is the inverse of the other. This is not supported by the empirical evidence (Hall, 1988). The general view in the empirical literature appears to be that (Laibson, 1998):

$$\frac{\partial \ln (c_{t+1}/c_t)}{\partial \ln R} < \frac{1}{\gamma}. \quad (11.61)$$

The hyperbolic discounting model can potentially account for (11.61). Writing the Euler equation under hyperbolic discounting, (11.53), under certainty and CRRA preferences, for any two consecutive time periods $t, t + 1$ we get

$$\begin{aligned} c_t^{-\gamma} &= R\delta [1 - (1 - \beta)\omega_{t+1}] c_{t+1}^{-\gamma}. \\ \Rightarrow \ln \frac{c_{t+1}}{c_t} &= \frac{1}{\gamma} [\ln R + \ln \delta + \ln [1 - (1 - \beta)\omega_{t+1}]]. \end{aligned} \quad (11.62)$$

Differentiating (11.62) with respect to $\ln R$ on both sides we get

$$EIS = \frac{\partial \ln (c_{t+1}/c_t)}{\partial \ln R} = \frac{1}{\gamma} - \frac{(1 - \beta)}{\gamma [1 - (1 - \beta)\omega_{t+1}]} \frac{\partial \omega_{t+1}}{\partial \ln R}. \quad (11.63)$$

If $\gamma > 1$, as suggested by several empirical studies (Laibson, 1998), then $\frac{\partial \omega_{t+1}}{\partial \ln R} > 0$. The intuition is as follows. As is well known, interest rate changes have opposing income and substitution effects, hence, the effect on consumption is ambiguous. However, when $\gamma > 1$, the income effect overweighs the substitution effect and so an increase in R increases consumption in period $t + 1$, hence, $\frac{\partial \omega_{t+1}}{\partial \ln R} > 0$. But this implies that $EIS < \frac{1}{\gamma}$, which is in accord with the empirical evidence.

11.3.7 Patience heterogeneity: covariation with wealth and age

Relatively young consumers expect rising income profiles in the future, so they might have high marginal propensity to consume (MPC). Consumers who are liquidity constrained might also have a high MPC. Conversely, old consumers and those who are not liquidity constrained are likely to have low MPC. Hence, the effective discount factor in hyperbolic models, see (11.48) and (11.53), varies over an individual's life and is likely to be high when individuals are old and have greater financial wealth. These results are consistent with the experimental evidence, see Green et al. (1999), but are difficult to reconcile with the EDU model.

11.4 Endogenous retirement decisions

In a model of multiple selves, retirement is a decision made by one's future selves. Current selves, however, have an important stake in that decision because the retirement decision affects the future profile of income, so, it must influence the amount of savings that current selves need to make for the future. What actions can current selves take to influence the retirement decisions of future selves? When are these actions effective? What is the role of labor market policy in the presence of such concerns? We now consider these issues in the three period model of Diamond and Köszegi (2003).

Denote the three periods of an individual's lifetime as periods 1, 2, 3. Each period corresponds to a distinct self of the individual; self $t = 1, 2, 3$ lives in period t . The individual has two kinds of decisions to make. The *savings decision*, i.e., how much to consume/save, and the *retirement decision*, i.e., when to retire?

It is mandatory for self 1 to work. Self 1 inelastically supplies one unit of labor and earns a wage, $y > 0$. To make the retirement decision tractable, we assume that only self 2 makes the retirement decision. Self 1 must always work, and self 3 must always be retired. Self 2 can either inelastically supply a unit of labor and earn wage income, y , or retire, in which case the only source of income is the savings made by self 1, s_1 . There is a fixed cost $e > 0$ to self 2 from working.¹⁰ For simplicity, assume (β, δ) preferences, $0 < \beta < 1$. Assume also that self 1 is *sophisticated*, i.e., has full awareness of the self-control problems of the future selves (see (11.11)).¹¹ We assume that the utility function is of the CRRA form.

$$u(c) = \frac{1}{1-\rho} c^{1-\rho}; \quad \rho > 0, \rho \neq 1,$$

where $c \geq 0$ is real-valued consumption. For CRRA preferences, $u' > 0$ and $u'' < 0$. The next lemma states a useful result.

Lemma 11.2 *Consider the maximization problem*

$$\text{Max}_{(c_1, c_2)} u(c_1) + \lambda u(c_2), \quad \lambda > 0,$$

subject to

$$c_1 + \frac{1}{R} c_2 = W,$$

where $R > 0$ and $W > 0$. When u takes the CRRA form, then there exists $\omega(R, \lambda) : \mathbb{R}_+^2 \rightarrow [0, 1]$ such that the optimal solution is $c_1^* = \omega(R, \lambda)W$. Furthermore, $\frac{\partial \omega}{\partial \lambda} < 0$.

The preferences of self 1 are given by

$$U_1 = \begin{cases} u(c_1) + \beta \delta [u(c_2) - e] + \beta \delta^2 u(c_3) & \text{if self 2 works,} \\ u(c_1) + \beta \delta u(c_2) + \beta \delta^2 u(c_3) & \text{if self 2 retires.} \end{cases} \quad (11.64)$$

As compared to self 2, self 1 cares relatively more about self 3. This is so because the discount factor for self 1 over the two time periods 2, 3 is $\frac{\beta \delta^2}{\beta \delta} = \delta$, while the discount factor of self 2 over time periods 2, 3 is $\beta \delta < \delta$. Hence, self 1 may prefer that self 2 does not retire. However, because self 2 puts relatively lower weight on the utility of self 3, he might be inclined to retire.

Can self 1 take any actions that may induce self 2 to not retire? One possibility is for self 1 to save a “sufficiently low” amount that induces self 2 to continue working. However, saving too little is also costly for self 1 because it reduces the income of self 3. If these costs outweigh the potential benefits, then self 1 does not wish to influence the retirement decision of self 2, through reduced savings. In this case, self 1 is forced to save “too much,” so that some of the savings are passed on to self 3.

The budget constraint of self 1 is given by

$$c_1 + s_1 = y \equiv W_1. \quad (11.65)$$

¹⁰ Including a fixed cost of working for self 1 is not important because we assume that he works regardless.

¹¹ Diamond and Koszegi (2003) consider, in addition, the case of partially naive agents.

The solution to this dynamic game of perfect information is by backward induction. The only source of income for self 3 is the savings made by self 2, s_2 . Savings earn a gross return R , hence, period 3 wealth is $W_3 = Rs_2$. Since $u' > 0$, and there are no bequests, self 3 consumes all available wealth, thus, his utility is given by

$$U_3 = u(Rs_2).$$

Self 2 has two possible choices to make, work or retire; we use superscripts w (work) and r (retire) on the relevant variables to denote these two decisions. In making these decisions, self 2 takes as given the savings, s_1 , made by self 1. Thus, we may write the utility function of self 2 as

$$\begin{aligned} U_2^w(c_2, c_3 | s_1) &= u(c_2) - e + \beta \delta u(c_3) & \text{if self 2 works,} \\ U_2^r(c_2, c_3 | s_1) &= u(c_2) + \beta \delta u(c_3) & \text{if self 2 retires.} \end{aligned} \quad (11.66)$$

The budget constraint of self 2, and the third period consumption, c_3 , under the two decisions is

$$\begin{aligned} c_2 + s_2 &\leq y + Rs_1 \equiv W_2^w & \text{and } c_3 = Rs_2 & \text{if self 2 works,} \\ c_2 + s_2 &\leq Rs_1 \equiv W_2^r & \text{and } c_3 = Rs_2 & \text{if self 2 retires.} \end{aligned} \quad (11.67)$$

Since $u' > 0$, the budget constraints bind. From the perspective of self 2, we may alternatively write these budget constraints in an intertemporal form as

$$c_2 + \frac{1}{R}c_3 = W_2^i, i = w, r. \quad (11.68)$$

where W_2^w, W_2^r are defined in (11.67).

Using Lemma 11.2 we know that under CRRA preferences, the optimal consumption of self 2 is given by

$$c_2^i = \omega_2(R, \beta \delta) W_2^i, i = w, r. \quad (11.69)$$

Denote the optimal level of savings of self 2 by $s_2^w(s_1)$ when he works, and by $s_2^r(s_1)$ when he retires. From (11.69), $s_2^i(s_1) = (1 - \omega_2(R, \beta \delta)) W_2^i, i = w, r$. The corresponding consumption levels of self 3 are given by $c_3^i = Rs_2^i(s_1), i = w, r$. The consumption levels, $c_2^w, c_3^w, c_2^r, c_3^r$ are all functions of the savings choice made by self 1, s_1 . Substituting these optimal levels of consumption in (11.66), we get the indirect utility of self 2 under the two decisions.

$$\begin{aligned} V_2^w(s_1) &= u(c_2^w) - e + \beta \delta u(c_3^w) & \text{if self 2 works,} \\ V_2^r(s_1) &= u(c_2^r) + \beta \delta u(c_3^r) & \text{if self 2 retires.} \end{aligned} \quad (11.70)$$

Self 2 prefers to work rather than retire if $\Delta V_2(s_1) = V_2^w(s_1) - V_2^r(s_1) \geq 0$, or

$$\Delta V_2(s_1) = [u(c_2^w) - e - u(c_2^r)] + \beta \delta [u(c_3^w) - u(c_3^r)] \geq 0. \quad (11.71)$$

In order to make the problem interesting we assume that

$$\Delta V_2(0) > 0, \quad (11.72)$$

thus, if self 1 leaves no income for self 2, then self 2 prefers to work. We also assume that there is some lower bound on the cost of effort, e , which ensures that¹²

$$\Delta V_2(y) < 0, \quad (11.73)$$

i.e., if self 2 inherits the entire income of self 1, he chooses to retire.

Differentiating $\Delta V_2(s_1)$ with respect to s_1 , and using (11.67), (11.69), we get

$$\Delta V'_2(s_1) = R\omega_2 [u'(c_2^w) - u'(c_2^r)] + \beta\delta R(1 - \omega_2) [u'(c_3^w) - u'(c_3^r)] < 0. \quad (11.74)$$

We use two facts to sign (11.74). First, for CRRA preferences, $\lim_{c \rightarrow 0} u'(c) = \infty$, hence, the first period savings choice of self 1, $s_1 \in [0, y]$. For any fixed $s_1 > 0$, (11.67) implies that $W_2^w(s_1) > W_2^r(s_1)$, thus, for CRRA preferences, $c_2^w > c_2^r$ and $c_3^w > c_3^r$. Second, for CRRA preferences, the instantaneous utility function is strictly increasing and strictly concave, so $u'(c_2^w) < u'(c_2^r)$ and $u'(c_3^w) < u'(c_3^r)$. Hence, $\Delta V'_2(s_1) < 0$. Furthermore, $\Delta V_2(s_1)$ is a continuous function over $[0, y]$. From (11.72), (11.73), and the continuity of ΔV_2 , there must exist some critical level of savings $s_1 = s^*$ such that for $s_1 < s^*$, $\Delta V_2 > 0$ (self 2 works) and for $s_1 > s^*$, $\Delta V_2 < 0$ (self 2 retires). At $s_1 = s^*$ we have that self 2 is indifferent between working and retiring. We use the tie-breaking rule that in this case, self 2 chooses to work. We summarize these results for later reference in Lemma 11.3 below, which shows how the savings decision of self 1 influences the retirement decision of self 2.

Lemma 11.3 *Self 2 works if $s_1 \leq s^*$ and retires if $s_1 > s^*$, where $\Delta V_2(s^*) = 0$. Furthermore, $s^* > 0$.*

The final stage of the backward induction process is to find the optimizing choice of self 1. Self 1 is aware that the optimal consumption values for periods 2 and 3, are either (c_2^w, c_3^w) , or (c_2^r, c_3^r) . However, self 1 temporally precedes self 2, hence, under hyperbolic discounting, he differs in the relative weight placed on future utilities. Using (11.64), self 1 evaluates the utilities in (11.70) as follows.

$$\begin{aligned} V_1^w(s_1) &= \beta\delta (u(c_2^w) - e) + \beta\delta^2 u(c_3^w) & \text{if self 2 works,} \\ V_1^r(s_1) &= \beta\delta u(c_2^r) + \beta\delta^2 u(c_3^r) & \text{if self 2 retires.} \end{aligned}$$

Hence, self 1 prefers self 2 to work, rather than retire, if $\Delta V_1(s_1) = V_1^w(s_1) - V_1^r(s_1) \geq 0$, or

$$\Delta V_1(s_1) = [u(c_2^w) - e - u(c_2^r)] + \delta [u(c_3^w) - u(c_3^r)] \geq 0. \quad (11.75)$$

Notice that (11.75) is similar to (11.71), except for the different discount rates used to discount the outcomes for period 3. Since, $0 < \beta < 1$, and $u(c_3^w) - u(c_3^r) > 0$, a comparison of (11.71) and (11.75) reveals that

$$\Delta V_1(s_1) > \Delta V_2(s_1) \text{ for all } s_1 \in (0, y). \quad (11.76)$$

¹² The lower bound on effort, $e = e_L$ is given by

$$[u(\omega_2 y(1 + R)) - u(\omega_2 yR)] + \beta\delta [u((1 - \omega_2)y(1 + R)) - u((1 - \omega_2)yR^2)] \leq e_L.$$

Let $s_1 = s^{**}$ be the savings level such that $\Delta V_1(s^{**}) = 0$; the existence of such a s^{**} can be shown in a manner similar to the existence of s^* , above. As in (11.72) and (11.73), we assume that $\Delta V_1(0) > 0$ and $\Delta V_1(y) < 0$.

Lemma 11.4 *Suppose that $\Delta V_1(0) > 0$ and $\Delta V_1(y) < 0$. The savings level, $s_1 = s^{**}$, at which self 1 is indifferent between the work/retire decision of self 2, is greater than the corresponding value of $s_1 = s^*$ at which self 2 is indifferent between the two decisions, i.e., $s^{**} > s^*$. Given our tie-breaking convention, self 1 prefers self 2 to work for any $s_1 \leq s^{**}$ and retire for any $s_1 > s^{**}$.*

Proof: $\Delta V_1(s_1)$ is a continuous function over $[0, y]$, $\Delta V_1'(s_1) < 0$, $\Delta V_1(0) > 0$ and $\Delta V_1(y) < 0$. Given (11.76), at $s_1 = s^*$ we have $\Delta V_1(s^*) > \Delta V_2(s^*)$. Hence, it must be the case that $s^{**} > s^*$.

Using Lemma 11.4, we summarize the difference in decisions between selves 1 and 2 in Lemma 11.5.

Lemma 11.5 *For any level of saving by self 1, $s_1 : 0 \leq s_1 \leq s^*$ both selves 1 and 2 prefer self 2 to work rather than retire. For $s^* < s_1 \leq s^{**}$ self 2 prefers to retire, but self 1 prefers self 2 to work. Finally, for $s_1 > s^{**}$ both selves 1 and 2 would prefer self 2 to retire.*

From Lemma 11.5, it is the region $s^* < s_1 \leq s^{**}$ that creates a conflict of interest between the two selves.

Suppose that self 1 does not have a commitment device to bind self 2 to either work, or retire. We then need to find the optimal savings of self 1 that serve as an indirect control. We solve the problem in two steps. We shall denote the utility of self 1 when self 2 works by $U_1^w(s_1)$. If self 2 retires, the utility of self 1 is denoted by $U_1^r(s_1)$. Denote the optimal levels of savings by self 1 over the interval $[0, y]$ by s_1^w and s_1^r , respectively.

For $0 \leq s_1 \leq s^*$, self 2 works, giving self 1 a continuation payoff of $\beta\delta u(c_2^w) + \beta\delta^2 u(c_3^w)$, hence, in this region, the optimization problem of self 1 is

$$s_1^{w*} \in \arg \max U_1^w(s_1) = u(y_1 - s_1) + \beta\delta [u(c_2^w) - e] + \beta\delta^2 u(c_3^w); 0 \leq s_1 \leq s^*. \quad (11.77)$$

For any $s_1 > s^*$, self 2 retires and the continuation payoff of self 1 is $\beta\delta u(c_2^r) + \beta\delta^2 u(c_3^r)$, hence, in this region, the optimization problem of self 1 is

$$s_1^{r*} \in \arg \max U_1^r(s_1) = u(y_1 - s_1) + \beta\delta u(c_2^r) + \beta\delta^2 u(c_3^r); s_1 > s^*. \quad (11.78)$$

The optimal savings choice of self 1 is given by

$$\begin{cases} s_1^{w*} & \text{if } U_1^w(s_1^{w*}) \geq U_1^r(s_1^{r*}), \\ s_1^{r*} & \text{if } U_1^w(s_1^{w*}) < U_1^r(s_1^{r*}). \end{cases} \quad (11.79)$$

We illustrate the equilibrium, heuristically, by using Figure 11.1. Each of the panels in Figure 11.1 plots hypothetical shapes of U_1^w and U_1^r over the unconstrained interval $[0, y]$; these curves reach their maximum values at s_1^w and s_1^r , respectively. However, the curve U_1^w applies only to the range $[0, s^*]$, while the curve U_1^r applies only to the range $(s^*, y]$; see (11.77) and (11.78). Over the interval $[0, s^*]$, for the particular curves drawn in Figure 11.1, $U_1^w(s_1)$ reaches its maximum at the boundary point, s^* . In Panel A of Figure 11.1, $U_1^w(s^*) > U_1^r(s_1^r)$, hence, self 1 optimally saves an amount s^* , so self 2 chooses to work. By contrast, in Panel B of Figure 11.1, $U_1^w(s^*) < U_1^r(s_1^r)$, so self 1 optimally saves an amount $s_1^r > s^*$, which induces self 2 to retire.




Figure 11.1 Diagrammatic illustration of the equilibrium in the Diamond-Köszegi model.

In Panel A of Figure 11.1, self 1 optimally undersaves ($s^* < s_1^w$) to force self 2 to work. This is termed as *strategic undersaving*. In Panel B of Figure 11.1, on the other hand, $s_1^w < s_1^r$, so self 1 oversaves. This is termed as *resigned oversaving*. In this case, self 1 knows that it is too costly to prevent self 2 from retirement, so resigns to saving an extra amount for the benefit of self 3.

We can also use Figure 11.1 to illustrate the effects of employment legislation such as minimum retirement age. Suppose that the original employment legislation is that one must work in period 2. However, a new legislation repeals the original legislation, so that the individual is now free to choose when to retire. What are the effects on savings and retirement?

Let us first consider the equilibrium savings and retirement decisions under the original legislation which binds self 2 to work regardless of the actions of self 1. The relevant curve in each panel is now the $U_1^w(s_1)$ curve. Self 1 optimally chooses the level of savings s_1^w and self 2 works. Under the new legislation, self 2 works only if it is optimal for him, so the equilibrium is as described in the absence of any employment legislation. Hence, in Panel A of Figure 11.1 the optimal savings are reduced from s_1^w to s^* and there is no change in the retirement decision. By contrast, in Panel B of Figure 11.1 the new savings are higher under the new legislation ($s_1^r > s_1^w$) and self 2 retires while under the older legislation he was required to work.

11.5 Procrastination and preproperation

Suppose that an individual must undertake a *single* task, or activity, *once* over T periods. In each period, the individual must decide whether to undertake or postpone the activity, conditional on not having undertaken the activity in the past. Once the activity is undertaken, the individual has no further action to take in any period in the future. Denote by $v_t \geq 0$, and $c_t \geq 0$, respectively, the reward and the cost of the activity in period t . Thus, the reward and the cost schedules are given, respectively, by $v = (v_1, v_2, \dots, v_T)$ and $c = (c_1, c_2, \dots, c_T)$.

One can imagine T selves, one for each period, that can undertake the activity. The self in any period t cannot bind the action of any future self in period τ , where $t < \tau \leq T$. Each self has quasi-hyperbolic (β, δ) preferences with $\delta = 1$, i.e., $(\beta, 1)$ preferences. In determining whether a future self will carry out the activity, the beliefs of each self are that any future self has $(\hat{\beta}, 1)$

preferences. Depending on the relation between β and $\hat{\beta}$, these beliefs can be of three possible types, given in (11.10), (11.11), and (11.13), i.e., time consistent, sophisticates, and naifs (we exclude partial naifs here).

The main question that we ask is: In which time period will each of the three types of individuals undertake the action? Do the answers depend on the temporal pattern of rewards and costs? If so, how? We follow O'Donoghue and Rabin (1999b) in exploring these issues further.

We denote by $U_t(\tau)$, the utility of self t if the activity is undertaken in some time period $\tau \geq t$. To illustrate alternative temporal patterns of rewards and costs, we distinguish between the following two main cases; results turn out to be very different for these cases.

1. **Immediate costs:** Suppose that self τ undertakes the activity. The case of *immediate cost* occurs if costs are incurred in period τ , but the reward is received in some future period. Examples include the following. Cleaning one's car, or garage, organizing one's desk, writing a paper, looking for a job, studying for exams, mowing one's lawn, and quitting an addiction such as smoking.

With $(\beta, 1)$ preferences, all future periods are weighted equally, so it is not necessary to specify the exact future period in which the reward is received. The intertemporal utility $U_t(\tau)$ is, therefore, given by

$$U_t(\tau) = \begin{cases} \beta v_t - c_t & \text{if } \tau = t, \\ \beta (v_t - c_t) & \text{if } \tau > t. \end{cases} \quad (11.80)$$

2. **Immediate rewards:** The case of *immediate rewards* occurs if rewards are received immediately, but costs are incurred in some future period. Examples include going away on vacation, watching a late night movie, taking a day off, and splurging on credit card debt. Under $(\beta, 1)$ preferences, the exact future period in which the cost is incurred is not important. Suppose that self τ undertakes the activity. The intertemporal utility of self t , $U_t(\tau)$, is given by

$$U_t(\tau) = \begin{cases} v_t - \beta c_t & \text{if } \tau = t \\ \beta (v_t - c_t) & \text{if } \tau > t. \end{cases} \quad (11.81)$$

The natural equilibrium concept in these problems is a slight modification of subgame perfection. One may term the strategies followed in such an equilibrium as *perception-perfect* strategies; see O'Donoghue and Rabin (1999b). These strategies require that individuals maximize the current payoff given their beliefs about how their future selves will behave. An individual's strategy is a complete plan of action for each of the T selves. Denote a strategy by $s^i = (s_1^i, s_2^i, \dots, s_T^i)$, where s_t^i is the action of self t , and the type of the player is indexed by $i = c, s, n$ for, respectively, time consistent, sophisticates, and naifs (see (11.10), (11.11), and (11.13)). Any s_t^i can take only two possible values in the set $\{Y, N\}$, where Y denotes "undertake activity in period t " and N denotes "do not undertake the activity in period t ."

Definition 11.16 A "perception-perfect strategy" for a type $i = c, n$ (respectively, time consistent, and naif) player is a strategy $s^i = (s_1^i, s_2^i, \dots, s_T^i)$ that satisfies $\forall t < T, s_t^i = Y$, if and only if, $U_t(t) > U_t(\tau)$ for all $\tau > t$. To rule out situations where the activity might never be undertaken, we assume that it is always optimal for the action to be taken in the last period, i.e.,

$$s_T^i = Y. \quad (11.82)$$

The definition of a perception-perfect strategy for time consistent and naifs is identical, although they lie at the two polar extremes in terms of their self-awareness. The reason is that a naif in any period believes that all future selves will behave like time consistent individuals.

Definition 11.17 A perception-perfect strategy for a sophisticate is a strategy $\mathbf{s}^s = (s_1^s, s_2^s, \dots, s_T^s)$ that satisfies $\forall t < T, s_t^s = Y$, if and only if, $U_t(t) > U_t(\underline{\tau})$ where $\underline{\tau} = \min_{\tau > t} \{\tau \mid s_\tau^s = Y\}$.

A sophisticate (Definition 11.17) knows that future selves will have identical self-control problems. So there might be more than a few future selves that actually succumb to temptation. The current self only needs to know the earliest future self that will succumb, because the activity is undertaken only once. By contrast, time consistent and naifs (Definition 11.16) imagine that all future selves have no self-control problems, so the plans that they make are time consistent. Hence, they only need to know if it is optimal to do it now, or later. In equilibrium, the current selves of time consistent will find that their beliefs come out to be true, but naifs often find that their beliefs do not. The examples below will further clarify the definitions.

The method of solution is by backward induction. We first determine the optimal action of self T ; then, the optimal action of self $T - 1$, conditional on the type-contingent beliefs of $T - 1$, and so on. To illustrate the equilibrium in the model, we give the following two examples, motivated by O'Donoghue and Rabin (1999b); each considers a four period model. In each example, we assume $(\beta, 1)$ preferences for individuals who follow quasi-hyperbolic discounting.

11.5.1 The case of immediate costs

Suppose that an individual wishes to reallocate her portfolio of assets that is currently suboptimal. If the individual optimally reallocates the portfolio, her reward is $\bar{v}$, in the following period. Under $(\beta, 1)$ preferences it does not matter which future period the reward $\bar{v}$ is received in. Let $\bar{v} > 26$, $T = 4$, and $\beta = 1/2$. The cost of not reallocating the portfolio optimally increases over time, perhaps on account of the cumulative returns forgone every period. Suppose that the cost schedule is given by

$$\mathbf{c} = (3, 5, 8, 13). \quad (11.83)$$

If self t undertakes the reallocation, she incurs an immediate cost c_t , and receives the reward $\bar{v}$ in the next period. Hence, under $(\beta, 1)$ preferences, her utility is

$$U_t(t) = \beta \bar{v} - c_t = \frac{\bar{v}}{2} - c_t. \quad (11.84)$$

Using (11.83), (11.84), the utility of the period 4 self from undertaking the action is

$$U_4(4) = \frac{\bar{v}}{2} - 13 > 0 \Leftrightarrow \bar{v} > 26, \quad (11.85)$$

which is true. Using (11.83), and (11.84), the lifetime utility of each of the selves from the portfolio reallocation is

$$U_1(1) = \frac{\bar{v}}{2} - 3; U_2(2) = \frac{\bar{v}}{2} - 5; U_3(3) = \frac{\bar{v}}{2} - 8; U_4(4) = \frac{\bar{v}}{2} - 13. \quad (11.86)$$

Let us now look at the decision made by each of the selves, conditional on their beliefs about the self-control problems of the future selves. Recall that $U_t(\tau)$ is the utility of self t , if the activity is undertaken in period $\tau \geq t$.

1. **Time consistent** ($\beta = \hat{\beta} = 1$): From Definition 11.16, these individuals undertake the activity in the period which gives the highest current payoff. Here,

$$U_1(1) = \bar{v} - 3; U_2(2) = \bar{v} - 5; U_3(3) = \bar{v} - 8; U_4(4) = \bar{v} - 13, \quad (11.87)$$

In period 1, the cost is the lowest, so the activity is undertaken immediately. One can check that a perception-perfect strategy is (Y, Y, Y, Y) because $U_t(t) > U_t(\tau)$ for all t , and for all $\tau > t$.

2. **Sophisticates** ($\beta = \hat{\beta} < 1$): We now use Definition 11.17 to analyze the decision of sophisticates. From (11.85), self 4 must choose Y (i.e., reallocate the portfolio). Anticipating correctly this decision, self 3 will choose N (i.e., not reallocate the portfolio) because

$$U_3(4) > U_3(3) \Leftrightarrow \beta(\bar{v} - c_4) > \beta\bar{v} - c_3 \Leftrightarrow 6.5 < 8. \quad (11.88)$$

Anticipating correctly the decisions of the subsequent selves, $\{N, Y\}$, self 2 will choose Y because

$$U_2(4) < U_2(2) \Leftrightarrow \beta(\bar{v} - c_4) < \beta\bar{v} - c_2 \Leftrightarrow 5 < 6.5. \quad (11.89)$$

Self 1 anticipates correctly the decisions of the subsequent selves to be $\{Y, N, Y\}$. Thus, his choice is between doing it now, or leaving it to be done by self 2 (the earliest self who will choose Y , in accordance with Definition 11.17). Therefore, self 1 chooses N because

$$U_1(2) > U_1(1) \Leftrightarrow \beta(\bar{v} - c_2) > \beta\bar{v} - c_1 \Leftrightarrow 2.5 < 3. \quad (11.90)$$

Hence, for sophisticates, the *perception-perfect strategy profile* of the four selves is $\{N, Y, N, Y\}$. The portfolio is reallocated in period 2, one period later than the time consistent.

3. **Naifs** ($\beta < \hat{\beta} = 1$): Naifs believe, incorrectly, that all their future selves are time consistent, and, thus, have no self-control problems. In particular, self t believes, incorrectly, that self $t + 1$ is a time consistent and will choose Y (see (11.87)). A simple calculation shows that the naifs in periods 1 through 3 will all choose N because

$$U_t(t) < U_t(t+1) \Leftrightarrow \beta\bar{v} - c_t < \beta(\bar{v} - c_{t+1}); t = 1, 2, 3. \quad (11.91)$$

Hence, naifs will do the task only in period 4 (see (11.85)).

In the case of immediate costs, therefore, lack of self-awareness about one's own future selves (typified most strongly for naifs) will result in greater procrastination. The reason is that lack of self-awareness leads one to underestimate procrastination in one's future selves. Sophisticates, on the other hand, have perfect foresight about the future, and they anticipate, correctly, that their future selves will procrastinate. Hence, being more aware of the self-control problems of their

future selves, the costs of procrastinating in the current period are higher. This mitigates their propensity to procrastinate; an idea originally recognized by Strotz (1956) and Akerlof (1991).

11.5.2 The case of immediate rewards

Suppose that $T = 4$, and $\beta = 1/2$. An investor faces a sequence of projects whose current value is given by the net reward schedule $\mathbf{v} = (3, 5, 8, 13)$. We assume that the rewards are net of the costs, so the cost schedule is $\mathbf{c} = (0, 0, 0, 0)$. The investment can be undertaken only once. The utility of any self is given in (11.81).

We examine below the investment behavior of time consistents, sophisticates, and naifs, using Definitions 11.16, 11.17. Condition (11.82), which requires that self 4 must always find it profitable to undertake the activity, is satisfied, because the net reward in period 4 is strictly positive.

1. **Time consistents** ($\beta = \hat{\beta} = 1$): Such a type will choose Y in the time period that maximizes the current net reward. For each of the first three selves in periods $t = 1, 2, 3$, we have $v_t < v_4$ (or $U_t(t) < U_t(4)$), hence, the project is undertaken in period 4 (Definition 11.16). Thus, the decision taken by the four successive selves is $\{N, N, N, Y\}$; this is also the perception-perfect strategy. Time consistents optimally procrastinate to the last period.
2. **Sophisticates** ($\beta = \hat{\beta} < 1$): The solution can be found by backward induction as follows. Self 4 will choose Y . Self 3 will undertake the investment because $8 > 13/2$. Anticipating this, selves 2 and 1 will also, recursively, choose Y because, respectively, $5 > 8/2$ and $3 > 5/2$. Hence, the *perception-perfect strategy profile* of the sophisticates is $\{Y, Y, Y, Y\}$. Using Definition 11.17, in equilibrium, a sophisticate will undertake the investment in period 1. Notice that the sophisticate *preproperates* too much when rewards are immediate. Such an individual correctly foresees that each of the future selves will have a self-control problem, and not leave it to the last period, where the net payoff is the highest.
3. **Naifs** ($\beta < \hat{\beta} = 1$): In this case, any self t believes, incorrectly, that self $t + 1$ is time consistent. Hence, selves 1, 2, 3 believe that self 4 will actually undertake the investment, and maximize the net rewards from the activity. It is easily checked that selves 1 and 2 will choose N because, respectively, $3 < 6.5$ and $5 < 6.5$. However, self 3 will choose Y because $8 > 6.5$, thus, the beliefs of selves 1 and 2 are incorrect. The perception-perfect strategy profile of the naifs is $\{N, N, Y, Y\}$, hence, naifs will undertake the investment in period 3.

Unlike the case of immediate costs, lack of self-awareness leads now to a better outcome. Naifs do much better than sophisticates because the latter know that their immediate future selves will give in to temptation, and so each self gives in to temptation as well. In general, relative to the decision of the time consistent types, lack of self-awareness leads one to preproperate, i.e., undertake an activity too quickly when they should, from the perspective of a social planner, actually wait. These results are not specific to the two examples considered here.

One can make these points in more general one activity models; the following two propositions summarize the results. Denote by τ^i , $i = c, s, n$, the time period when it is optimal for a type i individual to undertake the activity. The first proposition shows that for immediate costs, naifs procrastinate more than time consistents, while for immediate rewards, naifs preproperate relative to the time consistents

Proposition 11.19 Consider $(\beta, 1)$ preferences, $0 < \beta < 1$, and time $t \in \{1, 2, \dots, T\}$. The reward and the cost schedules are given, respectively, by $\mathbf{v} = (v_1, v_2, \dots, v_T)$, and $\mathbf{c} = (c_1, c_2, \dots, c_T)$. In the case of immediate costs, the lifetime utility, $U_t(\tau)$, is given in (11.80),

and in the case of immediate rewards, it is given by (11.81). The definitions of time consistent and naif individuals are given, respectively, in (11.10) and (11.13).

- (a) For the case of immediate costs, time consistent complete the task no later than the naifs.
- (b) For the case of immediate rewards, naifs complete the task no later than the time consistent.

Proof: (a) We need to show that if the naifs do it in any period, the time consistent do it too. Consider any time period t . Let t' be the time period where $(v_\tau - c_\tau)$ is maximized for all $\tau > t$. Naifs believe that all their future selves are time consistent. Since $\delta = 1$, they believe that the activity will be undertaken in period t' . Naifs will undertake the activity in period t only if the current reward from undertaking it under immediate costs, $\beta v_t - c_t$, is greater than the utility from having it undertaken it at time t' , $\beta (v_{t'} - c_{t'})$, or

$$v_t - \beta^{-1}c_t \geq v_{t'} - c_{t'}. \quad (11.92)$$

If (11.92) holds for $0 < \beta < 1$, then it also holds for $\beta = 1$, hence, time consistent individuals ($\beta = 1$) will also do it.

(b) Consider the case of immediate rewards. Here we need to show that if time consistent undertake the activity, then naifs undertake it as well. Consider any time period t , and let t' be as in part (a). Time consistent do it if

$$v_t - c_t \geq v_{t'} - c_{t'}. \quad (11.93)$$

Naifs do it in period t if $v_t - \beta c_t \geq \beta (v_{t'} - c_{t'})$, or

$$\beta^{-1}v_t - c_t \geq v_{t'} - c_{t'}. \quad (11.94)$$

If (11.93) holds, then (11.94) holds too, hence, naifs do it as well. ■

We now compare the behavior of sophisticates and naifs for the general case.

Proposition 11.20 Consider $(\beta, 1)$ preferences, $0 < \beta < 1$, and time $1, 2, \dots, T$. The reward and the cost schedules are given, respectively, by $v = (v_1, v_2, \dots, v_T)$, and $c = (c_1, c_2, \dots, c_T)$. In the case of immediate costs, the lifetime utility, $U_t(\tau)$, is given in (11.80), and in the case of immediate rewards, it is given by (11.81). The definitions of sophisticated and naif individuals are given, respectively, in (11.11) and (11.13). For the case of immediate costs and immediate rewards, sophisticates complete the task no later than naifs.

Proof: We need to show here that if naifs undertake an activity in any time period, then so must the sophisticates. From the definition of perception-perfect strategies (see Definition 11.16), naifs do it in period t rather than in any future period if

$$U_t(t) > U_t(\tau) \forall \tau > t. \quad (11.95)$$

Sophisticates do it in period t rather than in any future period if (see Definition 11.17),

$$U_t(t) > U_t(\underline{\tau}) \text{ where } \underline{\tau} = \min_{\tau > t} \{ \tau \mid s_\tau^i = Y \}. \quad (11.96)$$

(11.95) implies that doing the activity in period t must dominate doing it for any future period. In particular it must also then dominate doing it from any subset of the set of future periods. But this implies that (11.96) holds, so the result follows immediately. ■

11.5.3 A brief note on welfare comparisons

Given any two distinct strategies for any of the types, which of these strategies lead to higher welfare for the individual? The Pareto dominance criterion implies that a Pareto dominant strategy must give higher welfare for all selves of the same individual. This is a fairly strong requirement. O'Donoghue and Rabin (1999b) suggest using a weaker efficiency criterion, which is the maximization of *long-run utility*, W_0 , from the perspective of some ex-ante period 0 that precedes period 1.

$$W_0^i(U_1^i, U_2^i, \dots, U_T^i) = \sum_{t=1}^{t=T} \delta^t U_t^i, 0 < \delta < 1, \quad (11.97)$$

where the superscript i denotes the type of the player, such as time consistent, sophisticate, and naïf.

11.5.4 Some other features of present bias

Individuals who have time consistent preferences do not violate two properties, namely, *dominance*, and *independence of irrelevant alternatives*. First, we define these terms. Suppose that there exists some period τ with $v_\tau > 0$, and $c_\tau = 0$. In the context of the one-activity model, an individual violates “dominance” if she performs the activity in some period $\tilde{\tau} \neq \tau$ with $c_{\tilde{\tau}} > 0$ and $v_{\tilde{\tau}} = 0$.

Definition 11.18 (*Dominance*): An intertemporal strategy¹³ dominates another if it gives at least the same payoff in all periods, and a strictly greater payoff in at least one period.

The independence of irrelevant options implies the following. Suppose option a_k is chosen out of a set of options $A = \{a_1, a_2, \dots, a_n\}$, then it would still be chosen out of the set A/a_j where $k \neq j$. For the one-activity model, this leads to the following definition.

Definition 11.19 Suppose that $\mathbf{v} = (v_1, v_2, \dots, v_T)$ and $\mathbf{c} = (c_1, c_2, \dots, c_T)$ give the intertemporal pattern of rewards and costs over time periods $1, 2, \dots, T$. Let the individual choose to undertake the activity in period τ . Then the individual's behavior is “independent of irrelevant alternatives” if she still chooses to undertake the activity in period τ once the rewards and costs for some period $\tilde{\tau} \neq \tau$ are omitted.

Proposition 11.21 (O'Donoghue and Rabin, 1999b): Consider the single activity model. Let individuals have (β, δ) preferences. The following results hold for sophisticated and naïve individuals:

1. There exist payoff and cost vectors $\mathbf{v}, \mathbf{c}$, and assumptions about immediacy (i.e., immediate costs, or immediate benefits) such that an individual with (β, δ) preferences will violate dominance.

¹³ An intertemporal strategy in the one-activity model is a sequence $\{s_t\}_{t=1}^T$ where $s_t \in \{Y, N\}$.

2. There exist payoff and cost vectors v, c , and assumptions about immediacy (i.e., immediate costs, or immediate benefits) such that an individual with (β, δ) preferences will violate independence of irrelevant alternatives.

We leave the proof of Proposition 11.21 as an exercise, but use the following two examples to illustrate it.

Example 11.18 (O'Donoghue and Rabin, 1999b): Consider the case of “immediate rewards.” The utility function is, therefore, given by (11.81). Suppose $T = 3$, $v = (0, 5, 1)$, and $c = (1, 8, 0)$. Suppose that $\beta = 1/2$. Using Definition 11.18, undertaking the activity in period 1 is dominated by doing it in period 3. Consider first a sophisticated individual. If each self were to undertake the activity when it is their turn to make the decision, then, using (11.81), the respective payoffs are

$$U_3(3) = 1 > 0; U_2(2) = 5 - \frac{8}{2} = 1 > 0; U_1(1) = 0 - \frac{1}{2} = -\frac{1}{2} < 0. \quad (11.98)$$

We now use backward induction to solve for the perception-perfect equilibrium. Self 3 undertakes the activity because $U_3(3) > 0$. Self 2 will also undertake the activity because

$$U_2(2) > U_2(3) \Leftrightarrow 1 > \frac{1}{2}. \quad (11.99)$$

Self 1 is sophisticated, so he knows that the earliest future time at which the activity will occur is at time $t = 2$. Thus, he also undertakes the activity because

$$U_1(1) > U_1(2) \Leftrightarrow -\frac{1}{2} > \frac{1}{2}(5 - 8) \Leftrightarrow -\frac{1}{2} > -\frac{3}{2}. \quad (11.100)$$

The perception-perfect strategy profile for the sophisticates is $(s_1^s, s_2^s, s_3^s) = (Y, Y, Y)$. Thus, self 1 undertakes the action in period 1, even though it is dominated by undertaking the activity in period 3 (check that $U_1(1) < U_1(3)$). This action violates dominance. The reason is that self 1 guesses, correctly, that self 2 will undertake the activity (on account of his present bias). If we were to truncate the payoff and cost vectors to remove period 2, then self 1 will leave the activity to be undertaken by self 3. This violates the independence of irrelevant alternatives.

Example 11.19 (O'Donoghue and Rabin, 1999b): Consider now the case of “immediate costs.” Thus, the utility function is given by (11.80). Suppose that individuals are naive (see (11.13)), and $\beta = 1/2$. Let $T = 3$, $v = (1, 8, 0)$, and $c = (0, 5, 1)$. Using Definition 11.18, undertaking the activity in period 3 is dominated by doing it in period 1. If each self were to undertake the activity when it is their turn to make the decision, then the respective payoffs are

$$U_3(3) = \frac{1}{2}(0) - 1 = -1 < 0; U_2(2) = \frac{8}{2} - 5 = -1 < 0; U_1(1) = \frac{1}{2} - 0 = \frac{1}{2} > 0. \quad (11.101)$$

A naive individual believes, incorrectly, that all future selves are time consistent ($\hat{\beta} = 1$). Hence, self 1 believes (incorrectly) that self 2 will undertake the activity because $U_2(2) = 8 - 5 = 3 > 0$. For self 1, $U_1(1) < U_1(2)$ because $\frac{1}{2} < \frac{3}{2}$. Hence, self 1 will not undertake the activity. Self 2 would rather leave the activity to be undertaken by the period 3 self (because $-1 < -\frac{1}{2}$). Because the activity must be undertaken once, it is undertaken by self 3, which is a dominated action.

O'Donoghue and Rabin (1999a) consider the possibility of moral hazard issues in a principal-agent setting. Suppose that the cost realization of an agent, who uses hyperbolic discounting, is privately known to the agent. The principal only knows the distribution of the agent's costs. What if the principal were to observe a delay in completing the task? The delay could have been caused by procrastination on the part of the agent, or it could have been an efficient response to a high cost realization. Hence, if the principal imposed penalties on task completion that are too strict, then procrastination could be reduced, but at the cost of undertaking the task when costs are inefficiently high. O'Donoghue and Rabin (1999a) explore this trade-off under different assumptions on the level of sophistication of the agent. Under certain conditions, the optimal incentive scheme is for delay to be punished leniently early on, but more severely upon the passing of some deadline date.

11.6 Interdependence, multiple activities, and competition

We have considered procrastination and preproperation within the context of a one-activity model and a single individual. How do the results alter when there are multiple individuals and multiple activities (possibly sequential)? We follow Brocas and Carillo (2001) in dealing with these questions.

Suppose that projects are of a sequential nature, i.e., they have a time dimension associated with successive, interlinked, stages. Then, a delay in the current period increases future costs because if a prerequisite stage is not followed, then the completion of subsequent stages may be affected. This increases the cost of delay, and reduces procrastination. Similar arguments might be able to explain why some individuals accept an overload of work. Such acceptance can perhaps be viewed as a commitment device to ensure that delay on account of procrastination is reduced.

Inducing competition among agents through incentive contracts in which the *winner takes all*, can also be an important anti-delay device. If agent A procrastinates, and agent B does not, then the likelihood that agent A gets any reward is reduced. This increases the cost of delay. Furthermore, competition between two agents can also be viewed as a coordination game. One agent could be willing to delay a project, provided the other does too. Thus, one might get multiple equilibria where projects are either completed too soon, or too late, depending on the beliefs. We formalize some of these ideas in this section.

Consider two researchers with (β, δ) preferences who wish to undertake a separate project. If any project is undertaken in period t , then there is an immediate cost c , and a future benefit b in period $t + 1$. There are multiple selves of the same individual in different periods, and we refer to the self in period t as self t . We consider two possibilities.

(1) Each researcher undertakes a single project, and the two projects are *independent*. Thus, the payoff to a researcher from a project does not depend on the completion of the project by the other researcher.

(2) The two projects undertaken by the two researchers are *dependent*. In particular, the researcher who completes the project earlier is declared the winner, and receives all the benefits.

Proposition 11.22 (*Independent Projects*) Consider a situation where each researcher can undertake independent and identical projects. If

$$c \in (\underline{c}, \bar{c}), \text{ where } \underline{c} = \beta\delta b \frac{1-\delta}{1-\beta\delta}, \text{ and } \bar{c} = \beta\delta b, \quad (11.102)$$

then, at each date, each researcher completes the project with a probability $\lambda^* \in (0, 1)$, where

$$\lambda^* = \frac{1 - \delta}{\delta(1 - \beta)} \left(\frac{\beta\delta b}{c} - 1 \right) \in (0, 1). \quad (11.103)$$

λ^* is increasing in β, b , and decreasing in c .

Proof: Self t with (β, δ) preferences prefers doing the activity now rather than never if

$$-c + \beta\delta b > 0 \Leftrightarrow c < \bar{c} = \beta\delta b. \quad (11.104)$$

Self t prefers the activity to be done at date $t + 1$ rather than at t if

$$-c + \beta\delta b < \beta\delta(-c + \delta b) \Leftrightarrow c > \underline{c} = \beta\delta b \frac{1 - \delta}{1 - \beta\delta}. \quad (11.105)$$

Suppose that (11.104), (11.105) simultaneously hold, i.e., $c \in (\underline{c}, \bar{c})$. We consider the time-invariant probability λ of undertaking the project at date t . From (11.104), (11.105) neither of $\lambda = 0$ or $\lambda = 1$ is an equilibrium. In equilibrium, self t knows that all future selves undertake the project with probability λ , hence, self t is indifferent between doing, and not-doing the project if

$$-c + \beta\delta b = [\beta\delta(-c + \delta b)] [\lambda + (1 - \lambda)\lambda\delta + (1 - \lambda)^2\lambda\delta^2 + \dots]. \quad (11.106)$$

Equation (11.106) holds when $\lambda = \lambda^*$, given in (11.103). The comparative statics follow by straightforward differentiation of (11.103). ■

The intuition behind Proposition 11.22 is as follows. From (11.104), the current self of a researcher desires that the project be completed, but from (11.105) he wants to procrastinate. Hence, the optimal solution is an intermediate solution in that he completes the project with some probability. The comparative static results show that if the researcher puts greater weight on future outcomes (high β), or if future benefits, b , are high, then he procrastinates less. An increase in current costs, c , has the converse impact.

We now consider the case of *competition* among the researchers in the following sense. The two researchers undertake the projects separately, and the winner-takes-all, i.e., the first to complete the project gets the entire future benefit, b . If both complete the project at the same time, then both reap the benefit, b .

Proposition 11.23 (Interdependent projects; Brocas and Carillo, 2001) *Under competition between the two researchers in a winner takes all contest:*

- (i) If $c \in [\underline{c}, c^*]$, then both researchers complete the project in period 1 with probability 1, where

$$c^* = \frac{\beta\delta^2 b}{(1 - \beta)(2 - \delta - 2\sqrt{1 - \delta}) + \delta}. \quad (11.107)$$

- (ii) If $c \in [c^*, \bar{c}]$, then the probability that the project is completed each period is given by $p \in [\mu, 1]$, where $\mu > \lambda^*$, λ^* is defined in (11.103), and $\underline{c}, \bar{c}$ are defined in (11.102).

When the costs are low enough (below c^*), then both researchers undertake the project immediately. Since the incentive to procrastinate is low for any researcher in this case, the other

researcher realizes that he will miss out on the prize completely if he procrastinates. This works mutually to ensure early completion of the project. When costs are in the intermediate range, procrastination turns into a coordination problem. Each researcher is willing to procrastinate, provided the other does too. But this depends on mutual beliefs, so, not surprisingly, there is a multitude of equilibria. All symmetric equilibria then yield the same expected payoff to each agent. Introducing competition among the researchers induces earlier completion of the projects, relative to the case of independent projects because of the contractually increased costs of procrastination.

11.7 Optimal sin taxes

Sin goods are goods whose consumption by any temporal self of an individual creates a negative externality for future selves. Time consistent selves take full account of the effects of current consumption on future utility. However, under hyperbolic preferences, the multiple selves are biased towards the present, and do not take full account of the negative externality that they create for future selves. Hence, there is a role for corrective *Pigouvian taxes* to assist in the internalization of a harmful externality. We explore these issues in a simplified version of the model provided in O'Donoghue and Rabin (2006).

Consider two goods, a *composite good*, and a *sin good* (e.g., potato chips). At time t , the consumption of the composite good is z_t , and the consumption of the sin good is x_t . Markets are competitive, and the marginal cost of each good is normalized to unity. Hence, the price per unit of each good is also unity. The instantaneous utility of self t in any period takes the separable quasilinear form, given by

$$u_t = v(x_t) - h(x_{t-1}) + z_t, \quad (11.108)$$

where,

$$v' > 0, v'' < 0, h' > 0, h'' > 0. \quad (11.109)$$

Individuals derive utility from both, x_t and z_t , but they also derive disutility from the past consumption of the sin good, x_{t-1} , that is encapsulated in the *habit function*, h . The quasilinear form guarantees that all income effects in any period are absorbed by z_t , and x_t is sensitive only to the substitution effect. No borrowing or saving is allowed.

The intertemporal preferences take the $(\beta, 1)$ form. In particular, for self $t = 1, 2, \dots, T - 1$, these preferences are given by

$$U_t = u_t + \beta \sum_{\tau=t+1}^T u_{\tau}; \quad 0 < \beta < 1. \quad (11.110)$$

The time-invariant fiscal policy of the government in any period is given by the tuple (θ, α) , where $\theta \in [0, 1]$ is the per unit consumption tax on the sin good, and α is a lump-sum transfer given by the government to the taxpayer. The taxpayer's income in each period is time invariant, and denoted by I . Since savings or borrowings are disallowed, the budget constraint of the individual requires that in each period, expenditure is bounded above by income, i.e.,

$$(1 + \theta) x_t + z_t \leq I + \alpha. \quad (11.111)$$

Substitute (11.108) in (11.110), and isolate terms corresponding to the choice variables of self t . We can then rewrite (11.110) as

$$U_t = u^P(x_t, z_t) + \beta \sum_{\tau=t+2}^T u_\tau + (\beta v(x_{t+1}) + \beta z_{t+1} - h(x_{t-1})), \quad (11.112)$$

where

$$u^P(x_t, z_t) = v(x_t) - \beta h(x_t) + z_t. \quad (11.113)$$

The first term in (11.112), shown in (11.113) as the *private utility* of self t , lists all the terms that are influenced by the choices of self t . The initial level of x_0 is exogenously given.

Using (11.111) to (11.113), we can treat the decision problem of the individual as a sequence of time t problems, where $t = 1, 2, 3, \dots, T-1$. In (11.113), $\beta < 1$ is the source of self-control problems. For $\beta = 1$, there are no self-control problems, and we get the *social utility* of self t ,

$$u^S(x_t, z_t) = v(x_t) - h(x_t) + z_t. \quad (11.114)$$

Using (11.109), u^P and u^S are strictly concave in x . Using the terminology in Kahneman (1994), the distinction between (11.113) and (11.114) is analogous to that between *decision utility* and *experienced utility*.¹⁴ Since the decision problem at time t is independent of any other time period, we will, henceforth, drop time subscripts.

11.7.1 First best allocation

The first best allocation maximizes (11.114), subject to the *resource constraint*, which obtains by eliminating distortionary policy instruments (i.e., set $\theta = \alpha = 0$ in (11.111)). Thus, the resource constraint is given by

$$x + z \leq I. \quad (11.115)$$

Since marginal utility of the composite good is strictly increasing, the resource constraint binds. Substituting the resource constraint into (11.114), the unconstrained first best problem is

$$x^S = \arg \max u^S(x, z(x)) = v(x) - h(x) + (I - x), 0 \leq x \leq I. \quad (11.116)$$

Since u^S is strictly concave in x , and $x \in [0, I]$, the solution is unique. Assuming an interior solution, the first best solution is given implicitly by

$$v'(x^S) - h'(x^S) = 1, \text{ and } z^S = I - x^S. \quad (11.117)$$

11.7.2 Privately optimal choice

The privately optimal choice of the individual at time t is found by substituting (11.111) in (11.113). The unconstrained problem of the individual is

¹⁴ Decision utility is concerned with choices made by an individual while experienced utility encapsulates the individual's ex-post welfare.

$$x^P \in \arg \max u^P = v(x) - \beta h(x) + (I + \alpha - (1 + \theta)x), 0 \leq x \leq \frac{I + \alpha}{1 + \theta}.$$

u^P is strictly concave in x , and x belongs to a compact interval, so the solution is unique. Assuming an interior solution, the unique, privately optimal solution is given implicitly by¹⁵

$$v'(x^P) - \beta h'(x^P) = 1 + \theta, \text{ and } z^P = I + \alpha - (1 + \theta)x^P. \quad (11.118)$$

Let the solution to (11.118) be given by $x^P(\theta, \beta)$ and $z^P(\theta, \alpha, \beta) = I + \alpha - (1 + \theta)x^P(\theta, \beta)$. It follows by implicit differentiation that

$$\frac{\partial x^P}{\partial \theta} < 0; \quad \frac{\partial x^P}{\partial \beta} < 0. \quad (11.119)$$

A reduction in the degree of impatience (an increase in β) reduces consumption of the sin good. A higher value of β allows the individual to internalize better the negative effect of habit formation on the self in the next period. With quasilinear preferences, there are no income effects on x . Hence, the substitution effect implies that consumption of the sin good is decreasing in its price, $1 + \theta$.

The use of Pigouvian taxes to counter overconsumption on account of self-control problems has also been considered in Gruber and Köszegi (2001). In actual practice, cigarette taxes are alleged to be regressive with respect to income. Gruber and Köszegi (2004) point to a countervailing factor. The tax elasticity of consumption of cigarettes is higher for poorer individuals, hence, such taxes allow poorer individuals with present bias to engage in relatively better self-control, partially offsetting the negative effects of a regressive tax.

11.7.3 Implementation of the first best

We can choose the tax rate, θ , in such a way that the first best is implemented in private markets; this is the Pigouvian tax solution. Choose $\theta = (1 - \beta)h'(x)$, and substitute in (11.118) to get $v'(x) - \beta h'(x) = 1 + (1 - \beta)h'(x)$. Simplifying we get,

$$v'(x) - h'(x) = 1,$$

which, using (11.117), holds, uniquely, for $x = x^S$. Hence, appropriately designed sin taxes are able to implement the first best.

11.7.4 Population heterogeneity

Suppose that the only source of heterogeneity in the population is the extent of the self-control problem given by the magnitude of $1 - \beta$. As β increases towards 1 the self-control problem diminishes (for $\beta = 1$ we are back in an EDU framework, where optimal plans are time consistent). Suppose that the distribution function of β is given by $F(\beta)$, and the support of the density is given by $[\underline{\beta}, \bar{\beta}]$. We refer to an individual with $(\beta, 1)$ preferences as a β -individual. The demand

¹⁵ A sufficient condition for an interior solution is $v'(0) - \beta h'(0) > 1 + \theta$.

for the sin good by an individual is given by $x^P(\theta, \beta)$, so that the aggregate per capita, or average demand, $X(\theta)$, is given by

$$X(\theta) = E_F [x^P(\theta, \beta)] = \int_{\underline{\beta}}^{\bar{\beta}} x^P(\theta, \beta) dF(\beta),$$

where E_F is the expectation operator with respect to the distribution of β . The average tax proceeds are returned to each consumer as a lump-sum transfer, $\alpha(\theta)$, hence,

$$\alpha(\theta) = \theta X(\theta). \quad (11.120)$$

Given quasilinear preferences, income effects on account of α are fully absorbed into the demand for the composite good, z . We are now primarily interested in the comparative static effects of the policy parameters, θ, α . Hence, we suppress β , and summarize the privately optimal bundle of goods demanded by the consumer as $(x^P(\theta), z^P(\theta, \alpha))$.

We are interested in a welfare analysis below. Hence, we use the social utility u^S in (11.114), rather than the private utility u^P in (11.113). Substituting the consumption bundle $(x^P(\theta), z^P(\theta, \alpha))$ in (11.114), gives the indirect (social) utility function of a β -individual, $u^S(x^P(\theta), z^P(\theta, \alpha) | \beta)$. Substituting α from (11.120) gives u^S as a function only of θ , conditional on a value of β , i.e., $u^S(x^P(\theta), z^P(\theta, \alpha(\theta)) | \beta)$. The utilitarian social welfare function of society that gives equal weight to the social utility of any β -individual is then given by

$$W(\theta) = E_F [u^S(x^P(\theta), z^P(\theta, \alpha(\theta)) | \beta)], \theta \in [0, 1]. \quad (11.121)$$

Substituting (11.114), (11.115) in (11.121), we can rewrite social welfare as

$$W(\theta) = E_F [v(x^P(\theta)) - h(x^P(\theta)) - x^P(\theta) + I | \beta] \equiv E_F [\omega(\theta) + I | \beta], \quad (11.122)$$

where

$$\omega(\theta) \equiv v(x^P(\theta)) - h(x^P(\theta)) - x^P(\theta) \quad (11.123)$$

is maximized at $x^P(\theta) = x^S(\theta)$ (see (11.116), (11.117)).

We now show that in the absence of self-control problems ($\beta = 1$) there are no sin taxes. However, if at least some individuals have self-control problems, then the optimal tax, θ , on the sin good is strictly positive. A positive sin tax reduces the harmful externality that current selves cause on future selves through a higher current consumption of the sin good.

Proposition 11.24 *Under alternative assumptions on β , the optimal sin taxes are as follows.*

- (a) *If all consumers have $\beta = 1$, then the optimal tax $\theta^* = 0$.*
- (b) *If all consumers have $0 < \beta \leq 1$, and some individuals have $\beta < 1$, then the optimal tax $\theta^* > 0$.*

Proof: Implicitly differentiating (11.123), $\frac{\partial \omega}{\partial \theta} = \frac{\partial \omega}{\partial x} \frac{\partial x^P}{\partial \theta} = [v'(x^P(\theta)) - h'(x^P(\theta)) - 1] \frac{\partial x^P}{\partial \theta}$. Using (11.118) we can rewrite this as

$$\frac{\partial \omega}{\partial \theta} = [\theta - (1 - \beta) h'(x^P(\theta))] \frac{\partial x^P}{\partial \theta}. \quad (11.124)$$

(a) Suppose first that all consumers have $\beta = 1$. Then, from (11.124), $\frac{\partial \omega}{\partial \theta} = \theta \frac{\partial x^P}{\partial \theta}$. Hence, differentiating (11.122) we get

$$W'(\theta) = E_F \left[\frac{\partial \omega}{\partial \theta} \right] = \theta E_F \left[\frac{\partial x^P}{\partial \theta} \right]. \quad (11.125)$$

$W(\theta)$ is continuous over the compact interval $[0, 1]$. Using (11.119), $\frac{\partial x^P}{\partial \theta} < 0$, so $W'(\theta) = 0$ for $\theta = 0$, and $W'(\theta) < 0$ for $\theta > 0$. Thus, the optimal tax $\theta^* = 0$.

(b) Suppose that $\beta \leq 1$ for all consumers, and $\beta < 1$ for some consumers. For all consumers such that $\beta = 1$ we have

$$\beta = 1 : \left(\frac{\partial \omega}{\partial \theta} \right)_{\beta=1} = \frac{\partial \omega}{\partial x} \frac{\partial x^P}{\partial \theta} = \theta \frac{\partial x^P}{\partial \theta}. \quad (11.126)$$

Using (11.124), for all consumers such that $\beta < 1$,

$$\beta < 1 : \left(\frac{\partial \omega}{\partial \theta} \right)_{\beta < 1} = \frac{\partial \omega}{\partial x} \frac{\partial x^P}{\partial \theta} = [\theta - (1 - \beta) h'(x^P(\theta))] \frac{\partial x^P}{\partial \theta}. \quad (11.127)$$

$W(\theta)$ is a continuous function over a compact interval $[0, 1]$. When $\theta = 0$, $\left(\frac{\partial \omega}{\partial \theta} \right)_{\beta=1} = 0$, but $\left(\frac{\partial \omega}{\partial \theta} \right)_{\beta < 1} > 0$ ($h' > 0$ and $\frac{\partial x^P}{\partial \theta} < 0$), thus, we have $W'(0) > 0$. At $\theta = 1$, a simple calculation shows that $W'(1) < 0$. Hence, at any local maximum, we must have $\theta^* > 0$. There is no guarantee, unless we make further assumptions, that this equilibrium is unique. ■

How does the optimal sin tax vary with the extent of self-control problems within the population. Suppose that we have two distributions of β , $F(\beta)$ and $G(\beta)$, such that $F(\beta) \leq G(\beta)$, hence, F first order stochastically dominates G . The extent of the self-control problem is given by the magnitude of $1 - \beta$. Under distribution G , there is a relatively greater proportion of the population with lower values of β . Hence, distribution G is worse (relative to F) from the point of self-control problems.

Under certain conditions, O'Donoghue and Rabin (2006) show that the optimal tax is higher under distribution G . The intuition is that the negative externality created for one's future selves is higher under distribution G , hence, the consumption of the sin goods needs to be curbed to a greater extent. Insofar as higher taxes reduce the negative externality on future selves, they provide a commitment device to an individual who cares about lifetime consumption. Such an individual might actually be more happy after the imposition of sin taxes. Gruber and Mulainathan (2005) verify this insight by showing a positive link between excise taxes on cigarettes and self-reported measures of happiness of smokers from the US and Canada. No such link is found between other kinds of taxes and self-reported measures of happiness of smokers.

Further issues of Pareto optimal sin taxes and the distributional impact of β on optimal sin taxes are taken up in one of the exercises below.

11.8 Investment and leisure goods

Let us focus attention on only two kinds of goods, *investment goods* and *leisure goods*. Consumption of investment goods incurs current costs, and generates future benefits, e.g., attendance at

a health club. Leisure goods, on the other hand, are associated with current benefits, but future costs, e.g., credit card borrowing. Suppose that individuals exhibit hyperbolic preferences, and are either *sophisticated*, or *naive*. Sellers of investment and leisure goods are assumed to maximize profits. We are interested in the optimal behavior of firms for investment and leisure goods when they know that consumers use hyperbolic preferences. We now consider this question in the context of the model by DellaVigna and Malmendier (2004).

11.8.1 Investment goods

Suppose that there are three time periods, described below.

1. Period 0 (Participation decision): A monopolistic firm can sell one unit of output to a consumer. It offers a two part tariff (L, p) to a consumer, where L is a lump-sum fee, and p is the price per unit of the good. The consumer can either accept or reject. If the consumer rejects the contract in period 0, then firms make zero profits, and the consumer gets some reservation utility $\bar{u}$ in period 1. The consumer knows that her future cost, c , of consuming the good is drawn from some density $f(c)$; the corresponding distribution function is $F(c)$. For instance, if the good is a health club membership, then the cost, c , could include items such as costs of commuting to the health club, rearranging one's schedule with a partner/child, and missing favorite live television programs.
2. Period 1 (Current costs): At the beginning of period 1, the consumer learns her cost of consuming the good. If the consumer accepts the contract in period 0, then she has to pay the lump-sum cost, L in period 1. Now, she decides whether to consume the one unit, or not. If the consumer does not consume, after accepting the contract, then she gets a payoff equal to $-L$. If the consumer chooses instead to consume, then costs in period 1 equal $-L - c - p$.
3. Period 2 (Future benefits): If the consumer consumes the good in period 1, then she derives a benefit $b > 0$ in period 2, and if she does not consume in period 1, she derives zero benefits from the good in period 2.

THE CONSUMER'S PROBLEM

We assume that each period, $t = 0, 1, 2$, corresponds to a different self of the consumer. The risk neutral consumer has (β, δ) preferences, so the preferences of self 0 are given by

$$u_0 + \beta\delta u_1 + \beta\delta^2 u_2, \quad (11.128)$$

where u_t is the instantaneous utility function of self t . The standard case of time consistent preferences arises by setting $\beta = 1$. Such a consumer is referred to as an *exponential consumer*. Self 0 expects self 1 to have the preferences

$$u_1 + \hat{\beta}\delta u_2, \quad (11.129)$$

where $\hat{\beta}$ is the predicted value of β . A *sophisticated* self knows that $\beta = \hat{\beta} < 1$, a *naive* self thinks that $\beta < \hat{\beta} = 1$, and a partial-naif believes that $\beta < \hat{\beta} < 1$. We may consider $\hat{\beta} - \beta \geq 0$ to be a measure of *consumer overconfidence*. Thus, individuals who follow EDU, and sophisticated selves, have no overconfidence, while partial naifs and naifs have strictly positive overconfidence.

We now work out the consumer's decisions in successive periods.

1. Period 0: Using (11.128), the lifetime utility of self 0 who pays the lump-sum payment, L , and chooses to consume in period 1 is

$$0 + \beta\delta(-p - c - L) + \beta\delta^2b = \beta\delta(\delta b - p - c - L). \quad (11.130)$$

By contrast, the lifetime utility from incurring the lump-sum payment, L , but choosing not to consume is

$$0 + \beta\delta(-L) + 0 = -\beta\delta L. \quad (11.131)$$

From (11.130), (11.131), self 0 would like to consume if

$$c \leq \delta b - p. \quad (11.132)$$

Self 0 does not know the realization of cost, c . But he does know that if the subsequent realization of cost is no higher than $\delta b - p$, then she will consume. Given that the distribution function for cost is given by $F(c)$, the following result is obvious.

Lemma 11.6 *In period 0, the probability that the consumer will consume in the future is $F(\delta b - p)$. If the firm charges a higher price, p , then the probability that the consumer will consume is lower.*

2. Period 1: Self 0 believes that the utility of self 1 is given by (11.129). So she believes that the lifetime utility of self 1 from consuming is $(-p - c - L) + \hat{\beta}\delta b$, and the utility from not consuming is $-L$. Hence, self 0 believes that self 1 will consume if

$$c \leq \hat{\beta}\delta b - p. \quad (11.133)$$

Thus, self 0 believes that the probability of consumption in period 1 is

$$F(\hat{\beta}\delta b - p). \quad (11.134)$$

Lemma 11.7 *(Overconfidence) A measure of self 0's overconfidence is given by the difference between the forecasted and the actual consumption probabilities, i.e., $F(\hat{\beta}\delta b - p) - F(\beta\delta b - p)$, which is always non-negative because $\beta \leq \hat{\beta}$. For EDU decision makers, $\beta = \hat{\beta} = 1$, and for sophisticated selves $\beta = \hat{\beta} < 1$, so there is no overconfidence. Naifs ($\beta < \hat{\beta} = 1$) have the greatest degree of overconfidence, followed by partial-naifs ($\beta < \hat{\beta} < 1$).*

The lifetime expected utility of self 0 who accepts the contract is given by

$$\beta\delta(-L) + F(\hat{\beta}\delta b - p)(\beta\delta(-p - c) + \beta\delta^2b). \quad (11.135)$$

The first term gives the present value, in period 0, of the lump-sum payment, L . With probability $F(\hat{\beta}\delta b - p)$, consumption takes place in period 1, which generates, in present value terms, a period 1 cost of $\beta\delta(-p - c)$, and a period 2 benefit of $\beta\delta^2b$. Self 0 has some reservation utility $\bar{u} \geq 0$ that materializes in period 1, so she will accept the contract if

$$\beta\delta \left[-L + \int_{-\infty}^{\hat{\beta}\delta b - p} (\delta b - p - c)f(c)dc \right] \geq \bar{u}. \quad (11.136)$$

The LHS of (11.136) is just a restatement of the expression in (11.135). We shall refer to (11.136) as the *consumer's participation constraint*.

THE FIRM'S PROBLEM

The monopolist firm makes a take it or leave it offer, (L, p) , to the consumer in period 0 by incurring some set up cost $K \geq 0$. If the contract is accepted, the firm receives a lump-sum fee, L , in period 1. If the consumer decides to exercise the right to consume the one unit in period 1, then the firm receives a price, p , and incurs a marginal cost of production, $\alpha > 0$. The firm can borrow and lend at an interest rate of $r > 0$, and its discount factor is given by $\frac{1}{1+r} = \delta$. The firm knows, at $t = 0$, the distribution of the consumer's cost, $F(c)$. The expected profit of the firm, from the consumption of the good, subject to the consumer's participation in period 0, is

$$\Pi(p) = \delta [L - K + F(\beta\delta b - p)(p - \alpha)]. \quad (11.137)$$

The firm receives all payments in period 1, hence, all payoffs are discounted by the discount factor, δ . If the contract is signed, the firm receives a lump-sum payment, net of the setup costs. With probability $F(\beta\delta b - p)$, consumption takes place, and the firm receives the per unit price, p , net of its marginal cost of production, α . In writing (11.137), we make the strong assumption that firms know the true value of β for the consumer, thus, we replace $\hat{\beta}$ by β .

The firm maximizes its expected profit in (11.137), subject to the participation constraint of the consumer in (11.136). The participation constraint binds in equilibrium because the firm can always improve expected profits by increasing L . Substituting L from the binding participation constraint into (11.137) we get

$$\Pi(p) = \int_{-\infty}^{\beta\delta b - p} (\delta b - \alpha - c)f(c)dc + \int_{\beta\delta b - p}^{\hat{\beta}\delta b - p} (\delta b - p - c)f(c)dc - K - \bar{u}. \quad (11.138)$$

The firm chooses its price, p , in order to maximize $\Pi(p)$ given in (11.138). L is chosen residually to ensure that the participation constraint of the consumer binds. The first term on the RHS of (11.138) is the expected social surplus (sum of firm's profits, and the consumer's utility). The second term in (11.138) is the fictitious consumer surplus that results from underestimating one's future self-control problems. From Lemma 11.7, we know that the second term equals zero for exponential and sophisticated consumers, but it is strictly positive for partial-naifs and naifs.

Applying Leibniz's rule¹⁶ we can write the first order condition for the firm's profit maximization problem as

$$p^* - \alpha = -(1 - \hat{\beta})\delta b \frac{f(\hat{\beta}\delta b - p^*)}{f(\beta\delta b - p^*)} - \frac{F(\hat{\beta}\delta b - p^*) - F(\beta\delta b - p^*)}{f(\beta\delta b - p^*)}, \quad (11.139)$$

where p^* is the optimal price; we assume that a unique interior solution exists.

¹⁶ Leibniz's rule is given as follows. For a function, g , whose partial derivatives exist, and are continuous, we have

$$\frac{d}{dx} \int_{a(x)}^{b(x)} g(x, t)dt = g(b(x), t)b'(x) - g(a(x), t)a'(x) + \int_{a(x)}^{b(x)} \frac{d}{dx} g(x, t)dt,$$

assuming the derivatives $a'(x)$ and $b'(x)$ exist.

For exponential consumers, $\beta = \hat{\beta} = 1$. Substituting in (11.139), we get the traditional optimality condition that price equals marginal cost, $p^* = \alpha$. This ensures that the consumer undertakes the consumption decision only when $c \leq \delta b - p^* \Leftrightarrow \delta b - c - \alpha \geq 0$, i.e., when the net surplus is positive.

When $\beta < 1$, the following possibilities arise. For sophisticates, $\beta = \hat{\beta} < 1$, hence, the RHS of (11.139) equals $-(1 - \beta)\delta b < 0$, thus, $p^* < \alpha$, i.e., price is lower than marginal cost. A simple calculation shows that for $\beta < 1$ and for partial-naifs or for naive consumers, the RHS of (11.139) is also negative, so $p^* < \alpha$. The lump-sum payment, L , is set to ensure that the participation constraint binds. From (11.136), there is an inverse relation between the lump-sum payment, L and the price, p .

What is the intuition for the $p^* < \alpha$ result? Two distinct rationales (depending on whether we have a sophisticate or a naive consumer) can be given.

1. **Commitment rationale:** Individuals who are aware that their future selves have present bias know that their period 1 selves are less likely to consume the good. This is because period 1 costs are more salient to them relative to the period 2 benefit. The reduction in price (below marginal cost) increases the probability that they will consume in period 1 (see Lemma 11.6). This is captured by the first term on the RHS in (11.139). In effect, the sophisticate consumer is willing to pay a larger lump-sum cost, L , in return for a commitment device from the firm (lower price per unit of usage) that enhances future consumption.
2. **Consumer overconfidence:** A partially or fully naive consumer overestimates (in period 0) consumption in period 1 (see Lemma 11.7). Thus, the firm offers the consumer a discount on the per unit price, p , but a higher lump-sum payment, L (to ensure that the participation constraint binds). However, because the consumer will eventually consume less than anticipated, her ex-post welfare is lower with a contract that offers a price discount, but a higher lump-sum fee. In turn, the firm's profits are higher under this misperception of the consumer. One of the problems at the end of Part 3 formalizes this issue. The consumer overconfidence effect is directly proportional to the overconfidence of the consumer (see the second term in (11.139) and Lemma 11.7), which reflects the overestimation of future consumption.

11.8.2 Leisure goods

Leisure goods, such as credit cards, are characterized by current benefits and future costs. A slight reinterpretation of the investment goods model allows us to incorporate leisure goods. In the case of leisure goods, in period 2, reinterpret $b < 0$ as the future cost, while in period 1, reinterpret $c < 0$ as the current benefit (or negative cost). Hence, the distribution $F(c)$ is over a negative domain. Let us now reinterpret the two terms on the RHS of (11.139). The first term is positive (because $b < 0$). The second term is also positive because the overconfidence term $F(\hat{\beta}\delta b - p^*) - F(\beta\delta b - p^*) < 0$. Hence, the RHS is positive, and we are led to the result that $p^* > \alpha$, i.e., for leisure goods, price is greater than marginal cost.

The intuition is as follows. Unless she is naive, the consumer knows that she has self-control problems in the future, so she knows that she will overconsume in the future. The optimal $p^* > \alpha$ allows the firm to offer a commitment strategy to the consumer to reduce future consumption. However, insofar as future self-control problems are underestimated by consumers, they present a profit opportunity for the firm.

For exponential consumers, $\beta = \hat{\beta} = 1$, so, from (11.139), $p^* = \alpha$, as in the case of investment goods. Such consumers know that they have no self-control problems, and they are correct in these expectations, so there is no need to give them an incentive to underconsume. Exactly the same pricing outcome obtains for $b = 0$. In this case, there is no future temptation in period 1; this aligns the incentives of self 0 and self 1.

11.8.3 *Empirical results*

The main implication of current-biased preferences is that for investment goods, price is below marginal cost, while the converse is true for leisure goods.

To test the prediction for investment goods, DellaVigna and Malmendier (2004) examined the data from health clubs in the Boston area. Their final sample included 67 companies, which together own 97 clubs. They find that three main types of contracts were offered.

- (i) A monthly contract with a lump-sum start-up fee of \$129, and an average (across clubs) monthly fee of \$55.
- (ii) An annual contract with a start-up fee of \$64, and an average annual fee of \$625. These contracts do not allow a pay-per-visit feature.
- (iii) A pay-per-visit contract was also offered with average fee per visit of about \$11.

The pay-per-visit contracts are least prevalent, and only 2 out of the 67 companies offer them. The most common contracts offer a monthly fee, followed by those that offer an annual fee. On average, health clubs price their attendance at least \$3–\$6 below marginal cost, which is consistent with the predictions of the hyperbolic discounting model.¹⁷

Interestingly, in the 1950s, health clubs mainly offered the pay-per-visit contract. By the 1970's, such contracts had largely disappeared. One possible interpretation is that firms have learned gradually, over time, to offer contracts that maximize their profits. If consumers underestimate their future self-control problems, then they anticipate visiting the health clubs far more than they actually will. Hence, they are willing to trade off higher lump-sum fees in return for lower per-visit-fees of fixed monthly/annual costs. Firms benefit by receiving the higher lump-sum fees, and a lower than expected health club usage.

The second implication of the analysis is that for leisure goods, price exceeds marginal cost. To test this, DellaVigna and Malmendier (2004) examine evidence from credit card companies. They find that the interest rate on outstanding balances charged by credit card companies exceeds the prime rate by as much as 10%, thus, supporting the predictions of theory.

11.9 Other applications of hyperbolic discounting

In this section, we give a very brief and condensed treatment of some of the other evidence on hyperbolic discounting; the interested reader is encouraged to pursue the original sources mentioned below. Glaeser et al. (2003) provide a novel explanation for why Americans have become increasingly obese. They argue that, for technological reasons, the time delay between the onset of the desire to have food, and the actual consumption of food has gone down. For individuals who use hyperbolic discounting, small time delays matter greatly, on account of the desire for

¹⁷ The authors show that their data rules out other potential explanations, such as price discrimination and transaction costs.

immediate gratification. So, in the presence of the culture of instant meals, individuals tend to consume a greater amount relative to their long-term interests.

In the EDU model, consumers should perfectly smooth consumption over different pay periods. However, consistent with hyperbolic discounting, a *payday effect* is observed. Consumption declines between successive paydays. Huffman and Barenstein (2004) find an 18% decline in consumption spending between the first week of the monthly pay period and the last. Similar declines in consumption over successive payday periods are found by Shapiro (2005) for recipients of food stamp payments and Stephens (2003) for social security recipients.

Paserman (2008) and DellaVigna and Paserman (2005) exploit the distinctive features of the job search process (immediate costs and delayed benefits) to distinguish between the exponential and the hyperbolic model. There are two distinct sorts of costs in the job search process. In the first category are immediate and direct costs such as finding information about vacancies, writing résumés, making contact with employers, and attending interviews. The second set of costs arise from comparing the long-term benefits and costs of accepting a job, and deciding if to continue to search. The first set of costs might be particularly onerous to job seekers who use hyperbolic discounting (which makes current costs more salient relative to exponential discounting). Thus, there could be suboptimal job search, which seems to be borne by the evidence. It follows that a longer spell of unemployment might not just be due to pure luck, but also due to inefficient job search choices made by present-biased individuals. Thus, labor market policy that promotes job search is potentially welfare improving. Using field data, Paserman (2008) finds a significant degree of present bias for low and medium wage workers.

A related problem is one of accepting the *default option* when searching for alternatives that involve current costs, but future benefits. Choi et al. (2004) find support for this idea in the context of acceptance of the default option for the 401(k) retirement savings plan in the US. Individuals might procrastinate and not search for a better or cheaper alternative, hence, they may simply accept the default option set by their employers. This, of course, has immediate implications for the role of public policy in ensuring sensible default options.¹⁸

An important problem in labor economics is the long-term dependence on welfare payments by welfare recipients. One possible explanation for the long-term take-up of such programs is that recipients of welfare, on account of their current bias, might not spend enough time searching for a job. Hence, if their choice set is constrained, for instance, by legislating time bound welfare programs, then such a policy could provide these individuals with a commitment device to quit welfare and go back to work. Fang and Silverman (2009) find evidence of such present bias among single mothers with children that prevented them from going to work. However, in Fang and Silverman (2004) they find that the actual design of these time bound welfare programs prevented single mothers from becoming better off.

Using field evidence, Ariely and Wertenbroch (2002) examine the effect of deadlines on completion of class assignments. In their experiments, executive-education students have to write three assignments. In the first treatment, students can choose their own deadlines, and in the second, evenly spaced deadlines are exogenously imposed on the students. If students have a bias towards the present, then deadlines may ensure that they complete tasks on time. It turns out that there is a demand for deadlines. In the first treatment, instead of choosing to submit the assignments at the end, students use self-imposed deadlines. This is consistent with the idea that many students prefer to use deadlines as a commitment device. However, self-imposed deadlines are

¹⁸ See Benartzi and Thaler (2004) for examples of sensible default policy choices. See Part 8 of the book for a more detailed discussion of these issues.

not successful in preventing delay to the same extent as exogenously imposed deadlines. Hence, students in the first treatment are possibly partially naive in that they underestimate their future self-control problems and do not impose an efficient set of deadlines on themselves.

Wertenbroch (1998) considers why smokers purchase single packs of cigarettes when they could get quantity discounts on larger packs? One explanation is that there is a commitment value to restricting oneself to a smaller number of cigarettes. Furthermore, he finds that consumers are willing to pay a price premium to restrict themselves to smaller packages of their favorite chocolate chip cookies or potato chips.

Ausubel (1999) looks at variations in the response of customers, over a period of 21 months, to a change in the pre-teaser and post-teaser rates of interest offered by a credit card company. The control group were offered a 6.9% annual interest rate for the first six months, and a rate of 16% thereafter. The pre-teaser treatment group was presented with the altered figures of 4.9% for the first six months, and 16% thereafter. The post-teaser treatment offered 6.9% for the first six months, and 14% thereafter. Clearly, the interest rates offered in the post-teaser treatment should be at least as attractive as those in the pre-teaser treatment. However, contrary to this intuition, the increase in the take-up rate for the pre-teaser treatment was 2.5 times that for the post-teaser treatment. Ausubel's interpretation is that individuals exhibit a degree of naiveté. They think that their future selves will behave more in line with the standard textbook models, and will borrow more responsibly in the future. Hence, they do not find the post-teaser treatment to be an attractive one. But individuals seem to mispredict, and ex-post would have been better off, on average, accepting the post-teaser treatment instead. See also Shui and Ausubel (2004) for an extension to a dynamic stochastic model that replicates quantitative features of the real-world data. They estimate that the present bias parameter of hyperbolic discounting is $\beta = 0.8$.

Standard public economics provides the basis to tax cigarette consumption. Smokers create externalities for non-smokers, so the tax induces them to internalize the external costs that they impose on society. Gruber and Köszegi (2001) argue that this analysis does not take account of the internal costs that smokers impose on themselves. On account of a present bias, the current consumption of cigarettes by smokers is far greater than what their own patient long-run self would desire. Hence, if the optimal cigarette tax were to take account of this additional cost, then it would be even higher, relative to current magnitudes; see Section 11.7, above, for a theoretical model.

Duflo et al. (2011) find that 97% of farmers in Kenya intend to use fertilizers in the following season, but only 37% actually end up using fertilizers; this is sometimes known as the *intention-action divide*. They apply a *stochastic version of hyperbolic preferences* to the issue of the take-up of fertilizer use. They allow a fraction of the farmers to have stochastically varying degrees of impatience in a $(\beta, 1)$ model. For instance, farmers could have two levels of impatience, β_L, β_H , with respective probabilities $1 - p, p$ such that $\beta_L < \beta_H \leq 1$. However, the current selves of farmers could be naive, and overestimate the value of p (i.e., believe that in the future there is a greater likelihood that they will be more patient).

There is a small, fixed, cost of purchasing fertilizers. Due to present-biased preferences, farmers may be put off by this fixed cost, and not purchase the fertilizer. They may think, incorrectly, that because their future selves have a high value of β , they are less likely to procrastinate in buying fertilizer. In the event, all future selves may not buy the fertilizer despite it being optimal to buy from an ex-ante perspective. A subsidy on the price of the fertilizer may be optimal under these circumstances. In the actual intervention, home delivery of fertilizers that eliminates travel costs improves fertilizer use by 70%.

The success of the Grameen Bank in Bangladesh in extending microcredit (small-scale loans) has typically been explained in terms of traditional agency theory. It has been argued that the terms of the contract (e.g., group lending) increase the incentives of borrowers to repay and take fewer risks. Bauer et al. (2012) provide an alternative explanation based on present-biased preferences. They conjecture that individuals with self-control problems may find it particularly attractive to borrow from microcredit institutions for the following two reasons.

First, loan officers hold group meetings with borrowers. Hence, borrowers who do not repay are, in effect, publicly named and shamed. Second, the repayments involve immediate weekly repayments of the loan amount even before the investment bears fruit. Thus, borrowers must dip into their current income every week to start repaying. Both features are very attractive to a current self who believes that future selves may have self-control problems. In order to test this conjecture, the authors conduct lab experiments in the field. The discount rates from a random sample of 573 villagers from 18 villages in India are elicited. The main finding is that, as conjectured, subjects who exhibit present-biased preferences in the lab are also more likely to select microcredit institutions to borrow money. However, this result holds only for women, a finding that is similar for the case of the Philippines by Ashraf et al. (2006).

Exercises for Part 3

1. Prove Corollary 1.
2. (Manzini and Mariotti, 2007) Suppose that the utility function is given by $u(z, \tau) = \delta^\tau z^\alpha$, $0 < \alpha < 1$; z, τ is an outcome–time pair, and δ is a discount factor. Consider two alternative prizes $z, z' : z < z'$. Suppose that a, b are the compensations required by the decision maker to delay the respective rewards, z, z' , by one time period. Show that EDU implies that $\frac{z}{z+a} = \frac{z'}{z'+b}$, while the magnitude effect implies that $\frac{z'}{z'+b} < \frac{z}{z+a}$.
3. Show that the magnitude effect is not explained by hyperbolic discounting.
4. (Manzini and Mariotti, 2007) Suppose that we have a time interval $[0, \tau_3]$ that we split into three subintervals $[0, \tau_1], [\tau_1, \tau_2], [\tau_2, \tau_3]$. Suppose that the observed pattern of choice is cyclic, i.e., $(z_2, \tau_2) \succ (z_3, \tau_3) \succ (z_1, \tau_1) \succ (z_2, \tau_2)$, where (z_i, τ_i) denotes a prize of z_i received at the end of the interval $[\tau_{i-1}, \tau_i]$. Suppose also that in comparing two outcome–time pairs, the decision maker splits aggregate time into subintervals such that the common subintervals cancel out. So, for instance, $(z_2, \tau_2) \succ (z_3, \tau_3)$ implies $u(z_2)\delta_0^{\tau_2} > u(z_3)\delta_0^{\tau_2}\delta_{\tau_2}^{\tau_3}$; where δ_t^τ is the discount factor over a time interval $[t, \tau]$, $t \leq \tau$. Show that the observed cyclic choice is consistent with subadditive discounting.
5. (al-Nowaihi and Dhami, 2008a) Consider the interval discount function in Definition 10.3. Show that a continuous discount function, $D(r, t)$, is continuous in t .
6. (al-Nowaihi and Dhami, 2008a) Consider the interval discount function in Definition 10.3.
 - (a) (Time sensitivity): Let D be a continuous discount function. Suppose $r \geq 0$. If $0 < x \leq y$, or if $y \leq x < 0$, then show that $v(x) = v(y)D(r, t)$ for some $t \in [r, \infty)$.
 - (b) (Existence of present values): Let D be a discount function (continuous, or otherwise). Let $r \leq t$. Let $y \geq 0$ ($y \leq 0$). Then, show for some x , $0 \leq x \leq y$ ($y \leq x \leq 0$), $v(x) = v(y)D(r, t)$.
7. (al-Nowaihi and Dhami, 2009) Show that the following classes of popular utility functions are incompatible with increasing elasticity of the value function which is required to explain the magnitude effect (see Proposition 10.3).
 - (a) Constant relative risk aversion (CRRA), given by

$$v(x) = \frac{x^{1-\gamma}}{1-\gamma}, 0 < \gamma < 1.$$

- (b) Hyperbolic absolute risk aversion (HARA), given by

$$v(x) = \frac{\gamma}{1-\gamma} \left[\left(\mu + \frac{\theta x}{\gamma} \right)^{1-\gamma} - \mu^{1-\gamma} \right], \theta > 0, \mu > 0, 0 < \gamma < 1, x \geq 0.$$

- (c) Constant absolute risk aversion (CARA), given by

$$v(x) = 1 - e^{-\theta x}, \theta > 0, x \geq 0.$$

- (d) Logarithmic functions, given by

$$v(x) = \ln(1 + \theta x), \theta > 0, x \geq 0.$$

(e) Quadratic functions, given by

$$v(x) = \frac{1}{2}\mu^2 - \frac{1}{2}(\mu - \theta x)^2, \theta > 0, 0 \leq x < \frac{\mu}{\theta}.$$

(Hint: Show that for members of the CRRA class of value functions, $\epsilon_v(x) = \frac{xv'(x)}{v(x)}$ is constant. For members of the HARA, CARA, logarithmic, and quadratic classes of functions, $\epsilon_v(x)$ is declining. Hence, conclude that none of these families is compatible with Proposition 10.3.).

8. Prove Proposition 10.4.
9. (al-Nowaihi and Dhami, 2008a) For the SIE utility function in (10.10), derive expressions for $\rho_A = -\frac{v''(x)}{v'(x)}$, and $\rho_R = -\frac{xv''(x)}{v'(x)}$. Show that $\rho_A > 0$ and $\rho_R > 0$ for $x > 0$ and $\rho_A < 0$ and $\rho_R < 0$ for $x < 0$.
10. Prove Proposition 10.8.
11. Prove Lemma 10.1.
12. Prove Proposition 10.9
13. (Saito, 2011) Consider the model of Saito (2011) discussed in section 10.3.3. Preferences are said to be *temporally unbiased* if for $x, y \in X, d, d', s, t \in \mathbb{R}_+$

$$(x, t) \succeq_d (y, s) \Leftrightarrow (x, t) \succeq_{d'} (y, s).$$

Show that $\succeq^r$ satisfies the independence axiom of expected utility if, and only if, $\succeq_d$ is temporally unbiased.

14. (al-Nowaihi and Dhami, 2014) Prove Propositions 10.11, 10.12.
15. (al-Nowaihi and Dhami, 2014) Derive the values for the delay, speedup, and generating functions shown in Table 10.2.
16. (al-Nowaihi and Dhami, 2014) Prove Proposition 10.13.
17. (al-Nowaihi and Dhami, 2014) Prove Proposition 10.14.
18. (al-Nowaihi and Dhami, 2014) Prove Proposition 10.15.
19. (al-Nowaihi and Dhami, 2008a) Consider Thaler's apples example (Example 9.1). Take the reference point to be the status quo, i.e., no apple. Use the prospect theory utility function with $\lambda = 2.25, \gamma = 0.88$ (see Part 1 for a definition of this function). Consider two discount functions: the LP-discount function, and the RS-discount function (see Table 10.1). Choose the parameters: $\alpha = \beta = \tau = \rho = 1$.
 - (a) Show numerically that both discount functions explain Thaler's example.
 - (b) Show numerically that the LP-discount function exhibits declining impatience, but the RS-discount function exhibits constant impatience.
 - (c) Show numerically that the LP-discount function is additive, but the RS-discount function is subadditive.
 - (d) What can you conclude about the relative channels through which the two discount functions explain Thaler's example?
20. Prove Proposition 11.18.
21. Prove Proposition 11.21.
22. Using present-biased preferences explain the following phenomena.
 - (a) In the UK, there has been a massive growth of pawnbrokers on the high streets of cities offering loans at astronomical interest rates.

(b) Many Americans buy furniture on a rent-to-own plan that costs, over its lifetime, three times more than directly buying the furniture.

(c) There has been recent spurt in the growth of *payday loans*. These are loans that are borrowed between one payday, and the next payday, using the next paycheck as collateral. According to O'Donoghue and Rabin (2007), there is an 18% interest rate over a month, or typically over two weeks on these loans (what is the implied annual interest rate?). In a typical transaction, a person borrows \$300, and pays back \$354 when the next paycheck arrives in a month/two weeks time. It is estimated that in 2002, 8.5% of Americans took out such loans.

23. (Kaur et al., 2015) Kaur et al. designed an intervention in an Indian data-entry company that took the form of a “negative bonus” scheme. Under this scheme, workers were paid normal piece rate if they met *self-chosen* data entry targets, but penalized if they did not. The take-up of the scheme was voluntary. About 35% of workers chose to set non-zero targets. The negative bonus contracts increased output by an amount that was equivalent to the effect of increasing their piece rate by 33%, or that which would be achieved by a worker with a year's further education. How might we explain these results? Can you think of other domains of employer–employee relations where such policies might work, e.g., in the domain of worker absenteeism?
24. How do you explain the reasoning behind the following quote from Datta and Mulainathan (2014)? “Clocky is an alarm clock that makes it easier to stick to plans to wake up early. When you hit its snooze button, Clocky jumps off the nightstand and rolls into a corner. When it goes off again, you have no choice, but to get up. Similarly, a non-profit organization called StickK allows people to stake their own money on things they are trying to get themselves to do, such as give up smoking. Their funds are only returned to them if they are certified as having met their goal by a third party. So far, nearly 150,000 “commitment contracts” have been taken out on www.StickK.com.”
25. (Diamond and Köszegi, 2003) Consider the model of retirement under hyperbolic discounting presented in Section 11.4. Imagine the following policy intervention. Suppose that the original employment legislation was compulsory retirement in period 2. What happens if this law is now repealed, and retirement in period 2 is a voluntary decision?
26. DellaVigna and Malmendier (2006) examine contracts signed by clients at health clubs. Among those who sign a monthly contract, the average cost per visit is \$17 over six months. If these individuals had instead chosen to pay per visit they would have paid \$10 per visit. Furthermore, the monthly contract has an automatic renewal feature such that the contract can only be canceled if the individual visits the health club to cancel the contract, or cancels by mail. There is a lag of about 2.29 months on average in canceling the contract since the last visit to the club. There is also a positive correlation between the cancellation lag in the monthly contract, and the average price paid per visit. Can these findings be explained under exponential discounting? How can we use insights from models of present-biased preferences to explain these findings?
27. (O'Donoghue and Rabin, 2007) Suppose that there is a single activity that an individual must undertake once during a 7 day period. The pattern of rewards, v , and costs c , is given in Table 3.E1.

The two possible cases are *immediate costs* (costs incurred now, rewards in the next period), and *immediate rewards* (rewards incurred now, costs in the next period). The utility functions corresponding to these two cases are given in 11.80, and 11.81.

Table 3.E1 Pattern of rewards and costs over seven days

Completion day	1	2	3	4	5	6	7
Rewards	5	9	17	33	33	33	33
Costs	5	5	5	5	9	17	33
Net benefit	0	4	12	28	24	16	0

Source: O'Donoghue and Rabin (2007).

- (a) Which day will an individual who uses exponential discounting choose to complete the activity in the two cases of immediate costs and immediate rewards?
- (b) Consider an individual who has $(\beta, \delta) = (0.5, 1)$ preferences.
- (b1) In the case of immediate costs, prove the following. (i) Naifs procrastinate till day 7. (ii) Sophisticates perform the activity on day 5.
- (b2) In the case of immediate rewards, prove the following. (i) Naifs complete the activity on day 3. (ii) Sophisticates perform the activity on day 1.
- (b3) Compare the behavior of naifs with exponential discounters. For the case of immediate costs, naifs do the activity relatively later than exponential discounters, and for the case of immediate rewards, the converse happens. Why?
- (c) Suppose that individuals were asked, prior to day 1, when they would commit to do an activity. Show that all types of individuals (exponentials, sophisticates, naifs) will prefer to commit to day 4 in both cases (immediate costs and immediate rewards). Would all types of individuals see value in taking up the opportunity to commit if commitment can be made at some cost? Would there be circumstances where giving an incentive to the individual to make a commitment may not be optimal?
28. Consider the problem of a social planner who wishes to choose the consumption profile $c_0, c_1, \dots, c_T$ to maximize

$$U_0(c_0, c_1, \dots, c_T) = \log(c_0) + \beta \sum_{t=1}^T \delta^t \log(c_t), \quad 0 < \beta \leq 1,$$

where $0 < T \leq \infty$, subject to the budget constraint

$$c_t + k_{t+1} = k_t^\alpha, \quad 0 < \alpha < 1, \quad (3E.1)$$

where k_t is the capital stock, and the production function is given by $y(k_t) = k_t^\alpha$. The initial level of capital, k_0 , is given.

- (a) Derive the optimal solution when $\beta = 1$. Is the solution time consistent?
- (b) Now let $\beta < 1$. Show that the optimal solution is given by

$$k_1 = \phi k_0^\alpha, \text{ and } k_{t+1} = \alpha \delta k_t^\alpha, \quad t \geq 1,$$

where ϕ is a scalar that depends on α, β, δ .

- (c) What is the steady state capital stock when $\beta = 1$, and $\beta < 1$? How does the steady state depend on the impatience parameter, β ?
- (d) In your answer to part (b), suppose that $(\beta, \delta) = (0.5, 0.975)$, and $\alpha = 0.5$. Find, numerically, the optimal values of the capital stock, $k_1^*, k_2^*, k_3^*, k_4^*$. Now consider the solution

to the problem of the individual from the perspective of time $t = 1$, $U_1(c_1, \dots, c_T)$, subject to the constraint in (3E.1), and given the new initial capital stock k_1^* . Find, numerically, the new optimal values $k_2^{**}, k_3^{**}, k_4^{**}$. Compare the values k_2^*, k_3^*, k_4^* with the values $k_2^{**}, k_3^{**}, k_4^{**}$ from the perspective of time consistency of optimal plans.

29. Consider a quasi-hyperbolic model of discounting where $(\beta, \delta) = (0.5, 1)$. Consider a special case of immediate costs in which time is discrete, $0, 1, 2, \dots$, the reward schedule from some project is $\mathbf{v} = (0, 0, \dots, 0)$, and the time t cost of doing the project at t is $(1.5)^t$. The project must be completed by some deadline time T .
- At which date will a naif complete the project?
 - At which date will a sophisticate complete the project? Does your answer depend on whether the deadline T is odd or even?
 - Now consider a partial naif who expects the impatience parameter of the future selves to be $\hat{\beta}$. Show that if $\hat{\beta} < 2/3$, then the project is completed in period $t = 0$. When is the project completed if $\hat{\beta} > 2/3$?
30. (O'Donoghue and Rabin, 2007) Suppose that a student has to read a total of P pages in T days. Let p_t be the number of pages read on day $t = 1, 2, \dots, T$. The student wishes to choose the optimal sequence $\{p_1, p_2, \dots, p_T\}$, such that $\sum_{t=1}^T p_t = P$. It takes h hours to read p pages, and the reading technology is given by $h_t = p_t^\alpha$, $\alpha > 1$. Reading is onerous, hence, the disutility from spending h hours on reading at time t is given by $u_t = -h_t$. Let Eve be an exponential discounter with $\delta = 1$, and let Neil be a naive individual with $(\beta, 1)$ preferences such that $\beta < 1$.
- For Eve, what is the optimal sequence of pages to read $\{p_1^E, p_2^E, \dots, p_T^E\}$?
 - Let $\mu = \beta^{\frac{1}{\alpha-1}}$. Show that on day 1, Neil optimally reads p_1^N pages, where

$$p_1^N = \frac{\mu}{T-1+\mu},$$

and on any day $t > 1$ he optimally reads p_t^N pages where

$$p_t^N = \left(\prod_{k=1}^{t-1} \frac{T-k}{T-k+\mu} \right) \frac{\mu}{T-t+\mu}.$$

(c) Suppose that $P = T = 30$ and $\alpha = 2$. Determine, numerically, the optimal hours put in by Eve and Neil on each day. What is the total number of hours put in by each individual over the 30 days. Who puts in more hours in the aggregate over the 30 days, and why? Give the intuition behind Neil's optimally chosen hours for the first three days.

(d) Suppose that a teacher wishes for Neil to study more efficiently. The teacher comes up with a requirement that one page must be read every day. What is the outcome? What insight does such an intervention give us about the nature of optimal interventions?

31. (O'Donoghue and Rabin, 2007) Consider a model of harmful habit formation. At any time t , the individual derives utility from the consumption, x_t , of a harmful good (e.g., cigarettes, unhealthy meals, alcohol), and disutility from the level of consumption at time $t-1$, x_{t-1} , which captures the role of habit formation. The unit price at time t of the harmful good is $p_t = 1$, for all t . The utility function at time t is given by

$$u_t = \begin{cases} 15 \ln(1 + x_t) - x_t & \text{if } x_{t-1} = 0 \\ 20 \ln(1 + x_t) - x_t - (20 + 10 \ln(1 + x_{t-1})) & \text{if } x_{t-1} > 0. \end{cases}$$

The utility function incorporates, in a reduced form, the following features. (a) Consumption of the harmful good reduces composite consumption (hence $p_t x_t \equiv x_t$ is subtracted from the RHS). (b) If there is no addiction, $x_{t-1} = 0$, then the marginal utility from consumption is relatively lower, as compared to the case of addiction, $x_{t-1} > 0$. The last term in u_t when $x_{t-1} > 0$, reflects the harmful effects of addiction on current utility. Consider two individuals. Eve is an exponential discounter with $\delta = 1$, and Neil is a naive quasi-hyperbolic discounter with $\beta = 0.8$ and $\delta \approx 1$.

- (a) Suppose that we consider steady state behavior in which one is either addicted, or unaddicted. Will Eve choose to be addicted, or unaddicted? Would your result change if Eve had started off being addicted in her youth (off the steady state). What can you conclude about the role of history in this case?
 - (b) Does Neil prefer the addicted steady state, or the unaddicted steady state?
 - (c) Consider now the role of history for Neil who is not making a choice between quitting now, or never. Rather he is making a choice between quitting now, or tomorrow.
 - (c1) Suppose that Neil starts off being unaddicted. Then does he choose to consume the harmful good?
 - (c2) Suppose that Neil starts off being addicted. Then does he choose to consume the harmful good?
 - (d) Why does history play different roles for Eve and Neil?
32. (Brocas and Carillo, 2001; Carillo and Dewatripont, 2008) Suppose that an individual with multiple selves, one for each period, lives for three time periods $t = 0, 1, 2$. He has a choice to perform tasks only in period $t = 1$. He can perform τ tasks such that $\tau = 0, 1, 2$. Each task is identical and involves an immediate effort cost e , and a one-period delayed reward that has a value equal to 1. The individual has $(\beta, 1)$ preferences and $e \in (\beta, 1)$.
- (a) Show that self 0 prefers self 1 to perform $\tau = 2$ tasks in preference to $\tau = 1$ tasks which in turn is preferred to $\tau = 0$ tasks. Also show that self 1 prefers just the opposite, i.e., no tasks to one task to two tasks.
 - (b) Now suppose that self 0 can make a promise, at some cost, to a third party about the number of tasks that self 1 will perform. Promises have the feature that breaking a promise is costly. Denote by $[x, y]$ the completion of $\tau = x$ tasks when self 0 promises that y tasks will be completed (so $x = 0, 1, 2$, and $y = 0, 1, 2$). Falling short of one task (i.e., $[1, 2]$, or $[0, 1]$) has cost c_1 , while falling short of two tasks (i.e., $[0, 2]$) has cost c_2 such that $c_2 > c_1$.
 - (b1) Suppose that $c_1 < e - \beta$, and $c_2 - c_1 < e - \beta$. Show in this case that self 0 optimally decides not to make any promises.
 - (b2) Suppose that $c_1 < e - \beta$, and $c_2 - c_1 > e - \beta$. Show in this case that self 1 will fall short of one task for any promise, but self 0 prefers $[1, 2]$ to $[0, 1]$. What condition is required for self 0 to prefer $[1, 2]$ to making no promises at all?
 - (b3) Suppose that $c_1 > e - \beta$, and $c_2 < 2(e - \beta)$. Show in this case that self 1 fulfills one promise, but breaks two promises. The period 0 self, thus, prefers a mild commitment to be honored rather than a stiff one be broken.
 - (b4) Suppose that $c_1 > e - \beta$, and $c_2 > 2(e - \beta)$. Show that in this case, all promises are honored, and the equilibrium outcome is $[2, 2]$.
33. Consider a three period model of consumption. A separate self lives in each period. The utility function of self 1 is given by

$$U_1(c_1, c_2, c_3) = \log(c_1) + \beta \log(c_2) + \beta^2 \log(c_3).$$

The budget constraint of the consumer at time t is given by

$$c_t + A_{t+1} = A_t + y_t,$$

where, at time t , A_t is the level of assets, and y_t is the labor income. The initial level of assets is $A_1 = 1$. The labor income at times 1, 2, and 3 is $y_1 = 0$, $y_2 \in \left[1, \frac{1}{\beta}\right]$, and $y_3 = 1$. At each time t , the consumer faces liquidity constraints, so $c_t \leq A_t$. Suppose that the individual is a sophisticate. Derive the optimal path of consumption. Does consumption track income too closely relative to the neoclassical model where $\beta = 1$?

34. (DellaVigna and Malmendier, 2004) Consider the model of DellaVigna and Malmendier (2004) outlined in Section 11.8.
- (a) (Perfect competition) Suppose that instead of a monopolist firm, contracts are offered in period 0 by perfectly competitive firms. Assume that increased competition only affects the utility $\bar{u}$ of the consumer in the alternative activity, and raises it. Show that the solution p^* is unaffected. What, if any, are the effects on the lump-sum fee, L .
- (b) (Welfare issues) Substitute L from the binding participation constraint and p^* from (11.139) to obtain the indirect utility function of the consumer, and the profit function of the firm. Show that profits are increasing in the degree of naiveté $\hat{\beta} - \beta$ of the consumer. Also show that partially naive consumers have lower indirect utility relative to sophisticated consumers.
35. (O'Donoghue and Rabin, 2006): Consider the model of sin taxes in Section 11.7.
- (a) Suppose we have two distributions of β , $F(\beta)$, and $G(\beta)$, such that $F(\beta) \leq G(\beta)$ for all β , and $F(\beta) < G(\beta)$ for at least one value of β . Hence, F first order stochastically dominates G . Let θ^F and θ^G be the optimal taxes under the distributions, $F(\beta)$ and $G(\beta)$, respectively. Show that if $\frac{\partial}{\partial \beta} \left(\frac{\partial \omega(\theta)}{\partial \theta} \right) < 0$ (where $\omega(\theta)$ is defined in the text in (11.123)), then $\theta^F > \theta^G$.
- (b) Keeping the notation as in Section 11.7, suppose that the long-run utility of a β -individual is given by

$$L(\theta) = [(v(x^P(\theta)) - c(x^P(\theta)) - x^P(\theta)) + (I + \alpha(\theta) - \theta x^P(\theta)) | \beta].$$

Denote by $\mathcal{B}$ the compact support of the distribution of β . Then a tax θ_1 is *Pareto-superior* to a tax θ_2 if $L(\theta_1) > L(\theta_2)$ for all $\beta \in \mathcal{B}$. Furthermore, a tax is *Pareto-efficient* if there is no Pareto-superior tax. Suppose now that $\beta \leq 1$ for all consumers, and $\beta < 1$ for some consumers. Further, for all $\beta \in \mathcal{B}$ and for all points in the domain of v ,

$$v''' - \beta c''' \geq \frac{2c''}{c'} (v'' - \beta c'').$$

Show that there exists $\hat{\theta} > 0$ such that all taxes in the open interval $(0, \hat{\theta})$ are Pareto-superior to $\theta = 0$.

36. (Harris and Laibson, 2013) Consider the (β, δ) model of quasi-hyperbolic discounting. Suppose that we extend it to continuous time as follows. For any self born at time, $t = 0$, called self 0, the “present” lasts from period $t_0 = 0$ to $t_1 = 0 + \tau_0$, and the “future” lasts from t_1 to $t = \infty$. The current self has control over the present, and the future selves have control over the future. Self 1 is born at time t_1 . For Self 1 the present period lasts from t_1 to

$t_2 = t_1 + \tau_1$, and the future period from t_2 to ∞ . We can now define future selves recursively. For Self n the present period lasts from t_n to $t_{n+1} = t_n + \tau_n$, and the future lasts from t_{n+1} to ∞ . Each of the terms in the sequence $\tau_0, \tau_1, \tau_2, \dots$ is stochastic, and is drawn from an exponential probability distribution with a hazard rate $\lambda \in [0, \infty)$. Having distinguished between the present and the future periods for any self, discounting over these periods is, in essence, similar to the (β, δ) model. Consider self n born at time t_n . Then the discount function is given by

$$D_n(t) = \begin{cases} \delta^t & \text{if } t \in [t_n, t_n + \tau_n) \\ \beta\delta^t & \text{if } t \in [t_n + \tau_n, \infty). \end{cases} \quad (3E.2)$$

- (a) What are the values of β and λ for which the continuous time hyperbolic function in (3E.2) collapses to the exponential discount function?
- (b) Plot this discount function for a particular realization of $\tau_0 = 3.4$ for self 0.

REFERENCES FOR PART 3

- Abdellaoui, M. A., Diecidue, E., and Öncüler, A. (2011). Risk preferences at different time periods: an experimental investigation. *Management Science* 57(5): 957–87.
- Ahumada, H. A. and Garegnani, M. L. (2007). Testing hyperbolic discounting in consumer decisions: evidence for Argentina. *Economics Letters* 95(1): 146–50.
- Ainslie, G. W. (1975). Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychological Bulletin* 82(4): 463–96.
- Ainslie, G. W. (1992). *Picoeconomics*. Cambridge: Cambridge University Press.
- Ainslie, G. W. and Haendel, V. (1983). The motives of the will. In E. Gotthelf, K. Druley, T. Skodola, and H. Waxman (eds.), *Etiology Aspects of Alcohol and Drug Abuse*. Springfield: IL Charles C. Thomas, pp. 119–40.
- Ainslie, G. W. and Herrnstein, R. J. (1981). Preference reversal and delayed reinforcement. *Animal Learning and Behavior* 9(4): 476–82.
- Akerlof, G. (1991). Procrastination and obedience. *American Economic Review* 81(2): 1–19.
- Alessi, S. M. and Petry, N. M. (2003). Pathological gambling severity is associated with impulsivity in a delay discounting procedure. *Behavioural Processes* 64(3): 345–54.
- al-Nowaihi, A. and Dhami, S. (2006). A note on the Loewenstein-Prelec theory of intertemporal choice. *Mathematical Social Sciences* 52(1): 99–108.
- al-Nowaihi, A. and Dhami, S. (2008a). A general theory of time discounting: the reference-time theory of intertemporal choice. University of Leicester, Department of Economics. Working Paper 08/22.
- al-Nowaihi, A. and Dhami, S. (2008b). A note on the Loewenstein-Prelec theory of intertemporal choice: Corrigendum. *Mathematical Social Sciences* 52(1): 99–108.
- al-Nowaihi, A. and Dhami, S. (2009). A value function that explains the magnitude and sign effects. *Economics Letters* 105(3): 224–9.
- al-Nowaihi, A. and Dhami, S. (2013). A theory of reference time. University of Leicester. Discussion Paper in Economics 13/26.
- al-Nowaihi, A. and Dhami, S. (2014). Foundations and properties of time discount functions. University of Leicester. Discussion Paper in Economics.
- Anderhub, V., Güth, W., Gneezy, U., and Sonsino, D. (2001). On the interaction of risk and time preferences: an experimental study. *German Economic Review* 2(3): 239–53.
- Andersen, S., Cox, J. C., Harrison, G. W., et al. (2012). Asset integration and attitudes to risk: theory and evidence. mimeo.
- Andersen, S., Harrison, G., Morten, L., and Rutstroem, E. (2011). Discounting behavior: a reconsideration. University of Durham. Working Paper in Economics.
- Anderson, L. and Stafford, S. L. (2009). Individual decision-making experiments with risk and intertemporal choice. *Journal of Risk and Uncertainty* 38(1): 51–72.
- Angeletos, G.-M., Laibson, D., Repetto, A., Tobacman, J., and Weinberg, S. (2001). The hyperbolic consumption model: calibration, simulation and empirical evaluation. *Journal of Economic Perspectives* 15(3): 47–68.
- Ariely, D. and Wertenbroch, K. (2002). Procrastination, deadlines, and performance: self-control by precommitment. *Psychological Science* 13(3): 219–24.
- Ashraf, N., Karlan, D., and Yin, W. (2006). Tying Odysseus to the mast: evidence from a commitment savings product in the Philippines. *Quarterly Journal of Economics* 121(2): 635–72.
- Ausubel, L. (1999). Adverse selection in the credit card market. University of Maryland. Working Paper.
- Baker, F., Johnson, M. W., and Bickel, W. K. (2003). Delay discounting in current and never-before cigarette smokers: similarities and differences across commodity, sign, and magnitude. *Journal of Abnormal Psychology* 112(3): 382–92.
- Baucells, M. and Heukamp, F. H. (2010). Common ratio using delay. *Theory and Decision* 68(1–2): 149–58.
- Bauer, M., Chytilová, J., and Morduch, J. (2012). Behavioral foundations of micro-credit: experimental and survey evidence from rural India. *American Economic Review* 102(2): 1118–39.
- Benartzi, S. and Thaler, R. H. (2004). Save more tomorrow (TM): using behavioral economics to

- increase employee saving. *Journal of Political Economy* 112(S1): S164–S187.
- Benhabib, J., Bisin, A., and Schotter, A. (2010). Present-bias, quasi-hyperbolic discounting, and fixed costs. *Games and Economic Behavior* 69(2): 205–23.
- Benjamin, D. J., Brown, S. A., and Shapiro, J. M. (2006). Who is “behavioural”? Cognitive ability and anomalous preferences. Unpublished.
- Benzion, U., Rapoport, A., and Yagil, J. (1989). Discount rates inferred from decisions: an experimental study. *Management Science* 35(3): 270–84.
- Bernheim, B., Skinner, D., Skinner, J., and Weinberg, S. (1997). What accounts for the variation in retirement wealth among U.S. households? NBER. Working Paper 6227.
- Bickel, W. K., Odum, A. L., and Madden, G. J. (1999). Impulsivity and cigarette smoking: delay discounting in current, never, and ex-smokers. *Psychopharmacology* 146(4): 447–54.
- Bjork, J. M., Hommer, D. W., Grant, S. J., and Danube, C. (2004). Impulsivity in abstinent alcohol dependent patients: relation to control subjects and type 1/type 2-like traits. *Alcohol* 34(2–3): 133–50.
- Blanchard, T. C., Wolfe, L. S., Vlaev, I., Winston, J. S., and Hayden, B. Y. (2014). Biases in preferences for sequences of outcomes in monkeys. *Cognition* 130(3): 289–99.
- Bleichrodt, H., Rohde, K. I. M., and Wakker, P. P. (2009). Non-hyperbolic time inconsistency. *Games and Economic Behavior* 66(1): 27–38.
- Bretteville-Jensen, A. L. (1999). Addiction and discounting. *Journal of Health Economics* 18(4): 393–407.
- Brocas, I. and Carrillo, J. D. (2001). Rush and procrastination under hyperbolic discounting and interdependent activities. *Journal of Risk and Uncertainty* 22(2): 141–64.
- Brown, A. L., Chua, Z. E., and Camerer, C. F. (2009). Learning and visceral temptation in dynamic saving experiments. *Quarterly Journal of Economics* 124(1): 197–231.
- Browning, M. and Lusardi, A. (1996). Household saving: micro theories and micro facts. University of Copenhagen, Department of Economics. Discussion Paper 96-01.
- Carrillo, J. D. and Dewatripont, M. (2008). Promises, promises, . . . *Economic Journal* 118(531): 1453–73.
- Chabris, C. F., Laibson, D., Morris, C. L., Schuldt, J. P., and Taubinsky, D. (2008). Individual laboratory-measured discount rates predict field behavior. *Journal of Risk and Uncertainty* 37(2–3): 237–69.
- Chabris, C. F., Laibson, D. I., and Schuldt, J. P. (2007). Intertemporal choice. In S. Durlauf and L. Blume (eds.), *The New Palgrave Dictionary of Economics* 2nd edition. London: Palgrave Macmillan.
- Chapman, G. B. (1996). Expectations and preferences for sequences of health and money. *Organizational Behavior and Human Decision Processes* 67(1): 59–75.
- Chapman, G. B. (2000). Preferences for improving and declining sequences of health outcomes. *Journal of Behavioral Decision Making* 13(2): 203–18.
- Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2004). For better or for worse: default effects and 401(k) savings behavior. In: D. Wise (ed.), *Perspectives in the Economics of Aging*. Chicago: University of Chicago Press, pp. 81–121.
- Chung, S. H. and Herrnstein, R. J. (1967). Choice and delay of reinforcement. *Journal of the Experimental Analysis of Behavior* 10(1): 67–74.
- Coble, K. and Lusk, J. (2010). At the nexus of risk and time preferences: an experimental investigation. *Journal of Risk and Uncertainty* 41(1): 67–79.
- Coffey, S. F., Gudleski, G. D., Saladin, M. E., and Brady, K. T. (2003). Impulsivity and rapid discounting of delayed hypothetical rewards in cocaine-dependent individuals. *Experimental and Clinical Psychopharmacology* 11(1): 18–25.
- Cropper, M. L., Aydede, S. K., and Portney, P. R. (1992). Rates of time preference for saving lives. *American Economic Review* 82(2): 469–72.
- Dasgupta, P. and Maskin, E. (2005). Uncertainty and hyperbolic discounting. *American Economic Review* 95(4): 1290–99.
- Datta, S. and Mullainathan, S. (2014). Behavioral design: a new approach to development policy. *Review of Income and Wealth* 60(1): 7–35.
- DellaVigna, S. and Malmendier, U. (2004). Contract design and self-control: theory and evidence. *Quarterly Journal of Economics* 119(2): 353–402.
- DellaVigna, S. and Malmendier, U. (2006). Paying not to go to the gym. *American Economic Review* 96(3): 694–719.

- DellaVigna, S. and Paserman, M. D. (2005). Job search and impatience. *Journal of Labor Economics* 23(3): 437–66.
- Diamond, P. A. and Hausman, J. (1984). Individual retirement and savings behavior. *Journal of Public Economics* 23(1–2): 81–114.
- Diamond, P. and Koszegi, B. (2003). Quasi-hyperbolic discounting and retirement. *Journal of Public Economics* 87(9–10): 1839–72.
- Dixon, M. R., Marley, J., and Jacobs, E. A. (2003). Delay discounting by pathological gamblers. *Journal of Applied Behavior Analysis* 36(4): 449–58.
- Duflo, E., Kremer, M., and Robinson, J. (2011). Nudging farmers to use fertilizer: theory and experimental evidence from Kenya. *American Economic Review* 101(6): 2350–90.
- Eichhorn, W. (1978). *Functional Equations in Economics*. Boston: Addison-Wesley.
- Elster, J. (1979). *Ulysses and the Sirens: Studies in Rationality and Irrationality*. Cambridge: Cambridge University Press.
- Epper, T., Fehr-Duda, H., and Bruhin, A. (2011). Viewing the future through a warped lens: why uncertainty generates hyperbolic discounting. *Journal of Risk and Uncertainty* 43(3): 169–203.
- Fafchamps, M., Kebede, B., and Zizzo, D. J. (2014). Keep up with the winners: experimental evidence on risk taking, asset integration, and peer effects. University of East Anglia, CBESS. Discussion Paper 14–03.
- Fang, H. M. and Silverman, D. (2004). On the compassion of time-limited welfare programs. *Journal of Public Economics* 88(7–8): 1445–70.
- Fang, H. M. and Silverman, D. (2009). Time-inconsistency and welfare program participation: evidence from the NLSY. *International Economic Review* 50(4): 1043–77.
- Fang, H. M. and Wang, Y. (2010). Estimating dynamic discrete choice models with hyperbolic discounting, with an application to mammography decisions. NBER. Working Paper 16438.
- Fishburn, P. C. (1991). Nontransitive preferences in decision theory. *Journal of Risk and Uncertainty* 4(2): 113–34.
- Fishburn, P. C. and Rubinstein, A. (1982). Time preference. *International Economic Review* 23(3): 677–94.
- Fisher, I. (1930). *The Theory of Interest*. New York: Macmillan.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives* 19(4): 25–42.
- Frederick, S., Loewenstein, G., and O'Donoghue, T. (2002). Time discounting and time preferences: a critical review. *Journal of Economic Literature* 40(2): 351–401.
- Fudenberg, D. and Levine, D. (2009). Self-control, risk aversion, and the Allais paradox. Washington University in St. Louis. Working Paper.
- Gigliotti, G. and Sopher, B. (1997). Violations of present value maximization in income choice. *Theory and Decision* 43(1): 45–69.
- Gintis, H. (2009). *The Bounds of Reason*. Princeton, NJ: Princeton University Press.
- Glaeser, E. L., Cutler, D. M., and Shapiro, J. M. (2003). Why have Americans become more obese? *Journal of Economic Perspectives* 17(3): 93–118.
- Green, L., Fischer, E. B., Jr., Perlow, S., and Sherman, L. (1981). Preference reversal and self control: choice as a function of reward amount and delay. *Behavioural Analysis Letters* 1(1): 43–51.
- Green, L., Fry, A. F., and Myerson, J. (1994). Discounting of delayed rewards: a lifespan comparison. *Psychological Science* 5(1): 33–6.
- Green, L. and Myerson, J. (1996). Exponential versus hyperbolic discounting of delayed outcomes: risk and waiting times. *American Zoologist* 36(4): 496–505.
- Green, L., Myerson, J., Lichtman, D., Rosen, S., and Fry, A. (1996). Temporal discounting in choice between delayed rewards: the role of age and income. *Psychology and Aging* 11(1): 79–84.
- Green, L., Myerson, J., and Ostaszewski, P. (1999). Amount of reward has opposite effects on the discounting of delayed and probabilistic outcomes. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 25(2): 418–27.
- Gruber, J. H. and Koszegi, B. (2001). Is addiction “rational”? Theory and evidence. *Quarterly Journal of Economics* 116(4): 1261–303.
- Gruber, J. H. and Koszegi, B. (2004). Tax incidence when individuals are time-inconsistent: the case of cigarette excise taxes. *Journal of Public Economics* 88(9–10): 1959–87.
- Gruber, J. H. and Mullainathan, S. (2005). Do cigarette taxes make smokers happier? *B.E. Journals: Advances in Economic Analysis and Policy* 5(1): 1–45.

- Guysé, J. L., Keller, L. R., and Epple, T. (2002). Valuing environmental outcomes: preferences for constant or improving sequences. *Organizational Behavior and Human Decision Processes* 87(2): 253–77.
- Halevy, Y. (2008). Strotz meets Allais: diminishing impatience and the certainty effect. *American Economic Review* 98(3): 1145–62.
- Hall, R. E. (1988). Intertemporal substitution in consumption. *Journal of Political Economy* 96(2): 339–57.
- Harris, C. and Laibson, D. (2001a). Dynamic choices of hyperbolic consumers. *Econometrica* 69(4): 935–57.
- Harris, C. and Laibson, D. (2001b). Hyperbolic discounting and consumption. Mimeo, University of Cambridge.
- Harris, C. and Laibson, D. (2013). Instantaneous gratification. *Quarterly Journal of Economics* 128(1): 205–48.
- Harvey, C. M. (1986). Value functions for infinite-period planning. *Management Science* 32(9): 1123–39.
- Henderson, N. and Bateman, I. (1995). Empirical and public choice evidence for hyperbolic social discount rates and the implications for intergenerational discounting. *Environmental and Resource Economics* 5(4): 413–23.
- Hsee, C. K., Abelson, R. P., and Salovey, P. (1991). Relative weighting of position and velocity in satisfaction. *Psychological Science* 2(4): 263–6.
- Huffman, D. and Barenstein, M. (2004). A monthly struggle for self-control? Hyperbolic discounting, mental accounting, and the fall in expenditure between paydays. IZA. Discussion Paper 1430.
- Ikeda, S., Kang, M.-I., and Ohtake, F. (2010). Hyperbolic discounting, the sign effect, and the body mass index. *Journal of Health Economics* 29(2): 268–84.
- Judd, K. (2003). Existence, uniqueness, and computational theory for time consistent equilibria: a hyperbolic discounting example. Mimeo, Stanford University.
- Kahneman, D. (1994). New challenges to the rationality assumption. *Journal of Institutional and Theoretical Economics* 150(1): 18–36.
- Kahneman, D. and Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica* 47(2): 263–91.
- Kahneman, D. and Tversky, A. (2000). *Choices, Values and Frames*. Cambridge: Cambridge University Press.
- Kaur, S., Kremer, M., and Mullainathan, S. (2015). Self-control at work. forthcoming in *Journal of Political Economy*.
- Keren, G. and Roelofsma, P. (1995). Immediacy and certainty in intertemporal choice. *Organizational Behavior and Human Decision Processes* 63(3): 287–97.
- Kim, B. K. and Zauberman, G. (2009). Perception of anticipatory time in temporal discounting. *Journal of Neuroscience, Psychology, and Economics* 2(2): 91–101.
- Kinari, Y., Ohtake, F., and Tsutsui, Y. (2009). Time discounting: declining impatience and interval effect. *Journal of Risk and Uncertainty* 39(1): 87–112.
- Kirby, K. N. (1997). Bidding on the future: evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology: General* 126(1): 54–70.
- Kirby, K. N. and Herrnstein, R. J. (1995). Preference reversals due to myopic discounting of delayed reward. *Psychological Science* 6(2): 83–9.
- Kirby, K. N. and Petry, N. M. (2004). Heroin and cocaine abusers have higher discount rates for delayed rewards than alcoholics or non-drug-using controls. *Addiction* 99(4): 461–71.
- Kirby, K. N., Petry, N. M., and Bickel, W. K. (1999). Heroin addicts have higher discount rates for delayed rewards than non-drug-using controls. *Journal of Experimental Psychology* 128(1): 78–87.
- Kirby, K. N., Winston, G. C., and Santisteban, M. (2005). Impatience and grades: delay-discount rates correlate negatively with college GPA. *Learning and Individual Differences* 15(3): 213–22.
- Koopmans, T. C. (1960). Stationary ordinal utility and impatience. *Econometrica* 28(2): 287–309.
- Koopmans, T. C., Diamond, P. A., and Williamson, R. E. (1964). Stationary utility and time perspective. *Econometrica* 32(1): 82–100.
- Kosse, F. and Pfeiffer, F. (2013). Quasi-hyperbolic time preferences and their inter-generational transmission. IZA. Discussion Paper 7221.
- Kotlikoff, L. J., Spivak, A., and Summers, L. H. (1982). The adequacy of savings. *American Economic Review* 72(5): 1056–69.

- Laibson, D. (1994). Essays in hyperbolic discounting. MIT, Ph.D. Dissertation.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 112(2): 443–78.
- Laibson, D. (1998). Life-cycle consumption and hyperbolic discount functions. *European Economic Review* 42(3–5): 861–71.
- Laibson, D., Repetto, A., and Tobacman, J. (1998). Self-control and saving for retirement. *Brookings Papers on Economic Activity* 29(1): 91–196.
- Laibson, D., Repetto, A., and Tobacman, J. (2005). Estimating discount functions with consumption choices over the lifecycle. Mimeo, Harvard University.
- Loewenstein, G. F. (1987). Anticipation and the valuation of delayed consumption. *Economic Journal* 97(387): 666–84.
- Loewenstein, G. F. (1988). Frames of mind in intertemporal choice. *Management Science* 34(2): 200–14.
- Loewenstein, G. F. (2007). Exotic Preferences: *Behavioral Economics and Human Motivation*. New York: Oxford University Press.
- Loewenstein, G. F. and Prelec, D. (1992). Anomalies in intertemporal choice: evidence and an interpretation. *Quarterly Journal of Economics* 107(2): 573–97.
- Loewenstein, G. F. and Prelec, D. (1993). Preferences for sequences of outcomes. *Psychological Review* 100(1): 91–108.
- Loewenstein, G. F. and Sicherman, N. (1991). Do workers prefer increasing wage profiles? *Journal of Labor Economics* 9(1): 67–84.
- Logue, A. W. (1988). Research on self-control: an integrating framework. *Behavioral and Brain Sciences* 11(4): 665–79.
- Lusardi, A. (1998). On the importance of the precautionary saving motive. *American Economic Review* 88(2): 449–53.
- McAlvanah, P. (2010). Subadditivity, patience, and utility: the effects of dividing time intervals. *Journal of Economic Behavior and Organization* 76(2): 325–37.
- Malkoc, S. A. and Zauberman, G. (2006). Deferring versus expediting consumption: the effect of outcome concreteness on sensitivity to time horizon. *Journal of Marketing Research* 43(4): 618–27.
- Manzini, P. and Mariotti, M. (2006). A vague theory of choice over time. *Advances in Theoretical Economics* 6(1): 1–27.
- Manzini, P. and Mariotti, M. (2007). Sequentially rationalizable choice. *American Economic Review* 97(5): 1824–39.
- Margolis, H. (1982). *Selfishness, Altruism and Rationality: A Theory of Social Change*. Cambridge: Cambridge University Press.
- Mayer, T. and Russell, T. (2005). Income smoothing and self-control: the case of schoolteachers. *Economic Inquiry* 43(4): 823–30.
- Mischel, W., Grusec, J., and Masters, J. C. (1969). Effects of expected delay time on subjective value of rewards and punishments. *Journal of Personality and Social Psychology* 11(4): 363–73.
- Mitchell, S. H. (1999). Measures of impulsivity in cigarette smokers and non-smokers. *Psychopharmacology* 146(4): 455–64.
- Noor, J. (2007). Hyperbolic discounting and the standard model. Mimeo, Boston University.
- O'Donoghue, T. and Rabin, M. (1999a). Incentives for procrastinators. *Quarterly Journal of Economics* 114(3): 769–816.
- O'Donoghue, T. and Rabin, M. (1999b). Doing it now or later. *American Economic Review* 89(1): 103–24.
- O'Donoghue, T. and Rabin, M. (2001). Choice and procrastination. *Quarterly Journal of Economics* 116(1): 121–60.
- O'Donoghue, T. and Rabin, M. (2003). Self awareness and self control. R. Baumeister, G. Loewenstein, and D. Read (eds.), *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice*. New York: Russell Sage Foundation, pp. 217–43.
- O'Donoghue, T. and Rabin, M. (2006). Optimal sin taxes. *Journal of Public Economics* 90(10–11): 1825–49.
- O'Donoghue, T. and Rabin, M. (2007). Incentives and self control. R. Blundell, W. Newey, and T. Persson (eds.), *Advances in Economics and Econometrics: Volume 2, Theory and Applications (Ninth World Congress)*. Cambridge: Cambridge University Press, pp. 215–45.
- Ok, E. A. and Masatlioglu, Y. (2007). A theory of (relative) discounting. *Journal of Economic Theory* 137(1): 214–45.
- Paserman, M. D. (2008). Job search and hyperbolic discounting: structural estimation and policy evaluation. *Economic Journal* 118(531): 1418–52.
- Petry, N. M. (2001). Pathological gamblers, with and without substance use disorders, discount

- delayed rewards at high rates. *Journal of Abnormal Psychology* 110(3): 482–7.
- Petry, N. M. (2003). Discounting of money, health, and freedom in substance abusers and controls. *Drug and Alcohol Dependence* 71(2): 133–41.
- Petry, N. M. and Casarella, T. (1999). Excessive discounting of delayed rewards in substance abusers with gambling problems. *Drug and Alcohol Dependence* 56(1): 25–32.
- Phelps, E. and Pollak, R. A. (1968). On second best national savings and game equilibrium growth. *Review of Economic Studies* 35(2): 185–99.
- Pollak, R. A. (1968). Consistent planning. *Review of Economic Studies* 35(2): 201–8.
- Prelec, D. (1998). The probability weighting function. *Econometrica* 66(3): 497–528.
- Read, D. (2001). Is time-discounting hyperbolic or subadditive? *Journal of Risk and Uncertainty* 23(1): 5–32.
- Read, D. and Read, N. L. (2004). Time discounting over the lifespan. *Organizational Behavior and Human Decision Processes* 94(1): 22–32.
- Read, D. and Roelofsma, P. H. M. P. (2003). Sub-additive versus hyperbolic discounting: a comparison of choice and matching. *Organizational Behavior and Human Decision Processes* 91(2): 140–53.
- Reynolds, B., Richards, J. B., Horn, K., and Karraker, K. (2004). Delay discounting and probability discounting as related to cigarette smoking status in adults. *Behavioral Processes* 65(1): 35–42.
- Roelofsma, P. H. M. P. and Read, D. (2000). Intransitive intertemporal choice. *Journal of Behavioral Decision Making* 13(2): 161–77.
- Rubinstein, A. (1988). Similarity and decision making under risk (is there a utility theory resolution to the Allais paradox?). *Journal of Economic Theory* 46(1): 145–53.
- Rubinstein, A. (2003). “Economics and Psychology”? The case of hyperbolic discounting. *International Economic Review* 44(4): 1207–16.
- Saito, K. (2011). A relationship between risk and time preferences. Northwestern University CMS-EMS Working Paper 1477.
- Salois, M. J. and Moss, C. B. (2011). A direct test of hyperbolic discounting using market asset data. *Economics Letters* 112(3): 290–2.
- Samuelson, P. A. (1937). A note on measurement of utility. *Review of Economic Studies* 4(2): 155–61.
- Schelling, T. C. (1984). Self-command in practice, in policy, and in a theory of rational choice. *American Economic Review* 74(2): 1–11.
- Scholten, M. and Read, D. (2006a). Discounting by intervals: a generalized model of intertemporal choice. *Management Science* 52(9): 1424–36.
- Scholten, M. and Read, D. (2006b). Beyond discounting: the tradeoff model of intertemporal choice. London School of Economics and Political Science. Working Paper 22710.
- Shapiro, J. M. (2005). Is there a daily discount rate? Evidence from the food stamp nutrition cycle. *Journal of Public Economics* 89(2–3): 303–25.
- Shefrin, H. M. and Thaler, R. H. (1988). The behavioral life-cycle hypothesis. *Economic Inquiry* 26(4): 609–43.
- Shelley, M. K. (1993). Outcome signs, question frames and discount rates. *Management Science* 39(7): 806–15.
- Shui, H. and Ausubel, L. M. (2004). Time inconsistency in the credit card market (May 3, 2004). 14th Annual Utah Winter Finance Conference. Available at: <http://ssrn.com/abstract=586622>.
- Silva, F. J. and Gross, T. F. (2004). The rich get richer: students’ discounting of hypothetical delayed rewards and real effortful extra credit. *Psychonomic Bulletin and Review* 11(6): 1124–8.
- Stephens, M., Jr. (2003). “3rd of the month”: do social security recipients smooth consumption between checks? *American Economic Review* 93(1): 406–22.
- Strotz, R. H. (1955–1956). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* 23(3): 165–80.
- Tanaka, T., Camerer, C. F., and Nguyen, Q. (2010). Risk and time preferences: linking experimental and household survey data from Vietnam. *American Economic Review* 100(1): 557–71.
- Thaler, R. H. (1981). Some empirical evidence of dynamic inconsistency. *Economics Letters* 8(3): 201–7.
- Thaler, R. H. (1990). Anomalies: saving, fungibility and mental accounts. *Journal of Economic Perspectives* 4(1): 193–205.
- Thaler, R. H. and Shefrin, H. M. (1981). An economic theory of self-control. *Journal of Political Economy* 89(2): 392–406.
- Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4): 297–323.

- Tversky, A. and Koehler, D. J. (1994). Support theory: a nonextensional representation of subjective probability. *Psychological Review* 101(4): 547–67.
- Tversky, A., Slovic, P., and Kahneman, D. (1990). The causes of preference reversal. *American Economic Review* 80(1): 204–17.
- Varey, C. and Kahneman, D. (1992). Experiences extended across time: evaluation of moments and episodes. *Journal of Behavioral Decision Making* 5(3): 169–85.
- Varian, H. (1992). *Microeconomic Analysis*. 3rd edition. New York: W.W. Norton and Company.
- Viscusi, K. W. and Huber, J. (2006). Hyperbolic discounting of public goods. NBER. Working Paper 11935.
- Viscusi, K. W., Huber, J., and Bell, J. (2008). Estimating discount rates for environmental quality from utility-based choice experiments. *Journal of Risk and Uncertainty*. 37(2–3): 199–220.
- Vuchinich, R. E. and Simpson, C. A. (1998). Hyperbolic temporal discounting in social drinkers and problem drinkers. *Experimental and Clinical Psychopharmacology* 6(3): 292–305.
- Weber, B. J. and Chapman, G. B. (2005). The combined effects of risk and time on choice: does uncertainty eliminate the immediacy effect? Does delay eliminate the certainty effect? *Organizational Behavior and Human Decision Processes* 96(2): 104–18.
- Wertenbroch, K. (1998). Consumption self-control by rationing purchase quantities of virtue and vice. *Marketing Science* 17(4): 317–37.
- Zauberman, G., Kim, B. K., Malkoc, S. A., and Bettman, J. R. (2009). Discounting time and time discounting: subjective time perception and intertemporal preferences. *Journal of Marketing Research* 46(4): 543–56.

PART 4

Behavioral Game Theory

INTRODUCTION TO PART 4

In Part 1 we studied the behavior of humans when they engage in non-strategic interaction against nature (game against nature). By contrast, in Part 4 of the book, we consider multiperson strategic decision problems of individuals, firms, or organizations, in which nature is one of the possible players. In these problems, moves by nature, the choices of the players, and their beliefs interact to give rise to individual payoffs. It is assumed that readers are familiar, at a bare minimum, with the material in the appendix on game theory.

Classical game theory plays a hugely influential and important role in economics. It has forced economists to think carefully and formally about the players in their models, their strategies, the sequence of moves, and the payoffs. It has developed plausible equilibrium concepts in different strategic situations. This has allowed it to make precise, testable predictions. It is probably fair to say that the basic machinery of classical game theory underpins most results in modern economics in a mathematically elegant and parsimonious manner.

Part 4 is split into two chapters. Chapter 12 describes the evidence on human choices in strategic situations. The main emphasis is on experimental and field evidence that is designed to test the predictions of various equilibrium concepts in classical game theory. Chapter 13 considers models of behavioral game theory that are motivated by the evidence in Chapter 12. We also consider the relevant evidence on these models. These models range from those that are relatively well-established and backed by reasonably solid evidence to others that are promising, yet need further testing to enable us to make a more informed choice.

Chapter 12 is divided into several sections. Generally, in the early rounds of games and in unfamiliar games, human behavior is often not consistent with the predictions of classical game theory. It typically takes a substantial amount of experience and repeated play of the game before play converges towards the predictions of classical game theory, but this is often not the case in many environments (see also Part 5). Humans engage in many important decisions for only a limited number of times in their lives. For instance, when they choose a university degree, a marriage partner, a house, a consumer durable, or a pension plan. Many decisions made by firms are also of an infrequent nature. These include: capital restructuring, mergers, sunk costs in major machinery and equipment, and the choice of a new product. Even in decisions that are taken frequently, the economic environment is ever-changing and uncertain, bringing about an element of novelty in most decisions. For all these reasons, it is arguably the choices made in the early rounds of an experiment that may be critical in understanding strategic human behavior in many important domains.

We begin in Section 12.2 by looking at the evidence from the simplest games in which a unique Nash equilibrium is found by *iteratively deleting dominated strategies* in a small number of steps. Yet, the outcome in most experiments differs from Nash equilibrium if the number of steps is more than two or three.

In a fundamental result, Nash (1950a) showed that every finite game has at least one *mixed strategy Nash equilibrium* (MSE); Section 12.3 shows that the evidence on the play of equilibrium mixed strategies by individuals is mixed. In the majority of experiments, individual play does not conform to the predictions of an MSE, but the aggregate play is closer to an MSE. However, a fair number of experiments also reject MSE both at the individual and the aggregate levels.

Individuals do play mixed strategies but, typically, not in the proportions that are predicted by an MSE. There is too much serial correlation between strategies over successive rounds. Even when players are given a randomizing device, their play still suffers from serial correlation. There is evidence from professional sport that tennis and soccer players may choose an MSE, but this ability is not portable to other contexts.

In Section 12.4, we consider coordination games that have *multiple equilibria* and the aim is to see if players can coordinate on any of the equilibria. We examine the role of *focal points* and the conflict between focal points and *payoff-salience*. None of the existing selection principles such as *payoff dominance* or *risk dominance* can fully account for the data. Newer, behavioral selection principles, based on loss aversion may be promising. The chosen equilibria are also strongly *history dependent*. *Outside options*, *forward induction*, and *preplay communication* influence the degree of coordination, even in cases where classical game theory predicts that there should be no such effect.

Bargaining games are considered in Section 12.5. We consider the neoclassical theoretical framework and the relevant empirical evidence for *normative* and *positive sequential alternating offers models of bargaining*. We also consider models that do not fit into any of these categories, such as models of *self-serving bias* and *focal points*. The support for classical game theory under full information is relatively weak, and under *one-sided asymmetric information* it is, at best, mixed. The evidence does not appear to be consistent with models where players have purely self-regarding, risk-neutral, homogeneous preferences. However, under *two-sided asymmetric information*, which requires very high cognitive demands from the players, we find surprising support for classical game theory under *sealed bid mechanisms*. We also consider evidence from experiments that attempt to discover the *cognitive process* underlying choices in games using MOUSE-LAB experiments. Unless subjects are trained in the relevant game theoretic concepts, their actual pattern of *searches* and *lookups* is not consistent with the expected pattern from classical game theory.

Section 12.6 examines the evidence from *signaling games*. In several experiments, while play does not conform to the predictions of classical game theory in the early rounds of play, there is greater conformity in later rounds. But, typically, a significant fraction of the players still do not play in accordance with the predictions. A much replicated finding is that there is no support for refinements beyond the Cho–Kreps *intuitive criterion*.

Section 12.7 considers the experimental evidence on *correlated equilibria* in which a choreographer gives recommendations to the players to play particular strategies. In equilibrium, each player finds it in his interest to follow the recommendations, if others do so. When no non-binding preplay communication is allowed, then experimental results show that the recommendations are not followed. This is reversed in the presence of preplay communication. The main result is that the existence of a correlated equilibrium is a necessary, but not a sufficient condition for the recommendations to be followed.

Section 12.8 highlights experimental evidence from games in which the actions of players are either *strategic complements* (SC) or *strategic substitutes* (SS). The evidence suggests that the actions of players in both cases are positively correlated. This is suggestive of reciprocity considerations among players. We also consider the implications of SC and SS for macroeconomics. In the presence of such considerations, nominal rigidity can be explained by behavioral factors such as *money illusion* and *anchoring*.

Section 12.9 explores empirical evidence for a Walrasian competitive equilibrium, where strategic considerations, in the limit, vanish. The main message of this research is that there exists an institution (double auction) and a particular rule for changing prices (the market improvement

rule) that can implement a competitive equilibrium in relatively few rounds of play, even with a small number of players.

The evidence on human strategic choices in Chapter 12 has motivated the development of alternative theoretical models, which are often in better conformity with the empirical evidence as compared to classical game theory. Chapter 13 turns to a consideration of these models. We now outline the plan of this chapter.

A *perfect Bayesian Nash equilibrium* (PBNE) requires two conditions. First, given any history of the game, each player chooses his best reply, conditional on his beliefs at that stage of the game and the subsequent equilibrium strategies of the players. Second, beliefs are formed using Bayes' rule, wherever possible, and actions and beliefs are consistent with each other. The models considered in Chapter 13 relax these conditions individually or jointly.

Section 13.2 is devoted to the *quantal response equilibrium* (QRE). This relaxes the first condition of a PBNE. Players play *noisy best replies*, in the sense that all actions are played with a strictly positive probability but, conditional on the beliefs, the likelihood of playing actions that give a higher expected payoff is also greater. A QRE requires the cognitively demanding condition that players have the correct beliefs about the noisy play of others. An equilibrium requires the consistency of such beliefs with the actions of players, so the second condition of a PBNE holds. A QRE is able to organize much better the evidence from a range of games as compared to a Nash equilibrium.

Section 13.3 relaxes the second feature of a PBNE but keeps the first. Arguably this leads to a less cognitively challenging model. *Level- k models*, and the *cognitive hierarchy models* (CH), are the leading behavioral models in this category. In these models, there is a hierarchy of types among the players. Each type is more cognitively accomplished, relative to the lower types. The two models differ in the distributional assumptions on the lower types. Level- k models assume that each player believes that the other player is of the immediately lower type, while in CH models, a player of any type believes that there is a non-degenerate distribution of lower types. Such models are able to make precise predictions in a large variety of games that are consistent with the evidence. In many of these games, a Nash equilibrium is unable to account for the data; perhaps the *p-beauty contest* is the most popular example.

Section 13.4 describes several applications of level- k and CH models with particular emphasis on coordination games. The *market entry game*, considered in this section, was traditionally held up as a vindication of classical game theory because most experimental evidence was consistent with it (Daniel Kahneman described these results as *magical*). However, it turns out that level- k and CH models are able to account for this evidence equally well.

In classical game theory, players do not derive utility directly from the beliefs about others (*first order beliefs*) or beliefs about the beliefs of others (*second order beliefs*). Yet, human morality, and a range of emotions, such as guilt, and notions of justice, such as perceived fairness, seem to impart direct utility, based on second order beliefs. Since second order beliefs are endogenous, one cannot simply append the psychological payoffs to the material payoffs and use the framework of classical game theory. Rather, this requires new concepts and machinery that is embodied in *psychological game theory*. We consider the relevant equilibrium concepts for both normal form and extensive form games in Section 13.5. We also consider applications that include an explicit modeling of the role of *reciprocity* (see Part 2 of the book for the evidence), and to the role of *guilt* in a range of games.

Section 13.6 is devoted to some of the central ideas in Gintis (2009) that are sufficiently important to merit a section in their own right. He argues that there are no plausible *epistemic foundations* that imply *common knowledge* of rationality in classical game theory. Gintis also ascribes to

social norms and *conventions* the role of a *choreographer* in a *correlated equilibrium* that coordinates human behavior. He also cautions against necessarily interpreting Nash equilibria as social norms because many social norms, such as certain norms of cooperation, are often not Nash equilibria.

Section 13.7 explores a range of other behavioral models of game theory. Each of these models establishes a promising new direction, although the extent of psychological underpinnings of these models, and the evidence for them, varies. Equilibrium concepts based on these models, that we consider, include: *analogy-based equilibrium*, *subjective heterogenous quantal response equilibrium*, *evidential equilibrium*, *cursed equilibrium*, and *noisy introspection*.

The final section, Section 13.8, considers the behavioral economics of auctions. Particular attention is given to *common value auctions* that are characterized by a *winner's curse*. We consider the relevant theoretical explanations and the empirical evidence.

CHAPTER 12

The Evidence on Strategic Human Choice

12.1 Introduction

This chapter documents an extensive, and growing, body of experimental and field evidence that puts to the test the predictions of classical game theory. This evidence has enabled us to develop a better understanding of human behavior in strategic situations, and has guided developments in the emerging field of behavioral game theory that we consider in Chapter 13.

Of particular interest is the class of games which lead to a unique outcome when one iteratively deletes dominated strategies of players. Suppose that the unique outcome of a game can be found by successively deleting dominated strategies in j steps. Then the game is said to be solvable with j steps of iterated deletion of dominated strategies. In classical game theory, j can be any finite number. Section 12.2 tests this prediction for normal form and extensive form games, when $j = 2, 3$. These are some of the simplest games that one could use in order to test the predictions of classical game theory.

The evidence on normal form games with two to three steps of iterated dominance is considered in Sections 12.2.1 and 12.2.2. For $j = 1$, the predictions of classical game theory usually hold. However, when $j = 2$, a majority of the players violate the predictions. One possibility is that as j increases beyond 1, players may have lesser confidence in the rationality of others to engage in iterated deletion of dominated strategies. There are two problems with such a-priori arguments. First, if a player is rational, why does he believe that his opponent is irrational? Second, if opponents are believed to be irrational in such exceedingly simple games, then there are hardly any games where we can invoke common or even mutual knowledge of rationality. This severely restricts the scope of classical game theory. A range of other factors such as *resentment*, *assurance*, and *reciprocity*, may also be important in such games (Beard and Beil, 1994). When $j = 2$, extensive form games and their equivalent normal form representations elicit different behavioral responses. For instance, the frame (extensive or normal) may potentially influence the ability of people to spot dominated strategies (Schotter et al., 1994).

In Section 12.2.3, we consider iterated deletion of dominated strategies in extensive form games of perfect information. A particularly tractable and popular game in this class is the *centipede game*. Depending on the number of nodes in the centipede game, one can test for $j = 1, 2, 3, \dots$ steps of iterated deletion of dominated strategies. For centipede games that involve four to eight alternating moves by players, the evidence is not consistent with the subgame perfect outcome (McKelvey and Palfrey, 1992). In these games, players need more than three steps of iterated

thinking to solve the game, so it is not surprising that the cognitive demands on players are quite high.

One may test iterated dominance using special subject pools that are likely to be familiar with the logic of such reasoning; for instance chess players (Palacios-Huerta and Volij, 2008). In particular, when chess grandmasters play against other grandmasters, the authors report 100% conformity with the backward induction outcome of the centipede game. On the face of it, this is striking support for classical game theory. However, in carefully conducted experiments with professional chess players of high ability, Levitt et al. (2011) showed that 0% of the chess grandmasters conformed with the predictions of classical game theory. Furthermore, their performance was no better than the performance of the students.

In Section 12.2.4, we consider two very simple mechanisms, the Abreu–Matsushima and the Glazer–Perry mechanisms that are dominance solvable, respectively, in simultaneous and sequential games. However, the experimental evidence shows that the performance of these simple mechanisms is unsatisfactory (Sefton and Yavas, 1996; Katok et al., 2002). There is some evidence of learning, but even after several rounds of play, conformity with the predictions of the mechanism is relatively low. These results cast doubt on the implementability of more complex mechanisms.

Nash (1950a) showed that every finite strategic form game has an equilibrium in mixed strategies. Section 12.3 describes experimental evidence on mixed strategy Nash equilibrium (MSE) from a range of experiments. Many of the experimental results in this area are conducted on *zero-sum two-player games*. In these games, an MSE coincides with the maximin and the minimax outcomes,¹ a result that is often exploited in experimental games.

In Section 12.3.1, we first consider early evidence from a two-player zero-sum game in which each player has two strategies (O'Neill, 1987). The overall marginal frequency of play of each strategy is remarkably close to the MSE prediction. However, in an important paper, Brown and Rosenthal (1990) arrived at the opposite result, using exactly the same game. They showed that O'Neill's tests lack power to discriminate among alternative theories. In particular, the observed data in O'Neill (1987) can be explained by a range of plausible behavioral explanations. The statistical tests reject the hypothesis that individual data is identical to the MSE predictions. The current choice of the subjects in experiments is heavily influenced by their past actions as well as the past actions of the opponents. By contrast, theory predicts that over successive rounds, these choices should be independently and identically distributed (iid). There is little evidence of learning by the players.

In the *equilibrium in beliefs* interpretation of MSE, players do not actually play mixed strategies (Binmore, 1991). Rather, each player plays a pure strategy but the proportion of players who play each pure strategy is identical to the predicted MSE. Binmore et al. (2001) apply these ideas to eight different games (including the game in O'Neill, 1987) and allow for the possibility of learning. They find some evidence of convergence. However, the observed frequencies of play, although visually close to the MSE predictions, are statistically different. Section 12.3 then shows that *players do actually play mixed strategies, but not the ones predicted by MSE* (Shachat, 2002). This rejects the equilibrium in beliefs interpretation.

Mookerjee and Sopher (1994) reject MSE not only at the individual level but also at the aggregate level (unlike the findings of O'Neill (1987) and several others). Their verdict on MSE (p. 109) is that “it cannot even be treated as an acceptable as if theory.” Like Brown and Rosenthal (1990), but unlike Binmore et al. (2001), they find no tendency of play to converge to MSE.

¹ That the minimax and maximum outcomes are identical in these games was showed by von Neumann (1928).

Subjects randomize more successfully in game theoretic tasks, relative to decision theoretic tasks (Rapoport and Boebel, 1992; Budescu and Rapoport, 1994). It is likely that game theoretic interaction is more cognitively challenging, making it more difficult to remember past strategies. However, in each case, MSE fails.

Section 12.3.2 considers the evidence for MSE in the context of games with an explicitly economic context. We first consider experiments in a three-firm Hotelling competition model (Collins and Sherstyuk, 2000). The predictions of MSE are violated at the individual level and at the aggregate level. Rapoport and Amaldoss (2000) consider an R&D investment game that has a unique MSE. There is qualified support for the MSE at the aggregate level but rejection of MSE at the individual level. We also examine the related literature that considers similar classes of games.

In order to disentangle the separate effects of the players' desire to play an MSE and their inability to successfully make iid draws from the same distribution, in Section 12.3.3 we introduce a *mixed strategy device*. Such a device elicits the players' desired mixed strategy and then aids in making iid draws from the desired distribution (Bloomfield, 1994; Ochs, 1995; Shachat, 2002). When subjects are provided with a mixed strategy device, their play still does not conform to the predictions of MSE (Shachat, 2002). To be sure, players play mixed strategies, however, their mixtures are quite different from the MSE predictions. This evidence contradicts the *equilibrium in beliefs* interpretation of MSE.

Several papers have explored the ability of professional sports players to detect the play of an MSE. For instance, consider a tennis player randomizing serves to both sides of the opponent. Or a penalty kicker in football, randomizing kicks to both sides of the goalkeeper. One could then test for an implication of MSE by calculating if the win rates of players are identical from serving left or right or kicking to the left or right. In Section 12.3, we discuss the caveats to this line of research. For instance, this research tests for motor skills that are well-honed in human evolution. It is not clear if these results are portable to interesting economic contexts, say, investing in a portfolio of assets or choosing among pension plans. Furthermore, there are other aspects of the problem that are not easy to control for. For instance, in a tennis serve, some of the other relevant dimensions include the amount of topspin, sliced serves, or other bodily deceptions while serving.

In Section 12.3.4, we first report evidence that win rates on tennis serves for professional players are statistically similar for left and right serves (Walker and Wooders, 2001). However, there is serial dependence in the serves of even the top professionals that runs counter to the predictions of MSE. Interestingly, this serial dependence is not discovered by the opponents, who are, therefore, not able to exploit it. In a similar vein, there are statistically identical win rates from kicking to the left or to the right of the goalkeeper in professional soccer games (Palacios-Huerta, 2003). Section 12.3.4 also reports results on the portability of successful randomization experience in sports to other domains. Initial results claimed to have found evidence for such portability (Palacios-Huerta and Volij, 2008). However, this claim has been rejected in several carefully conducted studies (Wooders, 2010; Levitt et al., 2010).

In Section 12.4, we consider multiple equilibria that are ubiquitous in economics. For instance, in New Keynesian macroeconomic models, one of the equilibria is Pareto optimal (say, full employment) but the others are sub optimal (perhaps these involve a degree of unemployment). One of the central insights of this literature is that economic agents may not be able to coordinate on the Pareto optimal equilibrium. Such a situation is said to be one of *coordination failures*.

In game theory, multiple equilibria can arise in several ways. Camerer (2003a, p. 341) suggests a useful classification. In one class of games, all equilibria may have the same payoffs for the players. Consider, e.g., players who simultaneously choose locations to meet up in New York in

Schelling's classic example. In this case, players may use *conventions* or *focal points* to coordinate their actions. In a second class of games, there are asymmetric payoffs in the different equilibria, e.g., the battle of sexes (BoS) game. In this case, pre play communication, or the timing of moves may determine actions. In the third class of games, the equilibria are asymmetric but players are symmetric. An example is the stag-hunt game. Various selection principles such as *payoff dominance* or *risk dominance* may determine the outcome in this case.

In Section 12.4.1, we consider coordination failures in simple coordination games (Cooper et al., 1989). We begin, in Section 12.4.3, with *median action games*. In these games, the payoffs of players depend on their chosen action and on the median action of the group. The experimental evidence shows widespread coordination failures (Van Huyck et al., 1991). Furthermore, none of the selection principles such as payoff dominance or risk dominance are able to fully account for the data in any of the treatments. The equilibria are also strongly *history dependent*.

In choosing among multiple equilibria, one long-standing conjecture, dating back to Schelling (1960), is that people will choose a *focal* or a *salient* equilibrium. The relation between focal points and coordination is considered in Section 12.4.2. Mehta et al. (1994) consider a focal or a salient outcome in a wide variety of situations, such as coordinating on the name of a flower, city, year, color, car manufacturer, and a boy's name. When the payoffs of subjects depend on coordinating their choices with others (relative to making their own preferred choices), choices of players achieve relatively greater coordination, suggesting that *labels* carry important salience or focal information. Furthermore, non-Nash strategy profiles can also serve as focal points even when the game has Nash equilibria (Bosch-Domènech and Vriend, 2013).

In Section 12.4.4, we consider *weak link games*. These games are similar to median action games, except that the median is replaced by the minimum action over all players. For instance, in a sequential multi-stage production process, a slowdown at any stage, delays all subsequent stages. The results here are similar to those in median action games, i.e., widespread coordination failures, and an inability of any of the selection principles to fully organize the data (Van Huyck et al., 1990). The organization of groups as captured, for instance, in permanent versus random pairing of other group members, is a more important determinant of coordination than just the size of groups.

In Section 12.4.6, we consider issues of coordination in a stag-hunt game (Battalio et al., 2001). These authors consider three different kinds of games with identical multiple Nash equilibria. Define the *optimization premium* as the Nash equilibrium payoff minus the payoff from the next best action. The optimization premium does not alter any of the Nash equilibria, however, its size has an important bearing on the extent of coordination failures.

In section 12.4.7, we examine the effect of *outside options* and the role of *forward induction* in the class of two-player coordination games (Cooper et al., 1993). Even when classical game theory predicts that the outside option will have no bite, it alters the outcome to the advantage of the player who has the outside option.² There is also evidence for the role of *forward induction* and for a range of factors related to *timing* and *framing* effects that are not predicted to have any effect in classical game theory.

Section 12.4.8 considers the effect of preplay communication in achieving coordination in two-player games (Cooper et al., 1989, 1992, 1993, 1994). Two forms of communication are studied: One-way communication (1W) and two-way communication (2W). In the former, only one player is allowed to engage in *cheap talk* (sending costless messages) prior to the game, while in

² For instance, the outside option has no bite when the payoffs from the outside option are lower than what one might get from the continuation game.

the latter, both players can engage in cheap talk. In the battle of sexes game, 1W communication fosters coordination on the preferred equilibrium of the player who is allowed to engage in cheap talk. The opposite result is found for coordination games in which there is a conflict between payoff dominance and risk dominance, such as in stag-hunt games. In this case, 2W communication works very well. For instance, if both players can assure each other to play the risk dominant equilibrium through 2W communication, then this acts as an assurance to both players and facilitates coordination.³

Section 12.4.9 explores the powerful *role of history* in influencing the degree of cooperation in the *continental divide game* (Van Huyck et al., 1997). One of two possible equilibria in this case is determined entirely by the accidental choices in the early rounds of the experiment.

Section 12.4.10 considers the role of *financial incentives* such as bonuses in enhancing cooperation (Brandts and Cooper, 2006). We consider a *corporate turnaround game* in which a firm chooses a wage level and a bonus level. A team of workers then simultaneously choose their effort levels. The production technology is such that the minimum effort level in the team determines the final outcome. So, if workers cannot coordinate on a high effort level, the output produced and the bonus received are low. The introduction of a bonus enhances cooperation. However, the mere presence of the bonus, rather than its precise rate seems critical in promoting coordination. Furthermore, groups of players who had *strong leaders* (who make a large effort response to the bonus) do better at coordination.

Most of the literature has typically looked at payoff dominance and risk dominance as selection principles in coordination games. In Section 12.4.11, we consider a new selection principle based on *loss aversion* (Cachon and Camerer, 1996). In a median action game, they find that some of the improved coordination that is typically ascribed to forward induction actually arises on account of loss aversion.

In section 12.4.12, we ask the question if coordination can be “grown” in the lab by gradually increasing group size (Weber, 2006). It is surprisingly difficult to cultivate coordination in the lab. Coordination emerges successfully only under special conditions that include gradual growth of group sizes, successful coordination in small groups, and a knowledge of the history of such coordination.

Section 12.4.13 shows that inter-generational *public advice* under common knowledge can enhance coordination (Chaudhuri et al., 2009). However, private advice and public advice without common knowledge are not effective. Similar results are found by Schotter and Sopher (2003, 2007). Interestingly their findings are suggestive of *punctuated equilibria*, which we shall encounter again in Part 5 when we undertake the study of stochastic social dynamics. In a punctuated equilibrium, one of the coordination equilibria gets established for a few rounds (e.g., as a convention or a norm), only to be replaced by another, and yet emerge once again a few rounds later.

Each of the coordination equilibria is a Nash equilibrium, so players should play their part of the strategy profile, conditional on the other players playing their equilibrium strategy. We ask in Section 12.4.14 if the recommendations of an external arbiter to play a coordinated equilibrium will be followed (Van Huyck et al., 1992). It turns out that the recommendations are only followed if they do not conflict with some other salient principle such as equality of payoffs or payoff dominance.

Section 12.4.15 considers the role of *preplay communication* in the corporate turnaround game (Brandts and Cooper, 2007). They append to the corporate turnaround game a prior stage in

³ Mutual assurance could work, for instance, because humans may have a cost of breaking a promise (see Chapter 7).

which the manager and the workers can communicate with each other. They find that preplay communication significantly enhances coordination. Indeed, the coordination achieved through such communication is higher than that arrived at through purely financial incentives. Effective communication takes a particularly simple form. The manager requests higher effort from the workers, stresses the mutual benefits of doing so, and highlights that workers are well paid!

In Section 12.5 we consider bargaining games. We shall mainly focus on theories of bargaining in the mainstream economics literature.⁴ Within these theories we further restrict attention to those that have been sufficiently empirically tested; thus, we omit the theoretically important class of coalitional theories of bargaining.

Section 12.5.1 considers the *Nash bargaining approach*. Here, one begins with a set of axioms that summarize the desirable features of any bargaining outcome and then derives the unique bargaining outcome that satisfies the relevant axioms (Nash, 1950b). In this approach, there is no bargaining protocol, or rules of bargaining. Hence, tests of this approach are best conducted using *unstructured bargaining* in which parties bargain without explicit rules. Anonymity of subjects in experiments is typically preserved, unless one wishes to explore the implications of face to face bargaining.

We first briefly explain the Nash bargaining approach and the implications of relaxing some of its axioms and then consider the relevant evidence. Early experiments found a strong tendency for equal splits to be implemented (Nydegger and Owen, 1975; Rapoport et al., 1977; Siegel and Fouraker, 1960). However, these studies typically confounded the distinction between the utility and the monetary payoffs of players, as well as the amount of information that players have. The first clinching and satisfactory test of the model was conducted by Roth and Malouf (1979), and it gave support to the Nash bargaining solution. We also explore the role of other factors that are not consistent with any of the main existing normative theories of bargaining. These include the effect of risk attitudes (Murnighan et al., 1988), deadlines (Roth et al., 1988), fairness and relative information possessed by players (Roth and Murnighan, 1982), and sociological factors (Roth et al., 1981).

Section 12.5.2 considers structured models of bargaining in the tradition of Ståhl (1972) and Rubinstein (1982). In these models, there is a definite bargaining protocol, and a set of strict rules of bargaining. Two players (or parties) alternatively make offers and counter offers that can be accepted or rejected by the other player. If an offer is accepted, the game ends. If rejected, the game proceeds to the next round with the other player making a counter offer that can be either accepted or rejected, and so on in successive rounds. In these models, delay is costly. Delay costs may arise due to time discounting reasons, or due to other factors that reduce the size of the pie, such as unproductive strikes, or go-slow tactics in the workplace. Thus, players who have greater patience may be willing to hold out for a better settlement.

We begin by considering non-cooperative bargaining experiments under full information in Section 12.5.2. Early results indicated strong support for the alternative offers bargaining game (Binmore et al., 1985). However, several objections were raised about their experimental protocol. Accounting for these issues either failed to replicate the results (Güth and Tietz, 1988; Neelin et al., 1988) or found significantly weaker support for the theory, as in Ochs and Roth (1989), who also find that the subgame perfect outcome is a poor predictor of experimental data. Ochs

⁴ Other approaches are available in the social sciences. For instance, there is a great deal of literature on *negotiations research* that tends to be taught in business and law schools; see, for instance, Tsay and Bazerman (2009). A set of papers on a rich spectrum of alternative theories of bargaining in the social sciences is collected in Bolton and Croson (2012).

and Roth (1989) also make the surprising discovery, replicated elsewhere (Binmore et al., 1985; Neelin et al., 1988), that when players reject an offer, they often make counteroffers that give them an even lower share of the pie (disadvantageous counteroffers).

Several reasons have been given for the puzzling finding of disadvantageous counteroffers. It could be that players use heuristics that vary in the degree of strategic and fairness considerations, and players may switch between them depending on the context (Güth and Tietz, 1990). Another extremely plausible argument is that the players care not only about their own payoffs, but also their relative payoffs. Thus, players may reject a high offer because they get a relatively low share, yet the disadvantageous counteroffer is more attractive because it gives them a relatively higher share (Bolton, 1991). There is evidence that when such models are nested within a *quantal response equilibrium* framework, they predict well, out of sample (De Bruyn and Bolton, 2008).

We also consider a range of other bargaining issues under full information. The evidence does not seem to indicate experimenter demand effects on the bargaining outcome (Bolton and Zwick, 1995; Bolton et al., 1998). While there is still a substantial departure from the subgame perfect outcome in the presence of outside options (Binmore et al. 1989), there is greater support for the *deal me out solution* relative to the *split the difference solution*.⁵ When the infinite horizon version of the alternating offers bargaining game, as in Rubinstein (1982), is implemented by having a probability of termination of the game at each stage, then the predictions of the relevant theory are violated (Zwick et al., 1992).

Section 12.5.2 considers bargaining under one-sided and two-sided asymmetric information. In the alternating offers bargaining game, results are sensitive to whether parties have asymmetric information about each other or not. Delay is costly, yet one of the parties may be willing to hold out longer, to signal that it is more patient, in the hope of eliciting greater future concessions. Asymmetric information can be one-sided or two-sided. There is a relatively small number of studies on this subject, hence, it would be wise to resist general conclusions, yet there are some intriguing findings.

We first consider a dynamic model of one-sided asymmetric information in which a seller and a buyer bargain over a single unit of an object (Rapoport et al., 1995). The objective is to test the relevant theoretical predictions of this model relating to the time path of prices and the comparative static effects of parameters, such as discount rates (Fudenberg et al., 1985; Gul et al., 1986). The authors draw negative conclusions for the theoretical predictions from their experiments, and find that their subjects follow simple rules of thumb. A second model of one-sided asymmetric information that we consider is a model of strikes in a union–firm setting (Forsythe et al., 1991). While the model is more successful in explaining the results within groups of games where strikes should and should not occur, it is less successful in explaining differences among games in a group. The model also finds it difficult to explain the bargaining allocations among the players based on predictions of Myerson's *random dictator axiom*.

We then move on to two-sided asymmetric information games in which a buyer and a seller each have privately known and uniformly distributed valuations for a single indivisible unit of an object. They bargain under a *sealed-bid bilateral bargaining mechanism* in which players simultaneously submit bids. If the buyer's bid is higher, then trade takes place at the average of the two bids, otherwise both players get zero. Chatterjee and Samuelson (1983) derive the optimal

⁵ These terms are defined in greater detail below. The *split the difference solution* arises from the non-cooperative bargaining models. Essentially, each player gets his disagreement payoff plus a share of the surplus. The *deal me out solution* respects the fact that at each stage, a player will reject any offer below his outside option.

bidding strategies, while restricting to the class of linear bids by both players. Among the class of all Bayesian equilibria, the one in which both players use linear bidding strategies turns out to be the most efficient (Myerson and Satterthwaite, 1983).

On the face of it, these games are the most cognitively challenging of all bargaining games considered in this chapter. Further, given ambiguous support for simpler bargaining games, our prior might be that the data would reject the predictions arising from such games. Amazingly, the finding is just the reverse. We begin with a classical test of the sealed-bid mechanism (Radner and Schotter, 1989). The authors find that the bidding strategies of players are stepwise-linear as predicted. Furthermore, the strategies are different on both sides of the optimal cutoff points in the manner predicted by theory. However, buyers did not shave their bids as much as predicted by the theory.

These results have turned out to be surprisingly robust to changes in experimental protocol, such as random matching of players, more rounds of the experiment, and feedback on the history of the game (Daniel et al., 1998; Valley et al. 2002). Furthermore, the degree of efficiency obtained in the experiments is very high. In the presence of prebid communication, trades take place that are not even predicted to take place in the absence of communication, and the efficiency exceeds the Myerson–Satterthwaite upper bound (Valley et al. 2002). Further, face to face communication is even more conducive to bargaining agreements, relative to written communication only, which in turn is better than no communication at all. It seems that the opportunity to communicate, prior to bidding, allows players to coordinate their expectations and, in a large number of cases, they bid an identical price. Although communication is conducive to bargaining agreements, players are often not fully honest in revealing their private information (sellers overstate costs and buyers overstate valuations). However, players are less likely to lie when communication is face to face.

Section 12.5.3 considers unstructured bargaining that does not fall into any of the other categories considered above. *Self-serving bias* is a leading candidate to explain *bargaining impasse* under full information (Babcock and Loewenstein, 1997). A self-serving bias may occur if players are selective about the facts of the case and about events from their memory, in a manner that serves to underpin their own interests. On account of a self-serving bias, players may confuse between their self-interest and a fair outcome. Hence, they may push too hard for an essentially self-interested outcome, convinced that they are only asking for a fair outcome. This may lead to a bargaining impasse. For instance, we do observe protracted and bitter divorce settlements, and labor union strikes. Such a self-serving bias may also nicely complement and explain some of the earlier bargaining results too. In particular, it alerts us to the fact that asymmetric information is not a necessary condition for a bargaining impasse.

The second theory that we consider in Section 12.5.3 is Schelling's (1960) *tacit coordination* based on payoff-irrelevant cues. In particular, we focus on the empirical study by Isoni et al. (2013). They derive strong support for Schelling's insights, based on stylized, abstract, yet quite simple bargaining games. In these games, players make choices based on spatial cues. The payoffs from these games may be symmetric or asymmetric, yet players find it in their interests to use spatial cues as a coordinating or equilibrium selection device.

Section 12.5.4 considers tests designed to infer the cognitive process of subjects when they play experimental bargaining games. Game theory does not make any direct assumptions about the cognitive processes that underlies strategic choice. Traditionalists might argue that a study of these cognitive processes cannot confirm or reject the predictions of game theory. However, on the flip side, the evidence is so strong and persuasive that it would be undesirable to reject it on a-priori grounds.

Consider Rubinstein's alternative offer bargaining game played by two players over three rounds. The game is easily solved by backward induction. Given the solution concept, one expects players to search backwards. We expect them to first explore the last round data, such as payoffs and actions, then look at the second round data and finally, at the first round data. One may use MOUSELAB software to determine if the pattern of searches and lookups of subjects in experiments is backward, beginning with the third round data (Johnson et al., 2002).⁶ It turns out that most subjects search forward rather than backward. Traditionalists might argue that subjects could search forwards, memorize the data, and then solve the game backwards in their heads. Not only would this be very cognitively costly, evidence shows that when subjects are taught the logic of backward induction, they are more likely to end up searching backward; the outcomes in this case are closer to the predicted outcomes in classical game theory.

Section 12.6 considers the evidence on signaling games. The notion of an equilibrium in signaling games is quite complex and cognitively demanding. Given the evidence from simpler games, reviewed above, a priori it would seem that such equilibria are even less likely to be observed.

In Section 12.6.1, we consider the evidence from simple sender–receiver signaling games (Brandts and Holt, 1992). The only equilibria in these experimental games are pooling equilibria. Of the available pooling equilibria, the authors take as their theoretical prediction, the unique pooling equilibrium that survives the Cho–Kreps *intuitive criterion*. They find that one of the types almost always conformed to the equilibrium predictions. The other type conformed poorly in the early rounds but got better towards the end of the experiment. None of the “unintuitive” equilibria were played. The results are mixed but encouraging for traditional game theory. A well-replicated result is that there is little evidence that the data respects any of the refinements beyond the Cho–Kreps intuitive criterion (Banks et al., 1994).

Section 12.6.2 considers evidence from the *limit entry pricing game* of Milgrom and Roberts (1982) (Cooper et al., 1997b). Under asymmetric information about costs or market demand, firms engage in limit pricing if they restrict their quantity choice relative to the case of full information. By doing so, they attempt to signal bad news to a potential entrant (say, low demand) to induce the entrant to stay out. The payoffs in the experiments are chosen in such a manner that there are both separating and pooling equilibria. In the first few rounds, subjects try to play the full information equilibrium. As subjects gain experience in these games, they move towards a pooling equilibrium. However, convergence is slow. Even in the best sessions, a substantial fraction of the subjects still do not play according to the equilibrium predictions. In treatments that do not have a pooling equilibrium, subjects have a tendency to engage in pooling before a separating equilibrium emerges.

Section 12.6.3 considers models of finitely repeated asymmetric information games. The equilibrium predictions in this class of models (the KWMR model) derive from the results of Kreps and Wilson (1982) and Milgrom and Roberts (1982). Jung et al. (1994) put this model to the test in an entry deterrence game played over several rounds. Up to about round 31, the experimental data does not support the equilibrium predictions and existing incumbents are too accommodating of entry relative to the predictions. However, from rounds 31 onwards there is evidence of *predatory pricing*, i.e., firms take actions to deter entry. However, there are important departures from the predictions of the KWMR model.

Section 12.6.4 considers evidence from models of asymmetric information in corporate finance (Cadsby et al., 1990, 1998). Managers might wish to signal their private information on the prospects of their firm to uninformed investors by sending some costly signal (Myers and Majluf,

⁶ The early, pioneering work is by Camerer et al. (1994).

1984). For instance, managers could signal good prospects by relying more on borrowing relative to equity finance. This may explain, for instance, why firms are sometimes excessively leveraged. When costly signaling is not allowed, experimental play converges to the predictions of the signaling model quite rapidly (Cadsby et al., 1990). Such convergence is quicker when there is a unique equilibrium. When there are multiple equilibria, separating, hybrid, and pooling, then outcomes converge to the pooling equilibrium. In the presence of costly signaling, the findings are less positive for the classical predictions (Cadsby et al., 1998).

Contributions from large donors (leaders) to charities are often found to elicit enhanced contributions from others (followers). One possible explanation is that most donors have asymmetric information about the quality of the charity. The contributions of the leaders convey valuable information about quality. An alternative hypothesis is that players are motivated by “other-regarding preferences.” In this view, the initial leadership contributions elicit positive reciprocity among followers who then contribute more. In section 12.6.5, we consider experiments that distinguish between these two competing hypotheses (Potters et al., 2007). The signaling hypothesis is predicted to have bite in the asymmetric information treatment but not in the full information treatment. Reciprocity turns out to be important because even in the full information treatment, leadership contributions elicit more contributions from followers. However, under asymmetric information, the effect of leadership contributions on the contributions of followers is even more potent, suggesting an additional role for the signaling hypothesis.

Section 12.6.6 looks at the case of *multiple signals* in three classes of games, the prisoner’s dilemma game, the stag-hunt game, and the game of chicken (Duffy and Feltovich, 2002, 2006). In alternative treatments, they allow subjects to either engage in cheap talk, prior to the game, or allow them to observe the history of the game, or allow them to observe if previous cheap talk by players was matched by the intended actions. They find that when cheap talk was predicted to be credible, then the level of cooperation among the players was greater in the cheap talk treatment. In this case, *words spoke louder than actions*. When subjects receive additional information about the past play of opponents, there are two possibilities. One is that past intentions stated through cheap talk were matched by past actions and the subject is revealed to be honest. The other is that subjects are found to be dishonest. If currently stated intentions are identical to past intentions, then the past honesty of players becomes an important determinant of the current response of players. Hence, cooperation is either enhanced or reduced, depending on the level of inferred honesty.

Section 12.7 considers the experimental evidence on *correlated equilibria*. The idea behind a correlated equilibrium is that if all players in a game can condition their play on a commonly observed public signal, then they can improve their payoffs. Think, for instance, of the coordinating role of traffic lights (public signal). It is beneficial for everyone to “go” when the light is green and “stop” when the light is red. The *recommendation* to “go” when the light is green and “stop” when it is red is self-enforcing in the sense that if others are expected to follow the recommendations, then it is in the interests of each player to also follow them.

Cason and Sharma (2007) aim to see if players play a correlated equilibrium in which the payoffs of all players are greater, relative to any of the three Nash equilibria of their game. Subjects in experiments do not always follow the recommendations of a choreographer who recommends each player to play his part of the correlated equilibrium. Recommendations are more likely to be followed when humans play against robot opponents, as compared to human opponents, suggesting that subjects may doubt the rationality of others. When non-binding preplay communication is allowed, then the evidence is consistent with the play of recommended strategies (Moreno and Wooders, 1998). Duffy and Feltovich (2010) try to implement several strategy profiles that give

payoffs higher than, equal to, or lower than the Nash equilibrium payoffs. They find that a recommendation to play a correlated equilibrium moves play in the direction of such an equilibrium, relative to the mixed strategy Nash equilibrium of the game. Correlated equilibrium turns out to be a necessary, but not a sufficient condition for recommendations to be followed.

The concepts of *strategic complements* (SC) and *strategic substitutes* (SS) are fundamental in many areas of economics, such as coordination failures in macroeconomics, and network externalities. The actions of two players are SC if an increase in the action by one player increases the marginal payoff of the action of the other player. By contrast, in the case of SS, the marginal payoffs of players are decreasing in the actions of the opponents. Hence, one would expect the actions of players to be positively correlated in the SC case and negatively correlated in the SS case. Section 12.8 examines evidence from strategic environments characterized by SC and SS.

Section 12.8.1 shows that cooperation among players is higher when actions are SC rather than SS (Potters and Suetens, 2009). Play does not converge to the Nash equilibrium of the stage game. Interestingly, in both cases, SC and SS, the actions of the players are positively correlated. Reciprocity is one possible explanation. When a player takes an action that facilitates, or reduces, joint payoff maximization, the other player reciprocates likewise.

The work of Fehr and Tyran (2008) in Section 12.8.2 has deep and profound consequences for modern macroeconomic theory. They consider a monopolistic competition type model in which firms derive utility from their own price, the price level of others, and the nominal money supply.⁷ Two factors turn out to be important in explaining nominal rigidity; these are *money illusion* and *anchoring*. The strategic environment also turns out to be an important determinant of the extent of nominal rigidity. Prices are more rigid in the SC case relative to the SS case.

Competitive general equilibrium (CE) lies at the heart of modern economic theory. A CE is a vector of prices of all outputs and inputs such that all consumers maximize utility, all firms maximize profits, and all markets clear. The existence of a CE requires several strong assumptions, thus, it is important to experimentally test it. Theoretical treatments of competitive equilibrium require the existence of a fictitious actor, the Walrasian auctioneer, which makes it difficult to falsify the theory. Importantly, in a CE, no trading is allowed at any non-equilibrium price, and all trading immediately takes place when an equilibrium price is discovered by the Walrasian auctioneer. Two approaches have been taken in the literature, and we outline both in Section 12.9; these approaches create different proxies for a Walrasian auctioneer.

The first attempt to experimentally test for a CE was made in a seminal paper by Chamberlin (1948) that we consider in Section 12.9.1. Chamberlin (1948) put Harvard economics students in a room, assigned sellers their private costs and buyers their private valuations, and asked people to freely meet and strike bargains. This implementation of a CE would seem to many as eminently plausible. However, the experimental results reject a CE. In interpreting these results, Chamberlin (1948) questioned the very notion of a CE.

The second approach to testing for a CE was proposed by Smith (1962); we consider this in Section 12.9.2. Smith's innovation was to introduce the institution of a *double auction* and induced value elicitation. In this method, participants in their role as buyers and sellers are assigned privately known values and costs, and make public bids and asks. A critical component of the double auction institution is the *market improvement rule*. In successive rounds, sellers can only adjust their "ask price" downwards and buyers can only "bid" their price upwards until some trade takes place. Smith (1962, 1964) found evidence consistent with a CE. These results have been widely

⁷ This class of models should be familiar to most advanced undergraduate macroeconomics students. See, for instance, section 8.1 in Blanchard and Fischer (1989).

replicated; convergence to a CE is quick even when there are very few market participants, sometimes as few as three. However, Smith (1962, 1964) and Chamberlin (1948) allow for out of equilibrium trading and institutions such as the double auction method that are not a part of the theoretical framework. The Walrasian hypothesis, namely that prices increase more when excess demand is high, is rejected by Smith's experiments.

When the double auction experiment is played out with *zero intelligence robots* and the market improvement rule is imposed on their "bids" and "asks," a CE results in a few rounds with high economic efficiency (Gode and Sunder, 1993); we consider this in Section 12.9.3. In the textbook competitive equilibrium model, the likelihood of a CE being achieved is always greater under full information. Just the opposite result is found in experiments. Introducing full information retards rather than facilitates a CE. Perhaps the main message of the CE experiments is to demonstrate that an institution (double auction, augmented with a market improvement rule) exists that can implement a CE. One may wish to employ this institution on account of the desirable efficiency properties of a CE. While we do observe a similar institution in stock market trading, and in commodity markets, and in a few other selected domains, it is perhaps a puzzle why we do not observe it more widely. These experiments still implement a fairly simple version of a CE. When the setting is enriched by allowing for endogenous quality, then a CE does not obtain, as we showed in the introductory chapter and in Part 2 (Fehr et al., 1998).

12.2 Iterated deletion of dominated strategies

In this section, we consider two classes of games. Normal form games that can be solved by iterated deletion of dominated strategies and extensive form games of perfect information, such as centipede games, that can be solved by backward induction. For pedagogical reasons, it is expedient to defer a consideration of some other games in this category, such as *market entry games* and the *p-beauty contest*, to subsequent sections. The main empirical finding is that, while players are able to engage in a few steps of iterated dominance (usually one or two), higher steps of reasoning are rare.

12.2.1 Failure of two steps of iterated dominance

Beard and Beil (1994) gave one of the simplest, and earliest, demonstrations of the failure of iterated dominance using a game that was motivated by Rosenthal (1981). The baseline game is shown in its extensive form in Figure 12.1. The strategically equivalent normal form representation, with player 1 being the row player, is also shown.⁸ Players 1 and 2 move in sequence and choose an action from their respective strategy sets $\{L, R\}$ and $\{l, r\}$. The payoffs corresponding to each strategy profile are shown at the terminal nodes.

In the extensive form game, the unique outcome arising from backward induction is the strategy pair (R, r) . There is another Nash equilibrium, (L, l) , but it relies on a non-credible threat by player 2 to play his dominated action l , if given the move. Two steps of iterated dominance are needed to achieve the backward induction outcome. Strategy r weakly dominates strategy l . Mutual knowledge of rationality allows player 1 to delete strategy l . Rolling back the game tree, strategy R is the dominant strategy for player 1. Thus, the strategy pair (R, r) is the unique

⁸ By contrast, for a given normal form game there is no guarantee that one can write a unique representation in terms of the extensive form. Many extensive form games can correspond to a given normal form game.




Figure 12.1 The extensive form and the strategically equivalent normal form in the baseline treatment.

Source: Beard and Beil (1994). Reproduced with permission. Copyright, INFORMS, <http://www.informs.org>

prediction of classical game theory. A rejection of the equilibrium prediction may indicate the inability of players to perform two steps of iterated reasoning, or a lack of mutual knowledge of rationality. In the baseline treatment, 66% of subjects in their role as player 1, nevertheless, played L. It is possible that, while player 1s respect dominance, they are not confident that player 2s will also respect dominance.

Suppose that player 1 assigns probability p that player 2 will play r . Denote by q , the actual fraction of player 2s who responded to the play of R with r . Then, assuming subjects are randomly matched, player 1 is predicted to be indifferent among L and R if and only if

$$9.75 = 3(1 - p) + 10p \Leftrightarrow p = 0.97.$$

Thus, if $p > 0.97$, then, clearly player 1 would choose R over L (and L over R when $p < 0.97$). In the experiment, the fraction of player 2s who responded to the play of R by player 1s with the dominant action r , is $q = 0.83$. Since $0.83 < 0.97$ ($\Leftrightarrow q < p$) we cannot necessarily infer that 66% of player 1s who played L were irrational. This, does not, however, explain why 34% of player 1s played R.

Classical game theory typically assumes that people derive utility solely from their payoffs in the payoff matrix. To check for the influence of other factors such as resentment, risk, assurance, and reciprocity, Beard and Beil (1994) employ a range of other treatments that are reported in Table 12.1. We omit the baseline treatment (Treatment 1) that we have already described above. For all treatments in Table 12.1 (Treatments 2–7), $p < q = 1$, so, classical game theory predicts that player 1 should choose R over L. Thus, in the last column of Table 12.1, the predicted numbers in classical game theory are 0.0.

In Treatments 2 and 3, the *reduced risk treatments*, the payoff to player 1 from choosing L is reduced (to 9 and 7 respectively). This reduces the risk to player 1 from choosing R, making it more likely to be the chosen strategy. In Treatment 4, the *greater assurance treatment*, the payoff to player 2 from playing l is reduced. This should increase the confidence of player 1 that player 2 will not choose the dominated strategy, l . In Treatment 5, the *greater resentment treatment*, the payoff of player 2 from the strategy profile (L, l) or (L, r) is increased. Thus, if player 1 were to play R, then player 2 has to forgo a greater payoff, irrespective of her own strategy. Thus, the play of R by player 1 might make player 2 more resentful who, in turn, can then decide to play l instead of r in order to punish player 1. This is expected to increase the incentive of player 1 to play L. In Treatment 6, that combines *lower risk* for player 1 and *greater reciprocity* on the part of player 2, the incentive of player 1 to play L increases on both counts (lower risk and greater reciprocity). Finally, in Treatment 7, players face *higher stakes*. However, players receive their associated payoff with probability 1/6.

Table 12.1 Treatments 2–7.

Treatment	Number of pairs	Payoffs from various strategy combinations			% of Player 1s choosing strategy L
		(L, l) or (L, r)	(R, l)	(R, r)	
2	31	9, 3	3, 4.75	10, 5	0.65
3	25	7, 3	3, 4.75	10, 5	0.20
4	32	7, 3	3, 3	10, 5	0.47
5	21	7, 6	3, 3	10, 5	0.86
6	26	7, 5	5, 9.75	10, 10	0.31
7	30	58.5, 18	18, 29.4	60, 30	0.67

Source: Beard and Beil (1994).

The main finding from Table 12.1 is consistent with the intuition in Rosenthal (1981). Except for Treatments 3 and 6, where the choice of L drops to respectively 20% and 31%, a high proportion of player 1s play L. Since, for Treatments 2–7, $q = 1$ (i.e., player 2s always responded with the action r when player 1 played R), we should expect no play of L. But in half the cases, this prediction is not even supported by a majority of the subjects. One plausible conclusion from these results is that players fail to perform one to two steps of iterated dominance. These results were replicated by Goeree and Holt (2001); see one of the problems at the end of the chapter.

Consider the following two-player two-stage sequential bargaining game in Goeree and Holt (2001) that illustrates the failure of two steps of iterated dominance. The problem is to divide a pie that shrinks in size from \$5 to \$2 from the first to the second stage; this is a proxy for discounting the future. The reservation amount for each player in the experiment is zero. In the first stage, player 1 moves first to propose a split of \$5. If accepted by player 2, the proposed split is implemented and the game ends. If rejected, at the second stage, player 2 can make a counteroffer based on a pie of size \$2. Player 1 can either accept the counteroffer (payoffs being divided as in the counteroffer) or reject, in which case, both players get zero.

This is a dynamic game of perfect information, hence, it can be solved by backward induction. Suppose that the smallest unit of money is a cent and the reservation utility of each player is 0. Applying backward induction, at the second stage, player 2 should propose to keep \$1.99 for himself and offer \$0.01 to player 1. Player 1 should accept this offer because it exceeds the reservation amount. At the first stage, if player 1 offers player 2 a share $\$2 > \$1.99$, then player 2 should immediately accept. Hence, the unique backward induction outcome is for player 1 to keep \$3 and offer \$2 to player 2 at stage 1, which should be immediately accepted. Notice that the equilibrium offer of player 1 is the size of the initial pie minus the amount by which the pie shrinks. Thus, if the pie shrinks from \$5 to \$0.50, then, in equilibrium, player 1 keeps \$4.50 and offers \$0.50 to player 2, which should be immediately accepted.

Goeree and Holt (2001) use six cohorts of ten students each as their subjects in the experiment. In the first game, with a pie that shrinks from \$5 to \$2, the average amount kept by player 1 was \$2.83, not far from the equilibrium prediction of \$3. However, in the second game in which the pie shrinks from \$5 to \$0.50, the average amount kept by player 1 was \$3.38, which is substantially lower than the equilibrium prediction of \$4.50. Furthermore, rejections by player 2s were common. By rejecting the offer of player 1, player 2 can get no more than \$0.50, which is not rational.

The results described above are consistent with the pattern of results in similar games that had already emerged by the early 1990s (Davis and Holt, 1993). One possible explanation for the deviations from the backward induction prediction arises from the possibility that players have “other-regarding preferences.” For instance, if relative payoffs are important, then player 2 may reject a higher offer in the first stage, yet be willing to accept a lower amount at stage 2 because it is relatively more equal; see Goeree and Holt (2000) for a fuller consideration of these issues.

12.2.2 *Framing effects and failure of three steps of iterated dominance*

In classical game theory, one may expect the equilibrium outcomes to be identical if we consider games in extensive form or in their equivalent normal form representation (invariance). Schotter et al. (1994) ask if the two alternative game form representations are associated with different framing effects. For instance, spotting dominance may be relatively easier in one representation rather than the other. Figure 12.2 gives the extensive form representation for a two-player game and the strategically equivalent normal form. Subjects in the experiment are asked to play both representations of the game separately.

The backward induction outcome is (R, r) because strategy l is dominated by r for player 2. Folding back the game tree, strategy R dominates L for player 1. Hence, the equilibrium (R, r) requires two steps of iterated dominance. In the normal form version, 43% of player 1s choose R , while in the extensive form version, 92% choose R . One interpretation is that players have less confidence in the ability of player 2s to respect dominance in the normal form version relative to the extensive form version. Such expectations turn out to be justified. In the normal form version, 80% of player 2s choose the dominant strategy, r , while in the extensive form version, this increases to 98%. Thus, framing effects are critical in actual choice insofar as they may help to spot dominance. Framing effects can also arise from other sources. For instance, if only a verbal statement of the game is given, then people’s mental representation of the game may be neither of the normal or extensive form versions. This has implications for which equilibrium, if any, is likely to arise.

A second objective of Schotter et al. (1994) is to test whether individuals undertake three steps of iterated deletion of dominated strategies. Consider the extensive form game in Figure 12.3 with its strategically equivalent normal form. Eliminating weakly dominated strategies, the normal form game can be solved for the unique strict Nash equilibrium (M, M) . To see this: First eliminate the strictly dominated strategy B for player 1, then the (weakly) dominated strategies




Figure 12.2 An extensive form and the strategically equivalent normal form.

Source: Reprinted from *Games and Economic Behavior*, 6: 445–68. Schotter, A., Weigelt, K., and Wilson, C., “A laboratory investigation of multi-person rationality and presentation effects.” Copyright © 1994, with permission from Elsevier.




Figure 12.3 Checking for three steps of iterated deletion of dominated strategies.

Source: Reprinted from *Games and Economic Behavior*, 6: 445–68. Schotter, A., Weigelt, K., and Wilson, C., “A laboratory investigation of multi-person rationality and presentation effects.” Copyright © 1994, with permission from Elsevier.

T and B for player 2, then the dominated strategy T for player 1.⁹ Thus, three steps of iterated removal of weakly dominated strategies are required to achieve the Nash equilibrium (M, M). However, only 16% of the row players and 26% of the column players play M. 82% of the row players play T, the safe strategy. The vast majority (96%) of column players do not play their dominated strategy B. However, reasoning to three levels seems beyond most players. Now consider the extensive form game described in the right panel of Figure 12.3. The main difference from the extensive form game of Figure 12.2 is that, following the original move of L by player 2, both players play a simultaneous move battle of sexes game (BoS) in which each player chooses among one of two possible strategies, M and B. There are two subgame perfect outcomes.¹⁰ In the first equilibrium, player 1 plays T and ends the game, so that each player gets a payoff of 4. In the second equilibrium, player 1 plays L and the payoffs are (6, 3) corresponding to the play of (M, M) in the BoS game.

What outcome can we expect by using forward induction arguments? Suppose that instead of playing T, player 1 plays L and player 2 finds that she is unexpectedly given the move. What should player 2 infer? A rational player 1 must only have given up a payoff of 4 in the expectation of receiving a greater payoff, which only happens if both players play (M, M) in the BoS game. Anticipating this, player 2 should play L. The players may then be able to coordinate on the (M, M) outcome in the BoS game rather than the other pure strategy equilibrium in the BoS game given by (B, B). In the experimental data, while 100% of player 1s choose M in the BoS game, only 31% of column players do so. The vast majority of column players, 69%, play B instead. Possibly, correctly forecasting this, 70% of player 1s play T at the first node and end the game. In particular, the solution found by iteratively deleting weakly dominated strategies, (M, M), is played by only a small minority of players.

The play of T by player 1s is consistent with the prediction of subgame perfection. However, it does not necessarily imply that players do use backward induction. The payoff of 4 for player 1 is relatively high compared to the payoff of 0 that he/she could potentially achieve in three out of

⁹ Recall from the appendix on game theory that iterated deletion of weakly dominated strategies may also eliminate some Nash equilibria in the game such as (T, T) and (T, B). However, (M, M) is the only strict Nash equilibrium in this game.


¹⁰ See the appendix on game theory for the BoS game.

the remaining five cases. Player 1, could, for instance, be sufficiently loss averse and cognitively challenged, in which case it is not surprising that he may choose the safe payoff of 4 from playing T. More varied experimental designs are needed before we can say anything more definite.


12.2.3 Centipede games and failure of higher order steps of iterated dominance

Centipede games, introduced originally by Rosenthal (1981), are dynamic games of perfect information that have a finite number of players and strategies. In the typical game, there are two players who make alternating moves. At each move, a player can either play down, D, and end the game or play across, A, and pass on the move to the other player. Payoffs accrue at each terminal node.

Panels (a) and (b) of Figure 12.4 illustrate, respectively, the four-move and six-move games in McKelvey and Palfrey (1992). These games begin with the choice by player 1 from the actions A and D. If player 1 chooses A at any node, then player 2 also gets to make a choice between A and D. The total payoff to the two players doubles as one moves across to the right along successive nodes. For this reason, the Pareto optimal outcome is achieved if the player at the last node chooses A; this achieves respective total payoffs of \$8 and \$32, in the four-move and the six-move games. However, if the game were to end with the first player playing D at the first node, the total payoff to both players equals just \$0.50. Centipede games can be solved by applying backward induction. In each case, the backward induction outcome is for player 1 to play D at the very first move and end the game. This can be a disturbing outcome, even for classical game theorists.¹¹



Panel (a)



Panel (b)

Figure 12.4 Four-move and six-move centipede games.

Source: McKelvey and Palfrey (1992) with permission of The Econometric Society.

¹¹ See, for instance, Aumann (1992, 1995). Aumann (1995) showed that the backward induction outcome in which the player at the initial node plays D arises not due to rationality, or mutual knowledge of rationality but due to *common knowledge of rationality*. However, the importance of this result depends on how rationality and

Table 12.2 Conditional probability of choosing action D.

Type of game	P1	P2	P3	P4	P5	P6
4 move	0.07	0.38	0.65	0.75		
4 move (high stakes)	0.15	0.44	0.67	0.69		
6 move	0.01	0.06	0.21	0.53	0.73	0.85

Source: McKelvey and Palfrey (1992).

The rationality requirements for the backward induction equilibrium in the case of the centipede game are even more stringent compared to the other games that we have considered above. For instance, player 1 has to solve for four steps of iterated reasoning in the backward induction outcome in the four-move game. In the six-move game, this number rises to six. Insofar as the evidence generated by simpler games suggests that players fail to successfully iterate over a lower number of steps, it should not come as a surprise that the prediction of backward induction often fails for the centipede games.

McKelvey and Palfrey (1992) ran a total of seven sessions with students. No subject participated in more than one session; in each session, each subject participated in ten games. The first three sessions used the four-move game. The fourth session used a high stakes four-move game and sessions five through seven used the six-move game. The results across all sessions are shown in Table 12.2, where p_i is the probability of playing D at node i , $i = 1, 2, \dots, 6$, conditional on play having reached node i .

The results are as follows. (1) There is very little evidence of the backward induction outcome (i.e., play of D at the first node). Across the three types of games (four-move game, four-move game with high stakes, and the six-move game), the respective percentages of play of D at the first node are 7, 15, and 1. (2) The probability of playing D increases as we move closer towards the last node.¹² (3) Comparing corresponding nodes in the high stakes four-move and the six-move games, there is typically a higher probability of playing D in the four-move game, except for node 2.

Rapoport et al. (2003) ran three-player centipede game experiments; the game is shown in Figure 12.5. This experiment is not simply an extension of the McKelvey and Palfrey experiments to an additional player and much higher stakes (the payoffs are the dollar amounts of payoffs). Rather, there are several differences between the two sets of experiments that could have a bearing on the results. (1) Three players instead of two. All three players move in sequence and each player moves twice over the entire game. (2) The payoffs at the last node from moving across are zero for all players, while in the McKelvey and Palfrey experiments these payoffs are the highest possible in the game. (3) Subjects play the game over 60 rounds as compared to the 10 rounds in the McKelvey and Palfrey experiments. This allows for potentially much greater learning.

The results in Rapoport et al. (2003) are quite different from McKelvey and Palfrey (1992). (1) Averaged across sessions 1–4, about 40% of the play by player 1 at the first node is D. (2) The conditional probability of playing D at node i , p_i , is fairly stable across decision nodes 2 to 8.

common knowledge (or common belief) are defined, and there is disagreement over these concepts. These issues are treated in detail in Dekel and Siniscalchi (2015). I am grateful to Pierpaolo Battigalli for drawing my attention to these issues.

¹² There is only one exception to this finding. This was session 5 with Caltech students in which at the last two nodes, the conditional probability drops from 0.91 to 0.50.

	1	2	3	1	2	3	1	2	3	0
	A	A	A	A	A	A	A	A	A	0
D	D	D	D	D	D	D	D	D	D	0
10	2	4	80	16	32	640	128	256		
1	20	4	8	160	32	64	1280	256		
1	2	40	8	16	320	64	128	2560		

Figure 12.5 A three-player centipede game.

Source: Reprinted from *Games and Economic Behavior*, 43: 239–65. Rapoport, A., Stein, W. E., Parco, J. E., and Nicholas, T. E., "Equilibrium play and adaptive learning in a three-person centipede game." Copyright © 2003, with permission from Elsevier.

One conjecture for these results is that humans might be more hardwired to cooperate in dyadic relations (two individuals) rather than triadic (three individuals) relations. Another possibility is that there is some form of learning that explains a relatively greater play of D at the first node. There are, however, two possible caveats to the results. First, 60% of the choices still do not strictly conform to the backward induction prediction. Second, 60 rounds of learning in this simple environment might be too large for many kinds of actual human learning, e.g., buying a mortgage or an insurance, marriage, divorce, and buying a pension plan.

Fey et al. (1996) propose a constant-sum variant of the centipede game. Players divide a constant sum of \$3.20 at each node. The division at the first node is strictly equal. However, as one moves across towards the end node, the division becomes increasingly unequal. Thus, if players have "other-regarding preferences," then they should play D much earlier than comparable experiments. In their ten node and six node versions, 45% and 59% of the outcomes were, respectively, consistent with the backward induction outcome suggesting that other-regarding preferences may be a contributory factor.

Nagel and Tang (1988) considered a two-player, twelve-node, centipede game. Players are shown a strategically equivalent normal form game. Many of the results of McKelvey and Palfrey (1992) are replicated. The modal play indicates that players play D about halfway through the game. Only 0.5% of the choices are consistent with the strict backward induction outcome in which player 1 plays D at the first node. There is insufficient evidence to indicate that players converge, over successive rounds, towards the backward induction outcome.

Palacios-Huerta and Volij (2008) ran artefactual field experiments in which professional chess players play the centipede game. In their first experiment, chess players are matched with other chess players. In the second experiment, conducted in the lab, chess players play the centipede game with students and other chess players over multiple rounds under a random rematching design. In the first experiment, when the centipede game is played with chess players, the outcome of 69% of the games is consistent with the backward induction outcome. When chess grandmasters play against other chess players, without exception, they play the backward induction outcome in centipede games perfectly (i.e., choose D at the first node). In the second experiment, in repeated interactions of chess players against chess players, over ten rounds, 70% of the results are consistent with the backward induction outcome. When students play against chess players (respectively, other students) and acted as player 1, then 30% (respectively, 3%) of the games ended with the play of D at the first node. On the face of it, these experiments suggest that the strong refutation of backward induction that arises in experiments, typically conducted on student populations, may not apply to professionals. Indeed, when professionals who possess




Figure 12.6 A six-move centipede game.

Source: Levitt et al. (2011) with permission from the American Economic Association.

skills in backward induction play the centipede game, their performance provides evidence for classical game theory.

In a careful study, Levitt et al. (2011) carry out a replication of the artefactual field experimental design of Palacios-Huerta and Volij (2008) and reach diametrically opposite conclusions. In their experiments, they ask professional chess players to play two games.¹³ (i) A six-move centipede game in which the backward induction outcome is for the player starting the game to play D; see Figure 12.6. (ii) A “race to 100” game in which players sequentially choose numbers between 1–9 in one treatment and 1–10 in another. The numbers chosen in each sequential stage are added and the individual whose choice leads the aggregate sum to be 100 wins a prize, while the opponent gets nothing. This game can be solved by backward induction.¹⁴ The reason for introducing the second game is that it is not at all obvious that common belief in rationality implies backward induction in the centipede game, while such a link is clearer in the race game.¹⁵

The distinguishing feature of these experiments is the very high quality of the subject pool. Among chess players, ranking is based on the following main three-tiered hierarchy: Grandmasters, International Masters, World chess federation (FIDE) and United States chess Federation (USCF) Masters. Two further categories are introduced for chess players who do not have titles, yet they play at a high level. These two categories are based on Elo rating points for players who have, respectively, greater than and less than 2000 points. Each pair of players in the sample of 206 chess players played the following three games in random order: Centipede game, race to 100 in the 1–9 treatment and race to 100 in the 1–10 treatment.¹⁶

The results are as follows. Only 3.9% of the players respect the strict backward induction prescription to play D at the first node. In 45% of the cases, the play reaches the penultimate node (5th node) of the game. Conditional on the play having reached node 5, 37% of the players play A despite this being a dominated action. Two candidates for an explanation of such behavior are cognitive limitations and social preferences.

Table 12.3 shows the proportion of players who chose D at each of the six nodes, conditional on play having reached these nodes. For instance, conditional on play having reached node 4, across all chess players, 35.2% played D. Recall that Palacios-Huerta and Volij (2008)

¹³ Professional chess players were recruited from two international chess tournaments in 2008. Particular attention was focused on recruiting high ranked players.

¹⁴ This is easiest to see in the 1–9 treatment. If the second mover uses backward induction correctly, he will always win the game. For any number suggested by the first mover, the second mover needs to suggest a number that ends in a multiple of 10. So the following sequence will guarantee a win: 10, 20, ..., 90, 100.

¹⁵ For other experimental results on the race game, see Dufwenberg et al. (2010) and Gneezy et al. (2010). These papers also throw light on bounded rationality and learning issues in games.

¹⁶ Earnings were paid in cash, immediately, and privately, following all three games.

Table 12.3 Conditional probabilities of action D at each of the six nodes. Figures in brackets are the number of observations. The number 2000 refers to Elo ranking points. Source: Levitt et al. (2011).

	Node 1	Node 2	Node 3	Node 4	Node 5	Node 6
All chess players	0.039 (102)	0.102 (98)	0.193 (88)	0.352 (71)	0.587 (46)	0.632 (19)
Grandmasters	0 (16)	0.111 (9)	0.067 (15)	0 (7)	0.636 (11)	1 (2)
International Masters	0 (12)	0 (8)	0.083 (12)	0.625 (8)	0.625 (8)	0 (1)
Masters	0.063 (32)	0.077 (26)	0.259 (27)	0.318 (22)	0.455 (11)	0.5 (8)
> 2000	0.056 (18)	0.154 (26)	0.154 (13)	0.2 (15)	0.5 (6)	1 (2)
< 2000	0.042 (24)	0.103 (29)	0.286 (21)	0.526 (19)	0.700 (10)	0.667 (6)

claim that when chess Grandmasters play other chess players, they play D at the first node in 100% of the cases. However, in Table 12.3 not one of the 16 Grandmasters played D at the first node.

Only 11.11% of the Grandmasters played D at the second node and only 6.7% of the Grandmasters played D at the third node when play reached the third node. A very similar pattern holds for the International Masters who are in the next tier of the rankings. Overall, the results for this dataset decisively debunk the view that highly ranked chess players play the backward induction outcome in centipede games.

In the race to 100 in the 1–9 treatment, 60% of the players play the backward induction outcome. However, players do much worse in the 1–10 treatment; only 12.6% are able to play in accordance with strict backward induction. Players who play best the backward induction outcome in the relatively more difficult race to 100 in the 1–10 treatment are particularly poor at playing the backward induction outcome in the centipede game. Their performance in the centipede game was nearly identical to those who demonstrated little evidence of playing backward induction in the race to 100 in the 1–10 treatment.

Several explanations for the experimental results in the centipede game have been proposed. McKelvey and Palfrey (1992) consider the effect of a slight degree of incomplete information, namely a small probability that any player is altruistic rather than concerned only with own material reward, as in Kreps and Wilson (1982) and Milgrom and Roberts (1982). Altruistic players always play A in the centipede game. If rational players believe that some of their opponents are altruistic, this creates an incentive for them to play A for a while. However, the presence of altruistic players among the subjects in experiments was found to be relatively small (about 7% of the sample). Zauner (1999) provides yet another explanation in which players play with some error. In particular, the individual payoffs are additively and independently perturbed. Assuming a shrinking error rate explains some of the experimental results. Fey et al. (1996) use the *quantal response equilibrium* (QRE) model to explain the increased probabilities of playing D, as one approaches the end node. The *cognitive hierarchy* (CH) model can explain the difference

between the two-player and three-player games.¹⁷ A range of learning models has also been used to explain the results of the centipede game.¹⁸ These models are discussed later.

12.2.4 Evidence from mechanism design problems that are dominance solvable

In this section, we describe two simple mechanisms that rely on iterated deletion of dominated strategies in simultaneous and sequential play. These mechanisms are due to Abreu and Matsushima (1992) and Glazer and Perry (1996). We then examine empirical tests of these mechanisms due to Sefton and Yavas (1996) and Katok et al. (2002). Consider the following two-player game from Glazer and Rosenthal (1992) in which each player has two possible strategies, A and B.

	A	B
A	100, 100	0, 0
B	0, 0	200, 200

The game has two Nash equilibria (A, A) and (B, B). The equilibrium (B, B) is payoff dominant and risk-dominant, so it would seem to be salient.¹⁹ Suppose, that a social planner wishes to implement the Pareto-inferior outcome (A, A) instead. One possible story behind the choice of the social planner is the following. The row player is a shopping mall and the column player is a representative consumer. Strategies A and B are simply locations, say, the suburbs and the city center, respectively. Consumers and shopping malls would like to locate in the city center because of its auxiliary benefits.

A social planner might, however, be worried about the congestion created from such a move. From a social point of view, therefore, it may be preferable for the shopping mall to be located in the suburbs (choice A). The choice (A, A) gives lower private payoffs, of a 100 each, but it gives higher social payoffs that are internalized by the social planner but not by the players (hence, not shown in the payoff matrix). Abreu and Matsushima (1992) and Glazer and Perry (1996) propose two alternative mechanisms to implement the social planner's choice.

The idea in both mechanisms is to play the strategic-form stage game shown above over T time periods, and introduce a small fine f , if action B is played. Let us denote the row and column players, respectively, by players 1 and 2. Let x_{1t} and x_{2t} be, respectively, the strategy sets of the row and the column players at time $1 \leq t \leq T$. We focus on strictly dominated strategies, which, by definition, are pure strategies, so each player makes a binary choice (A or B). Denote the choices of A and B, respectively, by 0 and 1 so that $x_{1t} = \{0, 1\}$ and $x_{2t} = \{0, 1\}$.

In the stage game, shown above, each player gets a strictly positive payoff only if the strategies of both players match.²⁰ If $x_{1t} = x_{2t} = 0$, then the payoff at time $t \in \{1, 2, \dots, T\}$ to each player is $\frac{100}{T}$, and if $x_{1t} = x_{2t} = 1$, then each gets a payoff of $\frac{200}{T}$. If $x_{1t} \neq x_{2t}$, then each player gets a payoff of zero. Thus, the aggregate payoff to any player, $j = 1, 2$, over the T time periods is given by

¹⁷ See also Kawagoe and Takizawa (2012) who formalize this idea more fully.

¹⁸ These include reinforcement models of learning (Nagel and Tang, 1988), adaptive models of learning (Rapoport et al., 2003), and evolutionary models of learning (Ponti, 2000).

¹⁹ This is confirmed in pilot experiments of Sefton and Yavas (1996), where 92% choose B in a game with similar payoffs. See the game theory appendix for the definition of payoff dominance and risk dominance.

²⁰ The justification is that perhaps a coordinated effort is required for the project to be successful, as in the example of the shopping mall and the consumers, above

$$\pi_j = \sum_{t=1}^T \left[\left(x_{1t}x_{2t} \frac{200}{T} + (1 - x_{1t})(1 - x_{2t}) \frac{100}{T} \right) - x_{jt}f\phi_{jt} \right]. \quad (12.1)$$

The first term inside the square brackets is the payoff at any time t , as described above. The second term is the fine paid at time t . The variable ϕ_{jt} in (12.1) takes two possible values, 0 or 1. When $\phi_{jt} = 1$, the fine term is switched on, so that a fine is incurred from a particular choice of action (to be specified) and when $\phi_{jt} = 0$, it is switched off, so that no fine is incurred by any choice of action. We explain the term ϕ_{jt} in more detail below, following a fuller description of the two mechanisms. The main difference between the two mechanisms is that in the Abreu–Matsushima mechanism, the entire profile of actions by the two players is chosen simultaneously at the beginning, while in the Glazer–Perry mechanism, actions are chosen in sequence.

THE ABREU–MATSUSHIMA MECHANISM

Each player simultaneously submits the sequence of choices $\{x_{1t}\}_{t=1}^T$ and $\{x_{2t}\}_{t=1}^T$ where $1 \leq t \leq T$. The first player to play B according to the submitted choice pays a one-off fine, f . If both players make a choice of B in the same period, they both pay the fine. For the Abreu–Matsushima mechanism, ϕ_{jt} in (12.1) can be defined recursively. It is switched on till time t when someone first plays B (i.e., $x_{jt} = 1$) and pays the fine, but for all subsequent periods, i.e., $t + 1, \dots, T$, it is switched off.

Example 12.1 Suppose that $T = 10$ and $f = 30$. The strategies for the two players are given as follows. For player 1, $x_{1t} = B$ for all $1 \leq t \leq 10$. For player 2, $x_{21} = A$ and $x_{2t} = B$ for all $2 \leq t \leq 10$. The choices of both players are different for $t = 1$ but match for time periods $2 \leq t \leq 10$. Thus, the payoff in the first term in the square brackets in (12.1) for each player is $9 \times \frac{200}{10} = 180$. Since player 1 is the first player to play B, he also pays a fine $f = 30$, thus, the respective payoffs of player 1 and 2 are 150 and 180.

We now state the main result under the Abreu–Matsushima mechanism.

Proposition 12.1 Suppose that $f > \frac{200}{T}$ and players have common knowledge of rationality. The unique strategies of both players that survive iterated deletion of strictly dominated strategies are $x_{1t}^* = x_{2t}^* = 0$ for all $1 \leq t \leq T$.

Proof: Fix any strategy profile of player 2, $\{x_{2t}\}_{t=1}^T$. We first show that for player 1, the strategy profile $y_{11} = (0, x_{12}, \dots, x_{1T})$ strictly dominates $y'_{11} = (1, x_{12}, \dots, x_{1T})$. Since the actions under y_{11} and y'_{11} are identical for $2 \leq t \leq T$, the continuation payoffs from the two strategies, following the first period choices, are also identical. Thus, we simply need to show that the current payoff at $t = 1$ from playing $x_{11} = 0$ (choice A) strictly dominates that from playing $x_{11} = 1$ (choice B).

There are two possibilities, depending on the choice of the first period action of player 2. (a) $x_{21} = 1$. In this case, the current payoff to player 1 from playing $x_{11} = 0$ is 0, while the payoff from playing $x_{11} = 1$ is $\frac{200}{T} - f < 0$. (b) $x_{21} = 0$. In this case, the current payoff to player 1 from playing $x_{11} = 0$ is $\frac{100}{T} > 0$, while the payoff from playing $x_{11} = 1$ is $-f < 0$. Thus, y_{11} strictly dominates y'_{11} and the play of $x_{11} = 0$ is a best response of player 1 in period 1 to any strategy of the opponent.

A similar calculation also holds for player 2. Due to common knowledge of rationality, player 1 believes, correctly, that $x_{21} = 0$. Thus, the strategies $x_{11} = 1$ and $x_{21} = 1$ are deleted at the first step of the iterated deletion process and the play of $x_{11}^* = 0$ and $x_{21}^* = 0$ is a best response of each player in period 1 to any strategy of the opponent.

In the next stage, we consider the decisions from time $t = 2$ onwards. Hold fixed the strategy of player 2 at $\{x_{2t}\}_{t=2}^T$ and consider two strategy profiles for player 1: $y_{12} = (0, 0, x_{13}, \dots, x_{1T})$ and $y'_{12} = (0, 1, x_{13}, \dots, x_{1T})$. Proceeding as above, a simple calculation shows that y_{12} strictly dominates y'_{12} . Using common knowledge of rationality we can conclude that the strategies $x_{12} = 1$ and $x_{22} = 1$ are deleted at the second step of the iterated deletion process. Thus, the play of $x_{12}^* = 0$ and $x_{22}^* = 0$ is a best response of each player in period 2 to any strategy of the opponent.

Continue with this recursive process of iterated deletion of strictly dominated strategies to conclude that the unique outcome is $x_{1t}^* = x_{2t}^* = 0$ for all $1 \leq t \leq T$. ■

THE GLAZER-PERRY MECHANISM

Unlike the Abreu–Matsushima mechanism, in the Glazer–Perry mechanism, players move sequentially. In each period, player 1 makes his choice x_{1t} , which is observed by player 2 who then makes his choice x_{2t} . Thus, the sequence of choices over the T time periods is given by: $x_{11} \rightarrow x_{21}; x_{12} \rightarrow x_{22}; \dots; x_{1T} \rightarrow x_{2T}$. The player who is the “last” to choose B incurs a fine, f ; the sequential nature of moves ensures that only one player incurs the fine.

Example 12.2 Suppose that $T = 10$ and $f = 30$. The strategies for the two players are given as follows. For player 1, $x_{1t} = B$ for all $1 \leq t \leq 10$. For player 2, for all $1 \leq t \leq 9$ and $x_{210} = A$. The choices of both players match for time periods $1 \leq t \leq 9$ and are different for $t = 10$. Thus, the payoff in the first term in the square brackets in (12.1) for each player is $9 \times \frac{200}{10} = 180$. Since player 1 is the last player to play B he also pays a fine $f = 30$, thus, the respective payoffs of player 1 and 2 are 150 and 180.

The main result is contained in the next proposition.

Proposition 12.2 Suppose that $f > \frac{200}{T}$. The unique backward induction outcome in the Glazer–Perry mechanism is $x_{1t}^* = x_{2t}^* = 0$ for all $1 \leq t \leq T$.

Proof: Fix any strategy profile of player 1, $\{x_{1t}\}_{t=1}^T$. We first show that for player 2, the strategy profile $y_{2T} = (x_{21}, x_{22}, \dots, x_{2T-1}, 0)$ strictly dominates $y'_{2T} = (x_{21}, x_{22}, \dots, x_{2T-1}, 1)$. In the two strategies, y_{2T}, y'_{2T} , the choices are identical for the first $T - 1$ time periods. Thus, the choice between the two strategies depends only on the payoff arising in period T . Depending on the choice of player 1 in period T , x_{1T} , there are two possibilities. (a) Suppose that player 1 plays B, i.e., $x_{1T} = 1$. The payoff to player 2 from a choice of A, i.e., $x_{2T} = 0$, is 0, while the payoff from a choice of B, i.e., $x_{2T} = 1$, is $\frac{200}{10} - f < 0$ (because player 2 is the last player to choose B). (b) Suppose that $x_{1T} = 0$. The payoff to player 2 from a choice of $x_{2T} = 0$ is $\frac{200}{10}$, while the payoff from $x_{2T} = 1$ is $-f < 0$. In conjunction, (a) and (b) imply that y_{2T} strictly dominates y'_{2T} , so the best response of player 2 in period T to any choice of player 1 is $x_{2T}^* = 0$. Analogously, $x_{1T}^* = 0$.

Let us now consider time period $T - 1$. Fix the strategy of player 1 to be $(x_{11}, x_{12}, \dots, x_{1T-1}, 0)$. Consider two possible strategy profiles for player 2, $y_{2T-1} = (x_{21}, x_{22}, \dots, 0, 0)$ and $y'_{2T-1} = (x_{21}, x_{22}, \dots, 1, 0)$. As above, it is simple to show that the profile y_{2T-1} strictly dominates y'_{2T-1} . A similar calculation can be made for player 1 to show that the optimal choices at time $T - 1$ are $x_{1T-1}^* = x_{2T-1}^* = 0$.

Continue to carry out the backward induction process up to time $t = 1$ and conclude that the unique backward induction outcome is $x_{1t}^* = x_{2t}^* = 0$ for all $1 \leq t \leq T$. ■

Katok et al. (2002) give experimental results for the case $T = 10, f = 50$, for the payoffs given above. They asked one set of students to play the Abreu–Matsushima mechanism, and another set to play the Glazer–Perry mechanism, each over 17 rounds. It was ensured that subjects are not

rematched over the different rounds. The first 3 rounds were practice rounds where no monetary payments were made. By round 4, when incentives are introduced, there was no conformity of play with the prediction of the Abreu–Matsushima mechanism and only 7% of the play conformed with the prediction of the Glazer–Perry mechanism. There appears to be some evidence of learning over time and by round 17, the respective conformity with the predictions of the two mechanisms was 18% and 30%, which, nevertheless is quite low.

Sefton and Yavas (1996) tested consistency of play with the Abreu–Matsushima mechanism. In their experiments, the stage-game payoffs to the two players from (A,A) and (B,B) are, respectively, 480 each and 240 each. They tried various values of T such as $T = 4$, $T = 8$, and $T = 12$ and chose $f = 90$. For the three treatments $T = 4$, $T = 8$, and $T = 12$, $\frac{480}{T}$ equals, respectively, 120, 60, 40. With these payoffs, it can be shown that if $f > 480/T$, then the unique rationalizable outcome, determined by iteratively deleting dominated strategies, is for both players to play B ($x_{1t}^* = x_{2t}^* = 1$). The condition $f > 480/T$ is violated for the case $T = 4$ but it is satisfied for the other two cases, $T = 8$, $T = 12$. Thus, if individuals are fully rational, we should observe all-B choices in the cases $T = 8$ and $T = 12$ and all-A choices in the case $T = 4$ (paying the fine and playing B is profitable in this case).

The experimental findings are just the reverse of this prediction over the 14 periods of the experiment; see Figure 12.7. The percentage of all-B sequences increases gradually over the 14 rounds in the case $T = 4$ (the curve labeled 4T). On the other hand, for the cases $T = 8$, and $T = 12$ (curves labeled 8T and 12T, respectively), the percentage of all-A sequences declines steadily from a low initial level to a near-negligible level. What accounts for these results? For the two cases, $T = 8$ and $T = 12$, a distinct possibility is that subjects are trying to get the higher payoffs from coordinating on the all-B strategy profile as quickly as possible. They are, however, also keen to avoid the fine and so try not to be the first to play B. Irrespective of T , subjects switch to B in the third or fourth rounds and then continue to play B, thereafter.

These results appear to create a degree of skepticism about the possibility of implementation in even more complex mechanisms. Clearly, more empirical research is needed in this area.




Figure 12.7 Percentage of all-B play in various rounds.

Source: Reprinted from *Games and Economic Behavior*, 16: 280–302, Setton, M. and Yavas, A., “Abreu–Matsushima mechanisms: experimental evidence.” Copyright © 1996, with permission from Elsevier.

12.3 Mixed strategy Nash equilibria

Nash's (1950) fundamental result in game theory showed that every finite strategic form game has a *mixed strategy equilibrium* (MSE). The assumption that the players play mixed strategies is central in economic models, hence, it is important to test if people do so.

In an MSE, players must randomize in such a manner that their opponents are exactly indifferent between their pure strategies. It has been troublesome for classical game theory to justify why players should do so. Binmore et al. (2001) work on another interpretation of the mixed strategy equilibrium. In particular, they stress the role of *equilibrium in beliefs*. The *equilibrium in beliefs interpretation of mixed strategies* argues that individuals belong to different populations of players, and within each population, players only play identical pure strategies. Furthermore, the fraction of players in any population that plays a particular pure strategy, is identical to the probability with which that pure strategy is played in an MSE.²¹ In contrast to this interpretation, the evidence reviewed below shows that players do play mixed strategies, but typically not in the proportions that are predicted by an MSE. But this is not the only violation of an MSE; serial dependence in strategies over time is another violation. Yet, at the aggregate level, there is better conformity with an MSE as compared to individual-level data.

Harsanyi (1973) considered the limiting behavior in a set of games that have pure strategy equilibria, but the payoffs are perturbed. A mixed strategy equilibrium is the limiting strategy profile in this class of simple "purified" games. Yet, these results do not hold for a range of games that economists are interested in, such as repeated games, and principal-agent games. Hence this justification for the use of a mixed strategy equilibrium is quite limited in scope. For a detailed discussion of these issues, the interested reader can consult Gintis (2009, Chapter 6).

12.3.1 Do people play mixed strategy equilibria in games within a general context?

Much of the work on testing MSE has relied on finite (typically two-player) zero-sum games in which players have finite strategies (typically two to five strategies). Furthermore, the games considered have a unique MSE in order to abstract from issues of coordination. In these games, the *minimax* and *maximin* outcomes are identical and also coincide with the MSE.²²

The modern work on testing for MSE began with the paper by O'Neill (1987).²³ O'Neill's (1987) work illustrates many of the results found in subsequent research and, hence, there is merit in exploring this work in some detail.

O'Neill (1987) considered the zero-sum game shown in Figure 12.8. Each player has four possible strategies: Ace, 2, 3, and Joker. Players simultaneously choose their strategies. The payoffs are either +5 or -5 cents. The payoffs corresponding to the strategy profiles for which the row player is the winner are shown in a darker background. In the unique MSE, each player plays the strategies, Ace, 2, and 3 with probabilities 0.2 each and the strategy Joker with a probability 0.4. The

²¹ This is reminiscent of the polymorphic equilibria in evolutionary games; see Chapter 14 for the details.

²² Important work in this regard was done by Von Neumann (1928). See also the Appendix on game theory.

²³ For a review of some earlier work on testing for MSE in the 1960s, see Camerer (2003a, pp. 121–5). In these studies, subjects typically played against computerized opponents and had little idea of what the computerized opponent was trying to achieve. There are also concerns about the subjects having been misled (perhaps inadvertently) about the play of the computerized opponent. Subjects also had inadequate incentives to play. Nevertheless, the aggregate frequency of play in some of these studies was close to the mixed strategy equilibrium for the game.

	Ace	2	3	Joker
Ace	-5, 5	5, -5	5, -5	-5, 5
2	5, -5	-5, 5	5, -5	-5, 5
3	5, -5	5, -5	-5, 5	-5, 5
Joker	-5, 5	-5, 5	-5, 5	5, -5

Figure 12.8 O'Neill's zero-sum game to test for an MSE.

Source: O'Neill, Barry (1987). "Nonmetric test of the minimax theory of two-person zero-sum games." *Proceedings of the National Academy of Sciences of the United States of America* 84(7): 2106–9.

Table 12.4 Experimental results.

	Ace	2	3	Joker	Marginal row frequencies
Ace	0.044 (0.040)	0.043 (0.040)	0.043 (0.040)	0.091 (0.080)	0.221 (0.200)
2	0.046 (0.040)	0.038 (0.040)	0.092 (0.040)	0.092 (0.080)	0.215 (0.200)
3	0.049 (0.040)	0.032 (0.040)	0.037 (0.040)	0.085 (0.080)	0.203 (0.200)
Joker	0.086 (0.080)	0.065 (0.080)	0.051 (0.080)	0.158 (0.160)	0.362 (0.400)
Marginal column frequencies	0.226 (0.200)	0.179 (0.200)	0.169 (0.200)	0.426 (0.400)	–

Source: Brown and Rosenthal, 1990.

game is played repeatedly by 25 pairs of subjects for 105 periods for a total of 2625 joint moves. The results on aggregate play are shown in Table 12.4. Each cell in Table 12.4 shows the actual frequency of choices and below it, in brackets, the predicted frequency of the joint play of the strategies by both players in the MSE. The main features of the results are as follows.

1. The MSE implies that each player, in each round, draws a strategy from a multinomial probability distribution with parameters 0.2, 0.2, 0.2, and 0.4. The observed frequencies of *aggregate* play appear very close to the predicted MSE for all strategy profiles. Furthermore, the marginal frequencies of the play of each strategy for the row and the column players (last column and last row in Table 12.4) also seem to bear close proximity to the predicted frequencies.
2. Adding together the predicted MSE frequencies in cases where the row player wins (the shaded cells in Figure 12.8), we get a figure of 0.40. Hence, over a large number of rounds of play of the game, if players follow MSE, the row player should win 40% of the times. The aggregate win frequency of the row player was found to be 0.410.

These results, on the face of it, constitute a stunning validation of MSE predictions, and were interpreted as such by O'Neill (1987). Furthermore, these results, as we shall see, have been replicated for many similar zero-sum games.

In light of this, it is worth considering in some detail, the reappraisal of O'Neill (1987) by Brown and Rosenthal (1990) who interpreted O'Neill's (1987) findings as an unequivocal rejection of the MSE prediction.

1. *Overall winning percentages:* Suppose that the Joker card has greater *salience* relative to the numbered cards, which have, say, equal salience. Suppose that the row player picks out the Joker card with probability x and the three remaining cards with equal probability $\frac{1-x}{3}$. Likewise suppose that the column player picks the Joker card with probability y and the three remaining cards with equal probability $\frac{1-y}{3}$. Furthermore, to capture the salience of the Joker card, suppose that $x, y \in (0.2, 0.8)$. Then the probability that a row player wins in O'Neill's game is 38.3%, which is very close to the MSE prediction. Brown and Rosenthal (1990, p. 1069) write: "In short, the surface that relates winning percentages to independent strategy combinations is quite flat around the minimax strategy. As a result, tests of the minimax model that are based on observed winning percentages are unlikely to reject the minimax model, even when any member of a wide class of alternative models may in fact have generated the data. In statistical terminology, tests based on observed winning percentages for this game lack power."
2. *Frequencies of play of alternative strategy profiles:* One can just as easily generate predicted frequencies from non-minimax play that are actually quite close to the actual frequencies. For instance, suppose that players first try to convince opponents that their choices are purely random. Once the opponents are convinced of this, the players then play a best response to these beliefs of the opponents. Even if players played randomly 80% of the times and a best response only 20% of the times, it can be shown that the resulting frequencies are quite close to minimax. Like the first objection, for the game in Figure 12.8, tests based on matching observed frequencies of play of different strategy profiles have insufficient power to distinguish minimax from non-minimax play.
3. *Statistical closeness of choices to the predictions:* To the naked eye, the observed frequencies seem to have a good match with the predicted minimax frequencies in Table 12.4, but are they close enough in a statistical sense? For all the 16 cells corresponding to the 4×4 payoff matrix, a chi-square test rejects the minimax model at any level of significance. Thus, the observed frequencies could not have been statistically generated by minimax play.
4. *Individual choices:* It is important to note that the predictions of classical game theory are about the choice of strategies for individual players (not aggregate data). The prediction is that individual players will play their part of the MSE. It is easy to reject the hypothesis that individual play conforms to the minimax hypothesis. The chi-square test for the conformity of the choice of strategies of individuals with the minimax outcome gives a test statistic of 348.25 at 150 degrees of freedom. If the strategy choice of individuals was generated by minimax, then the probability of generating this value of the chi-square statistic is zero to at least 15 decimal places.
5. *Homogeneous choices:* One can reject the hypothesis that the data for all pairs of play came from the same multinomial distribution. There is far too much variation in choices of individual cards across all pairs of players.
6. *Effect of past choices on current choices:* When the game in Figure 12.8 is played over repeated rounds, players should draw on a multinomial distribution with parameters 0.2, 0.2, 0.2, and 0.4. Thus, across rounds, the iid (independently and identically distributed) assumption should hold. To test this, given the limited sample sizes, they look at the choice of the Joker card across time. If classical game theory is correct, then in any round, for any player, the choice of the Joker card should be independent of one's own and the opponent's previous choices of the Joker card. Brown and Rosenthal (1990) run logit regressions in which the dependent variable took a value of 1 if a choice of the Joker card was made and

0 otherwise. Among the independent variables are lagged values of the choice of the Joker card for each player.

The explanatory variables turn out to be jointly significant in explaining Joker choices by players. For 19 out of 25 pairs of players, the Joker choice can be statistically determined, at the 5% level, by the independent variables, for at least one player in the pair. The main reason was that players' past choices of the Joker card significantly determined their current choice of the Joker card. Thus, one can reject the hypothesis that the players' choices are uncorrelated with their own past choices.

The autocorrelation in the choices of players could arise either because (1) they are not using minimax, or (2) they are using minimax but are not able to generate iid draws from a underlying multinomial distribution. Brown and Rosenthal (1990) are not able to distinguish between the two explanations. Subsequent research, some of which we review below, has tried to address this question.

The past choices of the Joker card by the opponent also influence the current Joker choice of a player. If players believe that their opponents will play minimax, then this cannot happen. It is plausible that players perceive that opponents will play autocorrelated choices, so they attempt to benefit from them.

7. *Learning*: To examine issues of learning, one could examine the behavior of subjects in the experiment in the final 50 rounds after they had already played 70 rounds. There was no noticeable tendency for the subjects to play minimax in the final rounds.

Slonim et al. (2003) ask if the rejection of minimax is based on the specific features of O'Neill's (1987) game. They consider a range of two-player zero-sum games and sample randomly from these games. Like O'Neill, they find that strategies of players across successive rounds are interdependent. Unlike O'Neill, however, they find that the source of the problem is not too much randomization of choice, but too little. Slonim et al. (2003) find that evidence in support of MSE is more likely to be found in experiments when each of the pure strategies is played with an identical probability (equiprobable play).

Binmore et al. (2001) consider eight games (including the game in Figure 12.8) and look at the evolution of play in such games when they are played a sufficiently large number of times.²⁴ They find that there is temporal dependence of actions on past actions. However, they argue that subjects can only learn to play an MSE through a process of learning, and temporal dependence is evidence of learning. They interpret their results as convergence of observed "aggregate frequencies" to the MSE given sufficient learning opportunities (50–75 practice rounds followed by 150 rounds of the real game).

Mookherjee and Sopher (1994, 1997) consider four constant-sum games in which players could either win or lose, denoted respectively by W and $-W$. A win, W , means either a payment of Rs. 5 (in games 1 and 2), or a payment of Rs. 10 (in games 3 and 4). A loss, $-W$, denotes a payment of zero. Figure 12.9 shows the two 4×4 games, 1 and 3. In these games, each player has four possible strategies labeled 1 through 4. A payoff of pW , $0 < p < 1$ (as in the bottom right hand corner), denotes a lottery in which the player wins with probability p and loses with probability $1 - p$. In the mixed strategy equilibrium for the 4×4 game, players play their strategies, 1 through 4, with respective probabilities $\frac{3}{8}, \frac{1}{4}, \frac{3}{8}, 0$.

The 6×6 payoff matrix for games 2 and 4 is shown in Figure 12.10. Each player now has six strategies labeled 1 through 6. The 4×4 and the 6×6 games are identical because strategy 3 for

²⁴ In total, they ran 13 experiments with 12 subjects each of whom played 2 practice and 5 real games.

	1	2	3	4
1	W, -W	0, 0	0, 0	W, -W
2	0, 0	0, 0	W, -W	W, -W
3	0, 0	W, -W	$(1/3)W, -(1/3)W$	$(1/3)W, -(1/3)W$
4	0, 0	0, 0	$(2/3)W, -(2/3)W$	W, -W

Figure 12.9 The payoff matrix for games 1 and 3.

Source: Reprinted from *Games and Economic Behavior*, 19: 97–132, Mookherjee, D. and Sopher, B. “Learning and decision costs in experimental constant sum games.” Copyright © 1997, with permission from Elsevier.

	1	2	3	4	5	6
1	W, -W	0, 0	0, 0	0, 0	0, 0	W, -W
2	0, 0	0, 0	W, -W	W, -W	W, -W	W, -W
3	0, 0	W, -W	0, 0	0, 0	W, -W	0, 0
4	0, 0	W, -W	W, -W	0, 0	0, 0	0, 0
5	0, 0	W, -W	0, 0	W, -W	0, 0	W, -W
6	0, 0	0, 0	W, -W	0, 0	W, -W	W, -W

Figure 12.10 The payoff matrix for games 2 and 4.

Source: Reprinted from *Games and Economic Behavior*, 19: 97–132, Mookherjee, D. and Sopher, B. “Learning and decision costs in experimental constant sum games.” Copyright © 1997, with permission from Elsevier.

the row and column players in the 4×4 game is identical to the joint play of strategies 3, 4, and 5 for the row and column players in the 6×6 game. We indicate this by shading the appropriate region in Figure 12.10. The game was played for 40 periods by 10 pairs of subjects. The games were played repeatedly because the authors were also interested in exploring the learning implications for minimax play.

In each of the two games in Figures 12.9 and 12.10, strategy 2 is dominated by the last strategy (strategy 4 in the 4×4 game and strategy 6 in the 6×6 game). If we eliminate the dominated strategy, 2, the game reduces to the one in Rapoport and Boebel (1992). The unique, symmetric, mixed strategy equilibrium of the 6×6 game is for each player to play the six strategies 1 through 6 with the respective probabilities $\frac{3}{8}, \frac{1}{4}, \frac{1}{8}, \frac{1}{8}, \frac{1}{8}, 0$. Notice that the equilibrium is symmetric but the game is not.

The percentage choices of strategies in the four games are given in Table 12.5. The prediction of MSE is also shown in brackets in the row labeled Minimax. The results can be summarized as follows:

1. The actual frequencies of play are substantially different from the minimax predictions for all games (except for the column players in Game 4). This contrasts with the results of O'Neill (1987). As in Brown and Rosenthal (1990), using a chi-square test, Mookherjee and Sopher (1994) are able to reject the null hypothesis of equality between the actual and MSE predicted frequencies. They conclude (p. 109) that: “The rejection of the minimax theory at the aggregate level, and not just at the level of individual decision making, implies that it cannot even be treated as an acceptable as if theory.”
2. While minimax is rejected for all games, it does better than equiprobable play (i.e., playing all strategies with identical probabilities). This is not surprising because there is a clear dominated strategy in each game that is easy to spot.
3. As in Brown and Rosenthal (1990), minimax is also rejected on the grounds that players' strategies in various rounds are not independently and identically distributed.

Table 12.5 Percentage of subjects choosing various strategies in the 4×4 and 6×6 games.

Player	Choice 1	Choice 2	Choice 3	Choice 4	Choice 5	Choice 6
Game 1						
Row	28.75	20.50	44.00	6.75		
Column	42.25	29.25	25.25	3.25		
(Minimax)	(37.25)	(25.00)	(37.25)	(0.00)		
Game 2						
Row	42.50	24.00	4.50	8.25	17.75	3.00
Column	39.75	30.00	7.25	16.50	5.75	0.75
(Minimax)	(37.25)	(25.00)	(12.50)	(12.50)	(12.50)	(0.00)
Game 3						
Row	37.75	13.25	42.25	9.75		
Column	34.25	32.35	28.75	4.75		
(Minimax)	(37.25)	(25.00)	(37.25)	(0.00)		
Game 4						
Row	39.50	24.25	5.00	5.50	21.25	4.50
Column	33.75	23.75	12.50	16.75	10.25	3.00
(Minimax)	(37.25)	(25.00)	(12.50)	(12.50)	(12.50)	(0.00)

Source: Mookherjee and Sopher (1997).

4. Over time, there is no tendency of play to converge to minimax. However, subjects learn to avoid the dominated action as they play successive rounds of the game. As in some of the findings for coordination games, the choice of strategies during the early stages of the game established important precedents on the divergence between minimax and actual frequencies. This divergence tended to persist over successive rounds.

Rapoport and Boebel (1992) use the same experimental design in the 4×4 zero-sum game, as in O'Neill (1987), but extend it to a 5×5 zero-sum game in which each player has five actions. They reject the minimax solution (and so also the MSE) at the level of the individual and in the aggregate. This is not surprising as these games involve even greater cognitive effort on the part of subjects in experiments.

Budescu and Rapoport (1994) consider the two-player, 3×3 matching pennies game shown in Figure 12.11. Players simultaneously choose between three actions, 1, 2, and 3. In the event that all three actions match, the column player pays two pennies to the row player, otherwise the row player pays a penny to the column player. In the unique MSE, each player plays his strategies with an identical probability of one third.

There are two treatments. In the *dyadic treatment* (D), 60 subjects, who are paired randomly, played 200 rounds of the game in Figure 12.11. The games were played face to face, but communication or coordination was forbidden. In the random treatment (R), subjects were placed in a decision theoretic environment. They were given a dice with only the numbers 1, 2, and 3 on it. Each number was printed on two faces of the dice (so the numbers 4, 5, and 6 were absent). The probability of each number coming up on any single throw of the dice is one third, which is identical to the MSE that subjects are theoretically supposed to play in treatment D. Subjects in

	1	2	3
1	2, -2	-1, 1	-1, 1
2	-1, 1	2, -2	-1, 1
3	-1, 1	-1, 1	2, -2

Figure 12.11 A three-player matching pennies game.

Source: "Subjective randomization in one- and two-person games," Budescu, D. V. and Rapoport, A., *Journal of Behavioral Decision Making*, 7: 261–78. Copyright © 1994 John Wiley & Sons, Ltd.

treatment D were asked to simulate the outcomes on the dice when it is rolled 200 times in an unbiased manner. For any period, when the strategies match, the second player gave two pennies to the first player, otherwise the first player gave a penny to the second player. Under perfect randomization, we would obtain statistically identical outcomes in the two treatments, D and R.

It turns out that the actual choices made by players are influenced not only by their own past choices over the last two periods, but also by the last two choices of their opponent. A logit regression is run in which the independent variables are related to several variables, such as the past play of a player and the opponent, and past success of a particular strategy combination. A statistical test of the hypothesis that all coefficients on all independent variables are zero (which is required by minimax) is rejected for 51.7% of the subjects in treatment D. Thus, players' choices are not statistically independent over time. In treatment R, on the other hand, this figure rises to 73.3%. Thus, as in Rapoport and Budescu (1992) and Budescu and Rapoport (1994), players randomize better in game theoretic environments, relative to decision theoretic environments. One potential reason for this finding (relatively greater temporal dependence in decision theoretic tasks) is that the game theoretic task involves relatively greater use of limited cognitive facilities. Hence, subjects do not recall very well their own and the opponent's past actions, which paradoxically helps them to randomize better.

McCabe et al. (2000) test for the play of an MSE in a three-player matching pennies game. Each player is assigned a neighbor and can choose either heads or tails in the toss of a coin. Payoffs depend on a match with the neighbor's choice. The authors distinguish between two conditions. A *limited information condition* in which players are not given information about the entire history of play, and a *full information condition* in which they are given this information. In the full information condition, players are able to observe the actions of the other player at the end of each round. Hence, if they wished to, they could have all the information that they would need to apply sophisticated Bayesian reasoning. Evidence rejecting MSE is found in the limited information condition, but in the full information condition, on average, play comes close to the predictions of Nash equilibrium in some respects. However, the authors reject the hypothesis that players use sophisticated Bayesian learning to reach the MSE.

Shachat and Swarthout (2008) consider two games shown in Figure 12.12. The first (on the left) is a zero-sum symmetric matching pennies game, and the second game, on the right, is named as the *unprofitable game*.²⁵ The unique Nash equilibrium of the matching pennies game is for the row and the column players to play Left with probability $2/3$. Their respective equilibrium expected payoffs are respectively $2/3$ and $-2/3$. The unprofitable game, unlike most of the other games used to test the minimax hypothesis is not zero-sum, hence, the minimax and Nash solutions do not necessarily coincide. The minimax solution in the unprofitable game is for each player to play Right. The unique Nash equilibrium is for each player to use the equiprobable strat-

²⁵ The matching pennies game is taken from Rosenthal et al. (2003), while the unprofitable game was introduced by Shachat and Swarthout (2004).

	Left	Right
Left	1, -1	0, 0
Right	0, 0	2, -2

	Left	Right
Left	2, 0	0, 1
Right	1, 2	1, 1

Figure 12.12 The matching pennies and the unprofitable game.

Source: Shachat, Jason and J. Todd Swarthout, "A hidden Markov model for the detection of pure and mixed strategy play in games," *Econometric Theory*, volume 31(4): 729–52, reproduced with permission.

egy, i.e., play both strategies with probability one half. Interestingly, in the unprofitable game, both minimax and MSE give an identical expected payoff of 1. One may conjecture that players should not play MSE in the unprofitable game because by playing Right they can guarantee a payoff of 1, while MSE offers no improvement in the expected payoff yet it is more risky (Aumann, 1985).

The two games were played by 60 undergraduate students, evenly divided between the two games, for 200 periods. For the first 23 repetitions of the game, row and column players played against each other anonymously, entering their choices through a computer terminal. For the remaining periods, unknown to the subjects, their opponents were switched from humans to computer machines that played the MSE strategies. The results are as follows.

1. Statistically, z-tests can be used to reject the equality of the proportion of players who play Left with the predicted proportions (0.67 for the matching pennies and 0.5 for the unprofitable game). There is strong rejection of MSE in each case. A chi-square test rejects the null hypothesis that subjects use the same mixed strategy. There is substantial heterogeneity of choices across the subjects.
2. When statistical tests are conducted at the level of individual subjects, MSE is rejected for 41 out of 60 subjects. Serial independence of play is statistically rejected for 26 out of 60 subjects.

In a novel study, Okano (2013) asks if teams of two individuals are more likely to conform to the predictions of an MSE. He considers two games in which teams of two players, and individual players, play against each other. The two games are the game in O'Neill (see Figure 12.8) and the pursuer–evader game (see Figure 12.14); each game has a unique MSE. When data is aggregated across all players, the behavior of teams and individuals is similar and support for the play of MSE is not statistically significant. However, when disaggregated data is examined at the level of the team and the individual, the behavior of teams is in better accord with the predictions of MSE, as compared to that of individuals. This is an interesting result because the disaggregated data has typically provided lesser support for the predictions of MSE. However, the problem of serial correlation of choices remains, although, perhaps surprisingly, the serial correlation is lower for individuals.

12.3.2 *Do people play MSE in games within an economic context?*

The games in Section 12.3.1 were mostly abstract. In this section, we consider games that have an economic context but are nevertheless played with subjects in experiments. The main finding is that when data are examined at the aggregate level, there is mixed support for the MSE, but at the individual level, MSE fails.

Collins and Sherstyuk (2000) consider a model of spatial competition with three identical firms. Consumers are distributed uniformly over an interval of length $[0, 1]$ and they are simply




Figure 12.13 Distribution of locations for pooled data.

Source: "Spatial competition with three firms: an experimental study," Collins, Richard and Sherstyuk, Kalerina, *Economic Inquiry*, 38: 73–94. Copyright © 2007, John Wiley & Sons.

interested in buying from the nearest firm (say, to minimize transportation costs). In the classic two-firm model in Hotelling (1929), the Nash equilibrium in the choice of location for both firms is to locate identically in the middle, at location 0.5. However, in the case of three firms, it is not an equilibrium for all firms to locate in the middle. If they did so, then each would get one third of the market. However, any of the firms can now move infinitesimally in either direction to capture half the market. The equilibrium in the case of three firms is derived in Shaked (1982). We state his result without proof.

Proposition 12.3 (Shaked, 1982): *In a model of spatial competition with three firms and consumers distributed uniformly over the unit interval, each firm chooses a mixed strategy in which it randomizes uniformly over the set of locations in the interval $[0.25, 0.75]$. It is never optimal for any firm to choose a location outside this interval.*

In the experiments implemented in Collins and Sherstyuk (2000), the interval $[0, 1]$ was extended to $[0, 100]$, where the distance is in units of kilometers. A total of 9–12 subjects in the experiments participated in each of the four experimental sessions. Each session lasted for 35 periods, preceded by three practice rounds. Subjects played the role of firms in groups of threes, taking as given the following information. A single consumer is located at each integer location in $[0, 100]$. Each consumer had an inelastic demand of 10 units/period and chooses to buy from the closest firm. If two firms are located equidistant from the consumer, then the consumer buys 5 units each from them (provided that the third firm is not located any nearer, in which case it would buy from the third firm). There was random rematching of players from round to round. In any period, the three subjects participating in the experiment, simultaneously choose their location. The pooled results are shown in Figure 12.13.

Table 12.6 Results for the symmetric condition.

Investment	Empirical distribution			Equilibrium prediction
	Group 1	Group 2	Both groups	
Game L : low reward condition ($e=5, r=8$)				
$c_i=0$	0.116	0.222	0.169	0.125
$c_i=1$	0.116	0.117	0.116	0.125
$c_i=2$	0.096	0.079	0.088	0.125
$c_i=3$	0.084	0.153	0.118	0.125
$c_i=4$	0.094	0.086	0.090	0.125
$c_i=5$	0.494	0.343	0.418	0.375
Game H : high reward condition ($e=5, r=20$)				
$c_i=0$	0.135	0.147	0.141	0.050
$c_i=1$	0.065	0.044	0.055	0.050
$c_i=2$	0.053	0.054	0.053	0.050
$c_i=3$	0.040	0.066	0.053	0.050
$c_i=4$	0.044	0.094	0.069	0.050
$c_i=5$	0.662	0.594	0.628	0.750

Source: Rapoport and Amaldoss (2000).

For the pooled data, 80% of the choices of subjects are within the interval $[0.25, 0.75]$. The aggregate distribution of choices is bimodal with peaks at 40 and 70. Hence, the prediction of the relevant theory that the distribution should be uniform over $[0.25, 0.75]$ does not hold. This is also borne out by the statistical results. The difference between the observed and predicted frequencies is statistically significant. The hypothesis that the pooled sample is drawn from the Nash equilibrium distribution given in Proposition 12.3 is rejected by a Kolgomorov–Smirnov test at the 5% level.

The individual choice of locations of 36 out of 51 subjects are also inconsistent with the predicted Nash distribution at the 5% level. As in many other experiments that test for MSE, there is substantial variation in individual choices from the predicted choices. Insofar as most subjects locate in the interval $[0.25, 0.75]$, one might claim some partial vindication for MSE. But one has to be cautious in making such inferences. As Brown and Rosenthal (1990) showed in their critique of O'Neill's study, a range of other behavioral models might potentially rationalize the data in such a manner.

Rapoport and Amaldoss (2000) consider an investment game between two risk-neutral players in their roles as firms, indexed by $k = 1, 2$. Firms have identical, commonly known, endowments equal to e each and they simultaneously invest amounts $c_k \in \{0, 1, 2, \dots, e\}$ that are non-recoverable (think of these as R&D costs). This is known as the *symmetric condition*.²⁶ The firm that makes the higher investment wins a prize $r > e$ (say, patent rights). Conditional on the profit functions, in an MSE, players invest their endowment, e , with probability $\frac{r-e}{r}$ and choose all other investment levels smaller than e with an identical probability $\frac{1}{r}$.

²⁶ The authors consider two conditions. In condition 1, players are symmetric, so $e_1 = e_2 = e$. In condition 2, players are asymmetric and $e_1 > e_2$, so we may call firm 1 the strong player and firm 2 the weak player. Under both conditions there is an equilibrium in mixed strategies, but these equilibria differ. Here we describe results only for the symmetric case.

Given the results in simpler games in Section 12.3.1 and earlier results in Section 12.2, one may be skeptical of the ability of subjects to play the MSE. A total of 36 subjects were divided into two groups (Group 1 and Group 2) and they played the symmetric game, described above, for 80 successive periods. The results are shown in Table 12.6 for the two cases of a low reward condition ($e = 5$ and $r = 8$) and a high reward condition ($e = 5$ and $r = 20$).

At the individual level, MSE is rejected. However, at the aggregate level there is support for the MSE. The failure at the individual level is sufficient to reject the theory, because the predictions of theory are strictly at the level of the individual players.²⁷ Nevertheless, the success of MSE at the aggregate level is intriguing and deserving of further research. We summarize the findings in a bit more detail for the symmetric case and also the asymmetric case in which the endowments of the firms are different.

1. In the symmetric game, the investment strategy $c = e$ was chosen most often, which is consistent with the predictions. For the low reward game, Game L, in which $r = 8$, the actual average and predicted frequencies of play are statistically similar. For Game H, in which $r = 20$, the investment frequencies of Group 1 subjects statistically conform to the MSE predictions in the aggregate, but those of Group 2 do not. At the level of the individual data, the strategies chosen by players are not iid, contradicting the MSE prediction. A logit analysis shows that strategies depended on the opponent's play as well as the outcome in the previous period.
2. In the asymmetric game, the strategies of the strong players (whose endowment is relatively greater) support the MSE in the aggregate but the strategy choices of the weak players do not. Hence, even at the aggregate level, there is not unqualified support for MSE.²⁸

The thrust of these results, namely, support (albeit not unqualified) for the MSE at the aggregate level but rejection of MSE at the individual level continues in some of the other work that these authors are involved in.²⁹

12.3.3 *Do people play MSE when they have a mixed strategy device?*

Recall the sixth criticism made by Brown and Rosenthal (1990) in Section 12.3.1. The experimenter only observes the actual realization of actions of the two players but not the underlying probability distribution from which the actions may have been drawn. There are at least two reasons why subjects in repeated experiments may fail to produce the sequence of strategies that are consistent with an MSE. One possibility is that subjects simply do not play an MSE. The other possibility is that subjects would like to play MSE but are unable to generate iid sequences, say, on

²⁷ Unless, of course, one invokes the "equilibrium in beliefs" interpretation, which is not supported by the evidence (see above).

²⁸ Dechenaux et al. (2006) argue that the support for MSE is stronger in the Rapoport and Amaldoss (2000) model relative to the claim. Their main argument is that if alternative equilibria are taken into account, then the range of results for which the data will support the MSE prediction is wider. However, Rapoport and Amaldoss (2008) disagree.

²⁹ Rapoport and Amaldoss (2004) consider an all-pay first price auction. In such auctions, all bidders have to forfeit their bids. Think of players incurring sunk cost to win some prize or contract. The winner is rewarded but the losers are not able to recoup their investments. The findings here are similar to Rapoport and Amaldoss (2000). Furthermore, individual-level data can be reconciled using the Experience Weighted Attraction (EWA) model that we consider in Chapter 14. Amaldoss and Jain (2002) consider a model of patent race in which there is an asymmetric equilibrium in mixed strategies. The aggregate behavior is consistent with the predictions of MSE, but the individual behavior is very varied and is inconsistent with MSE. The EWA model with adaptive learning can account for the individual-level data.

account of cognitive limitations. To address this issue, recent work has provided players with a *randomizing device* in games that have a unique MSE. If there is a failure of MSE because subjects cannot generate iid draws from a multinomial distribution, then the provision of a randomizing device provides a clinching test.

A typical *randomizing device* or *mixed strategy device* works as follows.³⁰ Suppose that a player has two strategies, A and B. The player is then allocated X binary choices in total (e.g., $X = 50$ in the case of Bloomfield, 1994) to be distributed between A, B in the proportions $p = \frac{X_A}{X}$ and $1 - p = \frac{X_B}{X}$, where X_A, X_B are, respectively, the number of choices allocated to strategies A, B, and $X_A + X_B = X$. By using proportions in this manner, the subject reveals that he/she would like to play a mixed strategy in which A, B are played with respective probabilities $p, 1 - p$ in each round of play.

In each round of play, a machine then draws a temporally independent choice of strategy for the player from, say, a Bernoulli distribution with parameter p . Or possibly, a computer may be used to randomly play the mixed strategy choices of players against each other, over time. This allows a player to choose iid strategies over different rounds of play with exactly the intended probabilities. In the presence of this mixed strategy device, any observed correlation of actions over the various rounds cannot be due to an inability to generate iid play. Any departure must then be on account of players not wanting to play an MSE. Thus, the mixed strategy device allows us to observe the mixed strategy (Nash or otherwise) that is played by the players.

The point made in the previous paragraph is rather important. Some economists such as Binmore et al. (2001, p. 448) argue that “but we think it a mistake to demand that the players must actively randomize before the minimax theory can be said to be relevant to their play.” There are two issues here. First, a methodological issue about whether we need the assumptions of a model to be descriptive of the real world or whether we should just care about the predictions. Second, if we grant that the argument by Binmore et al. (2001, p. 448) is valid, then the availability of a mixed strategy device should settle the issue of whether players play minimax.

Shachat (2002) considers experimental results that arise with a mixed strategy device. He corrects for a drawback in the mixed strategy device that was implemented in the earlier work, such as that of Bloomfield (1994) and Ochs (1995).³¹ In Shachat (2002), subjects play the O'Neill (1987) game shown in Figure 12.8 but with the following differences. (i) Pairs of subjects are matched anonymously and play 60 repetitions unlike 105 in O'Neill (1987). (ii) Unlike O'Neill (1987), the subjects do not choose strategies face to face but make choices over computer terminals. (iii) The strategies are renamed as Green, Red, Brown, Purple in order to counter the possible salience of the “Joker” strategy. (iv) Stakes are increased. The row player receives 90 cents for a win and 0 cents for a loss; the corresponding figures for the column player are 60 cents and 0 cents. In the MSE, the stage game earnings of the row and the column player are identical and equal to 36 cents.

³⁰ See, for instance Bloomfield (1994), Ochs (1995), and Shachat (2002). Our usage of the term *mixed strategy device* follows Shachat (2002).

³¹ The drawback is related to the distinction between sampling with and without replacement. Suppose that we wish to implement a mixed strategy profile for a matching pennies game to be played over four rounds. Each player can place four balls in the urn. Since the MSE is to play heads and tails with probability one-half each, players should place two balls marked H and two balls marked T in the urn. A ball is now randomly drawn from the urn for each player, in each period, and the strategies marked on them are implemented in that round. However, in Ochs (1995), the balls are not replaced after each draw. Hence, the probability of drawing an H or a T is not 0.5 in each round. Thus, play in each round does not stay iid. With this caveat in the problems with their mixed strategy devices, both Bloomfield (1994) and Ochs (1995), in 2×2 games, find results similar to O'Neill (1987). Namely, that the marginal frequencies of play for the row and column players are close to the MSE predictions in the aggregate data.

Three treatments are conducted. Treatment I does not involve a mixed strategy device. Treatment II implements actions through a mixed strategy device.³² Treatment III is a disclosure treatment in which at the end of each round, the opponents' mixed strategy and the actual action choice are revealed.

For all three treatments, the results are similar to those in Table 12.4. We show the results in Table 12.7 for all three treatments, retaining the strategy names as in O'Neill (1987) for purposes of comparison. In each cell of Table 12.7, the results of Treatments I, II, and III are given in rows 1, 2, 3, while the final row gives the MSE prediction in brackets. The basic results are as follows. (i) Looking at the entries for each cell, the actual frequencies of play are close to the MSE predictions for all treatments, (ii) however, a chi-square test rejects the hypothesis that the observed frequencies are close to the MSE predictions at the 5% level, and (iii) furthermore, the data on win rates for the row players (not shown in Table 12.7) is also close to the MSE prediction of 40% as in O'Neill (1987) and Brown and Rosenthal (1990).

We now move on to other results in Shachat (2002).

Table 12.7 Frequencies of card choices in various treatments.

	Ace	2	3	Joker	Marginal row frequencies
Ace	0.044	0.043	0.047	0.068	0.202
	0.047	0.058	0.034	0.072	0.211
	0.061	0.047	0.037	0.078	0.222
	(0.040)	(0.040)	(0.040)	(0.080)	(0.200)
2	0.047	0.037	0.048	0.093	0.224
	0.047	0.055	0.045	0.072	0.219
	0.042	0.043	0.039	0.100	0.224
	(0.040)	(0.040)	(0.040)	(0.080)	(0.200)
3	0.050	0.042	0.031	0.091	0.214
	0.044	0.042	0.053	0.067	0.205
	0.042	0.047	0.034	0.087	0.210
	(0.040)	(0.040)	(0.040)	(0.080)	(0.200)
Joker	0.073	0.061	0.068	0.157	0.359
	0.080	0.073	0.065	0.147	0.365
	0.071	0.074	0.061	0.137	0.343
	(0.080)	(0.080)	(0.080)	(0.160)	(0.400)
Marginal column frequencies	0.214	0.183	0.193	0.409	–
	0.218	0.228	0.197	0.357	
	0.217	0.211	0.171	0.401	
	(0.200)	(0.200)	(0.200)	(0.400)	

Source: Shachat (2002).

³² Each experimental subject fills a *shoe* (device for holding multiple cards) of 100 cards with any composition of the four possible actions reflecting the mixed strategy that the subject wishes to play. The computer then shuffles the shoe in each round and randomly picks the card on the top. The strategies corresponding to the chosen cards are then played.

1. The row player can either win or lose in each round and the probability of winning a round is $p = 0.40$. Thus, we have a binomial distribution of win rates with parameter $p = 0.40$. Over n trials, the predicted variance of the number of wins for the row player is given by $np(1 - p)$. A chi-square test rejects the hypothesis that the sample variance and the predicted variance are the same; in particular, there is excessive sample variance relative to the predicted variance.³³ Similarly, one can statistically reject (1) the binomial probability implied by minimax for individual cards, and (2) the joint action profiles over actions for about 2/3 of the pairs in each treatment. Shachat (2002, p. 206) concludes: “Thus, without much hesitation, one can conclude that the presence of a device that allows subjects to easily randomize over their action sets does not bring the data of observed actions plausibly close to minimax play.”
2. Serial correlation is reduced relative to the findings of O'Neill (1987); however, it still remains. In particular, there is evidence that a player's choice of strategy depends on the previous play by the opponent; a clear indication that at least one of the players uses non-minimax play.
3. Many players seem to be using mixed strategies. However, very few players seem to be using mixed strategies as predicted by the MSE. In total, there were 3,720 decisions made in Treatments I and II. Of these, only 103 (2.8%) correspond to minimax play. On the other hand, equiprobable play in which each pure strategy is played with an identical probability, was more common and represented 13.2% of the choices. However, when subjects play mixed strategies, they play a wide variety of non-minimax strategies. Like Brown and Rosenthal (1990), Shachat (2002) finds that a variety of non-minimax strategies can lead to winning percentages for the row player that are close to the minimax prediction of 40%.

12.3.4 *Do professionals play mixed strategy equilibria?*

Formal economic models simplify all but the most essential aspects of reality in order to explain economic phenomena. Suppose that an economic model makes an MSE prediction that we wish to test in the field. These tests are confounded by at least two important factors. (1) Whether subjects in the field play mixed strategies, as predicted by the simple models. (2) Is the simple economic model a realistic description of the real world? By controlling for (2) in experiments, and by mapping experimental reality to payoffs, we are able to control one source of the confound. However, in experiments, we might not be able to control well for the first source of the confound (ecological validity), hence, giving rise to a trade-off from using experimental data.

One method of addressing the *ecological validity* concern with experimental tests of an MSE is to conduct the experiments with particularly suitable subject pools (artefactual experiments). Ideally, we are interested in a professional subject pool that has the relevant experience and skills.³⁴ Which subject pools and situations should one be looking out for to test for an MSE? The answer is not easy.

³³ The predicted variance is 14.4 for all three treatments and the sample variance for the three treatments is respectively 42.67, 27.4, and 24. The respective p-values are 0.000, 0.029, 0.025. The degrees of freedom are respectively 15, 16, 15.

³⁴ Broadly speaking, ecological validity refers to conducting an experiment in an environment which is suitable for testing some theory. For instance, if we want to detect paternal instinct in humans, then a laboratory experiment with students in which they assume roles of parents and children might not satisfy ecological validity, unless properly controlled for.

One possible source of field data that can be used to test the predictions of MSE is data from professional sports. Although this data source is getting increasing attention, one also needs to consider the potential caveats. Playing successful strategies based on experience in ball sports involves using motor skills that have a basis in our evolutionary history. As in Friedman's famous example, from the play of a successful billiards player, we cannot conclude that the player understands the complicated geometry that lies behind the actions. Rather, Friedman argued, the player behaves *as if* he/she was familiar with such geometry.

Could a successful billiards player solve an interesting economic problem that involved use of complicated geometry? One's a-priori hunch is that this is unlikely. The main reason is that while evolutionary skills may help us to succeed in ball sports, reflex actions, and hand-eye coordination, it is not clear if evolution has enabled us to solve a range of interesting economic problems. If we find that professional sports players use skills developed through evolution to play MSE in sports, we cannot necessarily conclude that the ability to play MSE transfers successfully to other economic problems. It would, of course, be very worthwhile to test this argument; we consider these issues below.

MINIMAX PLAY BY PROFESSIONAL SPORTS PLAYERS

It is not straightforward to use data from professional sports to test for an MSE. Walker and Wooders (2001) who are early pioneers in this field make the following observation (p. 1522):

As empirical tests of the minimax hypothesis, however, sports are generally inferior to experiments. In the classic confrontation between pitcher and batter, for example, there are many actions available (e.g., fastball, curve, change-up), and the possible outcomes are even more numerous (e.g., strike, ball, single, home run, fly ball, double-play grounder). Difficulties clearly arise in modeling such situations theoretically, in observing players' actions, and in obtaining sufficient data to conduct informative statistical tests.

With this caveat, Walker and Wooders (2001) test the minimax predictions, using as their observations the serves of top tennis players.

In tennis, a player can either serve to the left or to the right of an opponent. The opponent can simultaneously anticipate and also move to the left or to the right. Each serve has two outcomes: one player wins the point and the other loses. But, in actual tennis matches there are the following additional factors that are relevant, yet difficult to take into account. (1) Players can vary spin, angle, slice, or speed on the serve. (2) The receiver of a serve can stand far too much to the right (or left) and force the server to serve on the left (or right) or may engage in other body feints. (3) If the first service is faulty, then players are allowed a second serve. So, a two-dimensional strategy space (serve to the left or to the right) is not an innocuous simplification. Walker and Wooders (2001) consider payoff matrices of the following form.

	L	R
L	$\pi_{LL}, 1 - \pi_{LL}$	$\pi_{LR}, 1 - \pi_{LR}$
R	$\pi_{RL}, 1 - \pi_{RL}$	$\pi_{RR}, 1 - \pi_{RR}$

The row player (server) has two choices: serve left or serve right. The column player (receiver) has to guess which way the serve will fall, and in anticipation, move in that direction. Thus, the receiver also has two choices: overplay to the left or to the right. π_{ij} is the probability that the row player wins the point by serving in the direction $i = L, R$ when the receiver overplays in

the direction $j = L, R$. Analogously, $1 - \pi_{ij}$ is the probability that the column player wins under these conditions. Assuming that players care about winning, the payoff is directly proportional to the winning probability. Hence, in a reduced form manner, the payoffs are simply the probabilities of winning. This allows the authors to convert the problem of tennis serves into a two-player zero-sum game. The probabilities π_{ij} are allowed to vary from one match to another and during different stages of the match, such as deuce points and normal play.

In practice, it is difficult to know what π_{ij} should be. Hence, there is no available theoretical distribution with which one might compare the outcomes of actual play. So, Walker and Wooders (2001) test an implication of MSE, namely, that all pure strategies in the support of the equilibrium mixed strategy should have identical payoffs. In the tennis context, this means that serves to the left or to the right should have identical win rates. Data is obtained from videotapes of ten tennis matches from the major Grand Slam tournaments as well as the Master's championship conducted at the end of the tennis year; eight of these matches were the tournament finals.

The results are as follows. For individual matches, the Pearson test is used to test the null hypothesis that the win rates for left and right serves are identical. The value of the test statistic is 30.801, while the p-value is 0.852. Thus, the null hypothesis cannot be rejected at any level of significance. The predicted distribution (identical win rates on left and right serves) is compared with the actual distribution using the Kolmogorov–Smirnov test; the two distributions turn out to be statistically very close. However, the hypothesis of serial independence of first serves is rejected. Players switch their serves too frequently. This is in line with other research in which people find it difficult to generate random sequences. However, while there is serial dependence in play, it seems not to be exploited by the opponents. This result also holds for several other studies that we consider in this chapter; however, some contrary evidence is given in Scroggin (2007).

Palacios-Huerta (2003) tests the idea in Walker and Wooders (2001) for penalty kicks in professional soccer games. There may be some advantages of using data from penalty kicks over tennis serves in testing for minimax play. Issues of spin on the ball are less serious and there are no second kicks (unlike the second serves in tennis). On average, it takes just 0.3 seconds for the ball to travel to the goalpost. Hence, the goalkeeper makes an almost simultaneous decision to react. As in tennis, there are two possible directions: left and right for the kick and for the movement of the goalkeeper.³⁵ The only two possible outcomes are: goal or no goal. The vast majority of kickers are right footed and their natural shooting side is to the right of the goalkeeper. The natural shooting side of the left footed kickers is to the left of the goalkeeper. The goalkeepers take account of the kickers' natural side when they move.

Data is taken on 1,417 penalty kicks over the period 1995–2000 from professional football games. The payoff matrix is as in Walker and Wooders (2001). A little over 80% of the penalty kicks result in a goal. When the goalkeeper moves in the incorrect direction, nearly 100% of the kicks result in a goal. And when he moves in the correct direction, then 60% of kicks result in a goal. The null hypothesis, i.e., that scoring probabilities are identical on both sides of the goalkeeper is rejected at the 5% level for only 3 out of 42 players. Thus, the strategies of the vast majority of players are consistent with minimax. The Kolmogorov–Smirnov test shows that the null hypothesis of aggregate play being identical to the minimax predictions cannot be rejected. Finally, the null hypothesis that the choices of players are serially independent is rejected for only 2 of the 42 players. This result runs counter to the bulk of the evidence elsewhere that shows a

³⁵ The kicker can also kick straight to the center. However, these kicks constitute a very small part of the sample; 7.5% of kickers kick to the center and 1.7% of goalkeepers choose the center (i.e., stay still and not move either to the left or to the right).

serial dependence in actions. One possible explanation is that successive penalty kicks, unlike other experiments done in the lab, are spaced over a relatively large time interval. Chiapori et al. (2002) use a sample of 459 penalty kicks from the French and the Italian soccer leagues and reach conclusions that are similar to Palacios-Huerta (2003) in their support for the minimax hypothesis. However, they consider the role of heterogeneity in more detail.³⁶

The actions chosen by professional sports players may not rely on the same degree of rationality that is embodied in the play of MSE. For instance, based on field data, Bar-Eli et al. (2007) find that goalkeepers could save more penalty kicks if they stood still. Yet, we observe that they typically dive to the left or the right. They propose an explanation based on *norm theory*. It is the norm for goalkeepers to dive because in the event that a goal is scored, they feel worse from inaction (standing still) rather than from action (diving).

IS MSE PLAY TRANSFERABLE ACROSS CONTEXTS FOR PROFESSIONAL SPORTS PLAYERS?

Do experienced professional sports players play mixed strategies and can they transfer their experience to other contexts? Palacios-Huerta and Volij (2008) (henceforth, PH-V) give a strongly positive answer to the question of the portability of the successful play of mixed strategies from one context to the other, but this has been successfully overturned by the results in Wooders (2010). In a persuasive piece of empirical work, Levitt et al. (2010) (henceforth, LLR) reject portability of MSE across contexts.

LLR consider the following two zero-sum games. These games are similar to those in PH-V. The first game, the 2×2 pursuer–evader game, taken from Rosenthal et al. (2003), is shown in the upper panel of Figure 12.14. Players simultaneously choose among the two strategies, Black and Red. If the strategies do not match, the row player (pursuer) gets 1, while the column player (evader) gets 0. If the strategies match, then a six-faced dice is thrown. The payoffs for each player, either 0 or 1, then depend on the outcome of the throw of the dice, as shown in the figure. The MSE of the game is for each player to play Black and Red with respective probabilities $2/3$ and $1/3$.

PH-V consider a similar game that is based on the empirical success percentages from penalty kicks in soccer. The equilibrium mixing proportions in their case are $36 : 64$ for the row player (the pursuer) and $55 : 45$ for the column player (the evader). Insofar as mixing in the proportions $2/3 : 1/3$ is less cognitively challenging, the LLR study gives portability a better chance of success (but theory fails as we shall see below).

The second game used in LLR is a 4×4 game, taken from O'Neill (1987); see the lower panel of Figure 12.14. PH-V also use the same game, but strategies are named after colors (red, brown, purple, green). Each game was played a large number of times. For instance, in LLR, the 4×4 game was played 150 times with players switching roles between column and row players after 75 rounds.

PH-V use two sets of subjects in the experiment: Professional sports players and college students. LLR use undergraduate students and three sets of professional sports players: bridge players, poker players and major league soccer players.³⁷ Mixed strategies do not play an important

³⁶ Heterogeneity can arise, for instance, from the different abilities of the penalty kickers and the goalkeepers, and from the differing abilities to cope with stress, the game situation, and so on. Thus, in general, not all predictions at the level of specific kicker–goalkeeper pairs may survive aggregation to all players. On the other hand, one may examine the role of additional assumptions that allow for such aggregation to take place; this is the route followed in Chiapori et al. (2002).

³⁷ The bridge and poker players were recruited respectively at the 2006 North American Bridge Championship in Chicago, and the 2006 World Series of Poker in Las Vegas. The soccer players were recruited from three major league soccer teams (Los Angeles Galaxy, Chivas USA, and Real Salt Lake).

Black		Red		
Black	0, 1 (if the dice comes up 1 or 2) 1, 0 (if the dice comes up 3, 4, 5, 6)	0, 1 1, 0 (if the dice comes up 1 or 2) 0, 1 (if the dice comes up 3, 4, 5, 6)		
Red	0, 1			
Pursuer–Evader game				
	Club	Heart	Spade	Diamond
Club	1, 0	0, 0	0, 0	1, 0
Heart	0, 0	1, 0	0, 0	1, 0
Spade	0, 0	0, 0	1, 0	1, 0
Diamond	1, 0	1, 0	1, 0	0, 0
The 4 × 4 game				

Figure 12.14 The pursuer–evader game and the 4 × 4 game in Levitt et al. (2010).

Source: Levitt et al. (2010) with permission of The Econometric Society.

role in bridge, however, they are critical in the play of poker. Hence, poker players should have, in principle, sufficient familiarity with the basic idea. The results show that bridge players do not play MSE, hence, in what follows, the results for bridge players are not given.³⁸

Some of the subjects in experiments played against a pre-programmed computer. The computer used two alternative algorithms. (1) An *optimal algorithm* that is designed to play minimax for the first 15 periods and then, in subsequent periods, designed to exploit any suboptimal play by human subjects. (2) An *exploitable algorithm* that simply played all strategies with equal probabilities.³⁹ In LLR, all subjects in experiments (except professional soccer players) played the 2 × 2 game and the 4 × 4 game with humans and with the computer.⁴⁰ There is a unique mixed strategy equilibrium in both these games.

For the data to support the MSE predictions, the following must hold. (1) At the individual level, play should conform to the proportions predicted by MSE.⁴¹ (2) All pure strategies in the support of the mixed strategy should give identical payoffs. (3) In successive rounds of play, the actions of players should be serially uncorrelated. Table 12.8 shows the comparative results of LLR and PH-V for these four predictions for the 2 × 2 game.⁴²

³⁸ Levitt et al. (2010) also suppress the results for the bridge players. However, these can be found in the supplementary material for the paper or in an earlier working paper version of the paper dated 2007.

³⁹ Subjects were told in the first case that “we have programmed the computer to play the theoretically correct strategy in this game. In addition, any deviations that your play has from this correct style of play will be taken advantage of by the computer.” The instructions in the second case omitted any mention of optimal play by the computer.

⁴⁰ Because of time constraints, the soccer players played only the 4 × 4 game.

⁴¹ This is a sufficient condition for play at the aggregate level to conform to the proportions predicted by MSE.

⁴² Here is a full description of Table 12.8 from Levitt et al. (2010, p. 1421): “The columns correspond to the different subject pools tested, while the rows report results for each test. The last two columns report results for a similar experiment carried out by Palacios-Huerta and Volij (2008). Panel I shows p-values from Pearson’s chi-square test for goodness of fit of aggregate frequencies to minimax predictions. p-values for the marginal frequencies of the pursuer and evader are shown in the first two rows, while the third row shows p-values for combinations of plays by both players. The test uses one degree of freedom for the marginal distribution of play and three for the joint distribution. Panel II shows the percentage of individuals (or pairs) that we reject at the 5% level for this same chi-square test. Panel III presents the percentage of players for whom we can reject the null hypothesis of no serial correlation in actions, based on the runs test of Gibbons and Chakraborti (2003).”

Table 12.8 Results for the 2×2 pursuer–evader game.

	Source:	Levitt, List, and Reiley		Palacios-Huerta and Volij	
	Test:	College students	Poker players	Soccer pros	Soccer college
# of pairs		11	22	40	40
# pairs of roles		22	44	40	40
I. Minimax play at aggregate level					
Chi-square test for minimax play:					
pursuer (or row player)		<0.001	<0.001	<0.001	<0.001
evader (or column player)		<0.001	<0.001	0.374	<0.001
joint play		<0.001	<0.001	0.001	N/A
II. Minimax play at individual level					
Rejections at 5 percent:					
pursuer		59%	68%	5%	5%
evader		55%	52%	5%	10%
joint play		91%	75%	0%	5%
neither player		27%	41%	0%	N/A
III. Runs tests					
Rejections at 5 percent:					
for too few runs		23%	18%	5%	8%
for too many runs		10%	14%	0%	N/A
		14%	5%	5%	N/A

Source: Levitt et al. (2010).

The null hypothesis tested in Panel I in Table 12.8 (titled Minimax Play at Aggregate Level) is that the aggregate frequencies of play match the MSE prediction. The p -values needed for rejecting the null hypothesis at the aggregate level are shown in the table for the row player (pursuer), the column player (evader), and for joint play. The MSE predictions are rejected in all three cases at the $p < 0.01$ level for all subject pools: undergraduate students and the two sets of professional players (recall that we have suppressed the result for Bridge players). On all relevant yardsticks of comparison, for aggregate play, the LLR data from professionals is statistically indistinguishable from students. By contrast, in the PH-V data, the aggregate play of professional subjects is consistent with MSE.

Panel II in Table 12.8 (titled Minimax Play at Individual Level) focuses on individual-level data. It shows the percentage of players, in each category, for whom one can reject at the 5% level, the null hypothesis, namely that play conforms to MSE. It turns out that 90% of students and 75% of poker players do not conform to the predictions. In PH-V, by contrast, all professional players statistically conform to the MSE predictions.

Panel III in Table 12.8 (titled Runs Tests) shows the percentage of players for whom one can reject the null hypothesis of no serial correlation in play based on the runs test of Gibbons and Chakraborti (2003). These percentages are relatively small, and much smaller than the findings elsewhere, some of which we have reviewed above.

The findings for the 4×4 game are shown in Table 12.9 and parallel the findings of the 2×2 game in Table 12.8. There is a new panel in Table 12.8 titled “Equality of success rates . . .” This panel shows tests for the null hypothesis that the pure strategies in the support of the mixed strategy give identical payoffs and also that the level of these payoffs conforms to MSE predictions. For the LLR

Table 12.9 Results for the 4 × 4 game.

Source:	Levitt, List, and Reiley				Palaclos-Huerta and Volij	
Test	College students	Bridge players	Poker (overall)	Soccer pros	Soccer pros	Soccer college
# Pairs of people	12	11	26	16	20	20
# Pairs of roles	24	22	52	32	20	20
t-text for asymmetry of play by role	0.188	0.074	0.654	0.376	N/A	N/A
I. Minimax play at aggregate level						
Chi-square text for minimax play:						
Row player	0.320	0.362	0.253	0.001	0.956	0.890
Column player	0.008	0.006	0.001	<0.001	0.932	0.964
Joint play	0.105	0.030	0.008	<0.001	0.910	N/A
II. Minimax play at individual level						
Rejections at 5 percent:						
Row player	33.33%	36.36%	26.92%	28.13%	5.00%	10.00%
Column player	45.86%	36.36%	34.62%	15.63%	5.00%	10.00%
Joint play	37.50%	45.45%	30.77%	28.13%	10.00%	5.00%
III. Equality of success rates across strategies and to equilibrium rate						
Rejections at 5 percent	39.58%	45.45%	35.58%	31.25%	7.50%	7.50%
IV. Runs tests						
Rejections at 5 percent	27.08%	10.23%	9.13%	3.91%	2.50%	3.75%

Source: Levitt et al. (2010).

data, for nearly 2/3 of the students and the poker players, the null hypothesis can be rejected.⁴³ In the PH-V data, by contrast, we cannot reject non-conformity with the MSE predictions for any professional sports players and the rejection rate of students is only 10%.

In conjunction, these findings imply that LLR strongly reject minimax play for professional sports players in the lab. Levitt et al. (2010) write: “In sum, we are left to conclude that none of our subject pools, including professional soccer players, play these laboratory games in accordance with the predictions of minimax, even though there is evidence that they likely employ such strategies in the field . . . It appears that for our experimental subjects, what happens in the field stays in the field.” These results are diametrically opposite to the results reported in PH-V.

What accounts for the difference in results between LLR and PH-V? One possibility is that professional players would like to play minimax, however, they correctly anticipate that they are facing a student population that does not play minimax. Hence, a non-minimax strategy is a best reply. This is rejected by observing play against computers which are programmed to play minimax. When playing against a computer that is playing suboptimally, professional sports subjects do slightly better but still do not play minimax. LLR use post-experimental surveys which reveal that the main reason why players did not play minimax is that players did not place the lab games in the correct context. For instance, soccer players were not able to spontaneously make the link between penalty kicks and the MSE of lab games.

⁴³ The rejection rate for bridge players (not shown in the table) is 37%.

Wooders (2010) undertakes a re-examination of the data in PH-V. In PH-V, students and professional soccer players play two kinds of games. A penalty kick game and the original game in O'Neill (1987). Each of the games is played 200 times between pairs of players. The innovation in Wooders (2010) is to break the results down over the 200 rounds into two equal halves. Minimax predicts that the frequency distribution of the choices (arising under the MSE of each of the two games) should be identical across the two halves. This "stationarity property" is violated for professional soccer players who switch from underplaying relative to the MSE prediction in the first half to overplaying in the second half. Even more strikingly, the behavior of students is closer to the MSE prediction relative to the professionals.

It is fair to conclude that the data rejects portability of MSE play for professional sports players.

12.4 Coordination games

Coordination failures provide an important explanation of human behavior in a wide range of areas in economics (Schelling, 1960, 1978). Coordination problems underpin a range of macroeconomic models, particularly those that use imperfect competition. Models of network externalities and standards also rely on the solution to coordination problems. For instance, a technology or process or a standard might become useful and commercially viable only when a certain critical number of users coordinate on adopting it.

12.4.1 Evidence on coordination failures

Early evidence on coordination failures is provided in Cooper et al. (1989). They consider coordination in the following *battle of the sexes* (BoS) game.

	1	2
1	0, 0	200, 600
2	600, 200	0, 0

In this game, both players have two strategies, 1 and 2. The payoffs are given in points. Dividing the points by 1000 gives the probability of winning a prize.⁴⁴ This practice of giving payoffs in points also applies to the other papers, Cooper et al. (1990, 1993), that we consider below. The two pure strategy Nash equilibria are (1, 2) and (2, 1). There is also a mixed strategy Nash equilibrium in which each player plays strategy 1 with probability 0.25. This gives an expected payoff to each player from the mixed strategy of 150 points, which is worse than coordinating on any of the two Nash equilibria. The experimental evidence indicates that the actual observed frequency of play, when aggregated over all players was close to the mixed strategy equilibrium. The two pure strategy Nash equilibria arose less than 42% of the time. Hence, there were widespread coordination failures.

Cooper et al. (1990) provide evidence on coordination failures from two-player games in which each player has a choice between three pure strategies, 1, 2, 3. The generic game is shown in Figure 12.15. They consider four Type-I coordination games in which $z = 600$ and two Type-II games in which $z = 500$; in each case one of x or y is greater than 600.

⁴⁴ So, if the strategy profile (2, 1) is played, then player 1 wins the prize with probability 0.6 while player 2 wins it with probability 0.2.

	1	2	3
1	350, 350	350, 250	$x, 0$
2	250, 350	550, 550	$y, 0$
3	$0, x$	$0, y$	z, z

Figure 12.15 Type-I game ($z = 600$) and Type-II game ($z = 500$).

Source: Cooper et al. (1990) with permission from the American Economic Association.

These two types of games have the following features. (1) All games are symmetric coordination games. (2) The pure strategy Nash equilibria in each game are (1, 1), and (2, 2). (3) The two Nash equilibria can be Pareto ranked: (2, 2) Pareto dominates (1, 1). (4) In Type-I games, the cooperative outcome that maximizes joint payoffs is (3, 3). (5) In Type-II games, the cooperative outcome that maximizes joint payoffs is (2, 2). To enable assertions (4) and (5) to hold, there are appropriate upper limits on x and y .

The shaded part of the payoff matrix is common to all games played. In both types of games, strategy 3 is dominated, so it should have no effect on which Nash equilibrium is played. The results of the experiments are as follows.

(1) In Type-I games, most players play one of the two Nash equilibria, (1, 1), or (2, 2). In games 3, 4, which are Type-I games, subjects overwhelmingly play strategy 1 (in game 3, $x = 1000$, $y = 0$, and in game 4, $x = 700$, $y = 0$). Hence, subjects do not coordinate on the Pareto efficient equilibrium (2, 2). Thus, coordination failures arise in these games. (2) In Type-I games, the presence of the dominated strategy, 3, alters the equilibrium outcomes. This is indicated by the observed variation in the equilibrium outcomes when x, y are altered. (3) Why is strategy 3 important? One clue is the fact that the cooperative outcome arises when both players play strategy 3, despite strategy 3 being a dominated one (because, $x, y > 600$). Type-II games control for this confound by ensuring that strategy 3 no longer leads to the cooperative outcome; in these games, (2, 2), is the cooperative outcome. Hence, the only reason strategy 3 should have any influence in Type-II games is if players have beliefs that opponents will play a dominated strategy. This reduces the actual play of strategy 3 in experiments and also increases the incidence of the payoff dominant equilibrium (2, 2).

12.4.2 *Focal points and coordination*

In the presence of multiple equilibria, economic agents must form beliefs about which equilibrium is likely to be played. Coordinating on certain actions may give everyone a higher payoff, which makes them salient. It is also possible that some actions are salient in other respects, also termed as *focal points*; salient actions may serve as a coordinating device for players. Following the work of Schelling (1960), Mehta et al. (1994, p. 659) write: “Some labels are more prominent or conspicuous or salient than others; they ‘stick out’ or ‘suggest themselves,’ typically by virtue of analogies or associations of ideas which connect those labels to some aspect of the common experience, culture, or psychology of the players. Players tend to choose those strategies whose labels are salient. An equilibrium that results from such choices is a focal point.”

Mehta et al. (1994) distinguish between the concepts of *primary salience* (actions that are salient for a player) and *secondary salience* (coordinating on actions that are believed to have primary salience for the other players). In order to test for the significance of these concepts, they divided subjects in experiments into two groups: Group C (coordinating condition) and group P (picking condition). Group C’s payoffs depended on coordinating their actions with the actions of others,

Table 12.10 Checking for primary and secondary salience.

-
1. Write down any year, past, present, or future: _____
 2. Name any flower: _____
 3. Name any car manufacturer: _____
 4. Write down any day of the year: _____
 5. Name any British town or city: _____
 6. Write down any positive number: _____
 7. Write down any colour: _____
 8. Write down any boy's name: _____
 9. Complete the sentence: A coin was tossed. It came down: _____
 10. Complete the sentence:
The doctor asked for the patient's record. The nurse gave them to: _____
-

Source: Mehta et al. (1994).

Table 12.11 Responses compiled from Table 2 of Mehta et al. (1994). The numbers in brackets are the percentages of the players making that choice.

No.	Group P			Group C		
	Modal response	Next best response	c	Modal best response	Next best response	c
1	1971 (8.0)	1990 (6.8)	0.026	1990 (61.1)	2000 (11.1)	0.383
2	Rose (35.2)	Daffodil (13.6)	0.184	Rose (66.7)	Daisy (13.3)	0.447
3	Ford (53.4)	Rover (5.7)	0.290	Ford (81.1)	N.A.	0.659
4	25/12 (5.7)	10/12 (1.1)	0.005	25/12 (44.4)	10/12 (18.9)	0.238
5	London (15.9)	Norwich (12.9)	0.054	London (55.6)	Norwich (34.4)	0.238
6	7 (11.4)	2 (10.2)	0.052	1 (40.0)	7 (14.4)	0.206
7	Blue (38.6)	Red (33.0)	0.269	Red (58.9)	Blue (27.8)	0.422
8	John (9.1)	Fred (6.8)	0.002	John (50.0)	Peter (8.9)	0.264
9	Heads (76.1)	Tails (20.5)	0.618	Heads (86.7)	Tails (12.2)	0.764
10	Him (53.4)	Her (40.9)	0.447	Him (84.4)	Her (15.6)	0.734

while the actions of group P were independent of any coordination. Subjects were asked 20 questions. The first 10 questions are shown in Table 12.10. The remaining 10 questions are abstract and involve coordinating on patterns.

A total of 178 undergraduate students participated in the experiment. Group C was relatively more successful at coordinating actions as compared to group P. A *coordination index*, c , is used to measure the probability that individuals who are chosen at random will coordinate on their actions. If individuals make completely distinct choices, $c = 0$, and if they make identical choices, $c = 1$; values of c in between these two extremes denote partial coordination. The two most frequent responses to each of questions 1–10, for the two groups P and C, are shown in Table 12.11, along with the coordination index, c . In order to interpret the results, note that the date of the experiment is December 10, 1990, most participants would have been born in the years 1968–1971, and the subjects in experiments came from the University of East Anglia in Norwich, England.

The coordination index for group C was higher relative to group P for each of the 10 questions. The modal response was the same in 16 out of the 20 questions for both groups. This seems consistent with secondary salience. A large number of individuals choose their birth year or the current year for question 1. The most frequent response for towns was London (the capital), followed by Norwich (where the students study). The color blue is the most common in group P but red is most popular for group C; perhaps people think that red is thought of by others as a more salient representation of colorfulness. Hence, labels do convey important information that can help coordinate actions. However, despite the high degree of salience in the questions, the coordination index is, on average, less than 60% for the first 10 questions for group C. This is suggestive of a degree of coordination, but also of coordination failures.

Bosch-Domènech and Vriend (2013) consider the role of non-Nash strategy profiles in serving as focal points in the presence of Nash equilibria. They consider a simultaneous move two-player normal form coordination game in which the row player chooses a row $(r_1, r_2, \dots, r_n)$ and the column player chooses a column $(c_1, c_2, \dots, c_n)$. A typical game, the 100 treatment, is shown in Figure 12.16; each cell in the payoff matrix shows an identical payoff (0 or 100) for each of the two players. Corresponding to any strategy played by the row player, there are two strategies of the column player that result in a coordinated payoff of 100 for both players. Lack of coordination results in a payoff of zero for both.

There are several Nash equilibria in the game; each results in a payoff of 100 to each player. Hence, one cannot use *payoff dominance* as a selection criterion.⁴⁵ Neither does *risk dominance* help to choose among any of these Nash equilibria.⁴⁶ Players are randomly matched at the end of each round of play, hence, other relevant criteria such as *precedent* cannot help either. No preplay communication is allowed, neither is any feedback given to the players. Finally, strategies carry no *label salience* and the game does not suggest any particular *frame*, *analogy*, or *context*. This rules out the most obvious ways of using classical focal point criteria.

The authors then take one of the Nash equilibria in Figure 12.16, corresponding to a strategy profile (r_j, c_k) and replace the payoff of 100 for each player by some number $x < 100$, leaving the rest of the payoffs unchanged. This gives rise to three “ x treatment” stage games, corresponding to $x = 46$, $x = 87$, and $x = 99$; these are denoted, respectively, by 46 treatment, 87 treatment, and 99 treatment. The original game may, thus, be termed as the 100 treatment. Notice that in any x treatment stage game, the strategy profile (r_j, c_k) no longer constitutes a Nash equilibrium because any one of the players can defect to another strategy that guarantees a payoff of 100. The authors ask, if in any x treatment stage game, the non-Nash strategy profile (r_j, c_k) can serve as a focal point. If the answer turns out to be in the affirmative, then this would demonstrate the power of non-Nash strategy profiles to serve as focal points.

In each x treatment, the subjects in experiments first play a one-shot x treatment stage game with relatively higher stakes and unlimited time. While no feedback is provided, the one-shot game allows players to use deductive reasoning to coordinate their actions. This is followed by 50 rounds of the play of the x treatment stage game with random opponents (the 50-round game). At the end of each round, players are informed of their payoffs but not the strategy of the other player. Each of the rounds is played for real money. The numbers shown in Figure 12.16 are experimental payoff points; 10 payoff points correspond to a payment of 1 euro in the one-shot game and to

⁴⁵ The empirical evidence in Colman et al. (2008) supports the choice of payoff dominant profiles over Pareto-inferior unique Nash equilibrium profiles.

⁴⁶ See the appendix on game theory for the concepts of payoff and risk dominance. Also see below, several coordination games where these selection criteria apply, hence, lending themselves to empirical testing.

	0	0	0	0	0	0	0	0	100	0	0	100	0	0	0
	0	0	0	0	0	100	0	0	0	0	0	0	0	0	100
	0	0	0	100	0	0	0	100	0	0	0	0	0	0	0
	0	0	0	0	100	0	100	0	0	0	0	0	0	0	0
	0	0	0	0	0	0	0	100	0	0	0	0	100	0	0
	0	100	0	0	0	0	0	0	0	0	0	0	0	100	0
	0	0	100	0	0	0	0	0	0	0	0	0	100	0	0
	100	0	0	0	0	0	0	0	0	0	0	0	0	0	100
	0	100	0	0	0	0	0	0	0	0	100	0	0	0	0
	0	0	0	0	0	0	0	0	0	0	0	100	0	100	0
	0	0	100	0	0	0	0	0	0	100	0	0	0	0	0
	0	0	0	0	100	0	0	0	100	0	0	0	0	0	0
	100	0	0	0	0	0	0	0	0	0	100	0	0	0	0
	0	0	0	0	0	100	0	0	0	100	0	0	0	0	0
	0	0	0	100	0	0	100	0	0	0	0	0	0	0	0

Figure 12.16 Coordination game (100 treatment).

Source: Reprinted from *Journal of Economic Behavior and Organization*, 94: 52–67, Bosch-Domènech, A. and Vriend, N. J., "On the role of non-equilibrium focal points as coordination devices." Copyright © 2013, with permission from Elsevier.

0.05 euros in any round of the 50-round games. Thus, a greater degree of coordination among the players at the end of 50 rounds will also be reflected in higher earnings.

The average earnings in euros at the end of 50 rounds for the 46 treatment, 87 treatment, 99 treatment, and the 100 treatment are, respectively, 8.39, 12.27, 14.82, and 5.32. Clearly, the degree of coordination in the x treatments is much higher as compared to the 100 treatment (that has no focal points). The degree of coordination in the 100 treatment was no higher than what one would predict based on random choice of strategies for the row and the column players. Hence, it did not arise from the players using any particular strategy as a focal point; this result holds for the one-shot game and for the repeated play of the game over 50 rounds. Further, the degree of coordination increases as x increases successively from 46 to 87 to 99.

In the one-shot game, there is marginally higher coordination in each of the x treatments relative to the 100 treatment but this is not statistically significant. However, in the 50-round games, this difference attains statistical significance. Thus, players learn to coordinate better with experience. Consider the results over rounds 41–50. Relative to the 100 treatment, expected payoffs are 1.8 times higher in the 46 treatment, 3.1 times higher in the 87 treatment, and 3.5 times higher in the 99 treatment. This increased coordination, reflected in higher expected payoffs, arises from subjects coordinating on the non-Nash focal point. The results over rounds 41–50 show that the non-Nash focal point was used by 45% of the subjects in the 46 treatment, 71% of the subjects in the 87 treatment, and 79% of the subjects in the 99 treatment. Thus, the focal point is used more often as the focal payoff increases.

The results of Bardsley et al. (2010) suggest that it would be very difficult to construct a unified theory of focal points. In two different coordination games they find that coordination is best

described by the *cognitive hierarchy model* in one and by *team reasoning* in the other.⁴⁷ A more formal treatment of focal points in the context of level- k models is given in the next chapter.

12.4.3 Coordination failures in median action games

Van Huyck et al. (1991) document coordination failures in *median action games*; these are static games of full information where players simultaneously choose their actions and have common knowledge of the rationality of the other players. Median action games capture the desire of players to conform to an action or a norm in the sense that deviations of their own actions from the median action are payoff-reducing. These experiments also test for the importance of equilibrium selection principles such as payoff dominance, risk dominance, and the role of history in influencing coordination.

Consider a large number of players, n where n is an odd number. Players, $i = 1, 2, \dots, n$, simultaneously choose actions, $s_i \in \{1, 2, \dots, \bar{s}\}$, where $\bar{s}$ is the maximum value of the positive integer that can be chosen. Denote the median action by M . We focus on pure strategies only. The payoff of player i is given by

$$\pi(s_i, M) = aM - b(M - s_i)^2 + c; \quad a > 0, b > 0. \quad (12.2)$$

Since players choose their actions simultaneously, at the time of choosing their actions, they do not know the median action. Player i expects the median action to be $E_i[M]$. The optimum action for each player that maximizes (12.2) is $s_i = E_i[M]$, which reflects the desire of players to conform to the median action. The calculation of $E_i[M]$ involves an infinite regress. Each player, $i = 1, 2, \dots, n$, must form beliefs about $E_i[M]$, and beliefs about the beliefs of the expected median action, and so on.

Suppose that players could somehow coordinate their actions and beliefs. Then, we are looking for a symmetric pure strategy Nash equilibrium that maximizes the payoffs of all players. The optimum action for each player is to choose $s_i^* = \bar{s}$, so that the median is $M = \bar{s}$. That this is a (strict) Nash equilibrium is easily confirmed from the calculation,

$$\pi(s_i, \bar{s}) < \pi(\bar{s}, \bar{s}) \quad \text{for all } s_i \neq \bar{s} \text{ and for all } i = 1, 2, \dots, n. \quad (12.3)$$

To test for various selection principles, Van Huyck et al. (1991) use three different parameterizations of their problem. In each case, the maximum value taken by any player's action is $\bar{s} = 7$.

1. In game-I, $a = 0.1$, $b = 0.05$, and $c = 0.6$; the associated payoff matrix is shown in Table 12.12.
2. In game-II, when $s_i = M$, payoffs are as in game-I (with $a = 0.1$, $b = 0.05$, and $c = 0.6$). But when $s_i \neq M$, the payoff is zero. Hence, the payoff matrix for game-II is as in game-I but all off-diagonal elements are zero.
3. In game-III, $a = 0$, $b = 0.05$, and $c = 0.7$; the associated payoff matrix is shown in Table 12.13.

Consider now the predictions for each game based on various selection criteria when we restrict attention to a pure strategy symmetric Nash equilibrium. In game-I, there is a conflict

⁴⁷ See Part 5 for an exposition of the cognitive hierarchy model. *Team reasoning* implies that when individuals take actions in a group they might take actions that are optimal from the point of view of an imaginary manager/social planner who maximizes some team objective. This may lead to a coordinated choice of actions; see Sugden (1993, 1995) for an extensive discussion of team reasoning and Colman et al. (2008) for the experimental evidence.

Table 12.12 Payoff matrix for game-I.

		Median action M						
		7	6	5	4	3	2	1
Your choice of action s_i	7	1.30	1.15	0.90	0.55	0.10	−0.45	−1.10
	6	1.25	1.20	1.05	0.80	0.45	0.00	−0.55
	5	1.10	1.15	1.10	0.95	0.70	0.35	−0.10
	4	0.85	1.00	1.05	1.00	0.85	0.60	0.25
	3	0.50	0.75	0.90	0.95	0.90	0.75	0.50
	2	0.05	0.40	0.65	0.80	0.85	0.80	0.65
	1	−0.50	−0.05	0.30	0.55	0.70	0.75	0.70

Source: Van Huyck et al. (1991).

Table 12.13 Payoff matrix for game-III.

		Median action M						
		7	6	5	4	3	2	1
Your choice of action s_i	7	0.70	0.65	0.50	0.25	−0.10	−0.55	−1.10
	6	0.65	0.70	0.65	0.50	0.25	−0.10	−0.55
	5	0.50	0.65	0.70	0.65	0.50	0.25	−0.10
	4	0.25	0.50	0.65	0.70	0.65	0.50	0.25
	3	−0.10	0.25	0.50	0.65	0.70	0.65	0.50
	2	−0.55	−0.10	0.25	0.50	0.65	0.70	0.65
	1	−1.10	−0.55	−0.10	0.25	0.50	0.65	0.70

Source: Van Huyck et al. (1991).

between payoff dominance and security. The payoff dominant action for each player is $s_i = \bar{s} = 7$. However, the maximin or secure action for a player is $s_i = 3$.⁴⁸ In game-II, all off-diagonal elements are zero, so there is no secure option because the minimum payoff corresponding to each action is zero. However, there is a payoff dominant choice, $s_i = \bar{s} = 7$. Thus, unlike game-I there is no conflict between security and payoff dominance, so the prediction, based on the selection principles, is that $s_i = \bar{s} = 7$ should be observed. In game-III, all payoffs along the diagonal are identically equal to 0.70. Thus, payoff dominance has no bite. Security (the maximin strategy) selects $s_i = 4$ in this case.

Groups of nine players played the games over several rounds, so that players could learn to coordinate. The experimenter revealed the median to all players at the end of each round, but not the individual actions. The results are as follows.

1. The first period results are particularly interesting because players' actions cannot depend on any history of play. Rather, in choosing their actions, players can only resort to *deductive*

⁴⁸ This is easily seen from Table 12.12. The minimum payoffs corresponding to each of the actions 1 to 7 are respectively −0.5, 0.05, 0.50, 0.25, −0.10, −0.55, and −1.10. The maximum of these payoffs is 0.50, which corresponds to $a_i = 3$.

reasoning, about what others are likely to play. In game-I, where there is a conflict between payoff dominance and security, 15% choose $s_i = 7$ (the payoff dominant action) and another 15% choose $s_i = 3$ (the secure option). In game-II, the prediction $s_i = \bar{s} = 7$ is satisfied by only 52% of the players. In game-III, the prediction $s_i = 4$ is satisfied by only 41% of the players. Hence, none of the selection principles individually apply very well and there are widespread coordination failures.

2. In the first period, the individual choices of actions are quite dispersed, suggesting that players' beliefs about others are dispersed too. As Van Huyck et al. (1991, p. 897) observe: "The data are inconsistent with any theory of equilibrium selection that assumes that, because a decision maker will derive his prior probability distribution over other decision makers' pure strategies strictly from the parameters of the game, all decision makers will have the same prior probability distribution. Instead, most subjects predict that some participants will be more optimistic or pessimistic than themselves."
3. The actions and beliefs of the players in the first period seem consistent. More optimistic players choose higher actions and more pessimistic players choose lower actions.
4. In repeated play of the games, there seems to be a strong effect of history on subsequent play. The period 1 median was nearly always observed again in period 10. The other results from the first period play also replicate themselves in repeated play, namely, no single selection principle performs well, and there is substantial dispersion in the choices of individual players.

Van Huyck et al. (1993) extend the experiments in Van Huyck et al. (1991) with an important twist. They auction the rights to play the median action game among players. If players have diverse beliefs, then an auction may select the players whose beliefs are most optimistic about the coordination behavior of others. Each player may then coordinate more, reasoning that other players who have incurred resources in the auction will attempt to recoup their resources by facilitating higher coordination in the median action game. This is a form of forward induction. As expected, the degree of coordination is much higher. Further, the speed of convergence to a coordinated outcome is impressive. Crawford and Broseta (1998) build on the earlier work in Crawford (1995) to provide a unified theoretical explanation of Van Huyck et al. (1991, 1993) based on an adaptive learning model. In this model, players have heterogeneous beliefs and they engage in unsophisticated learning based on their observations of the actions of other players. Forward induction effects play a key role in the analysis.

12.4.4 *Coordination failures in weak link games*

Consider the following static game of full information in Van Huyck et al. (1990). Suppose that there are n players. Each player simultaneously chooses some action $s_i \in \{1, 2, \dots, 7\}$, $i = 1, 2, \dots, n$. Denote the minimum action by $\theta = \min \{s_1, s_2, \dots, s_n\}$. Then, the payoff of each player in Treatments A, B is given by the following two possible payoff functions.

$$A : \pi(s_i, \theta) = 0.60 + (0.10)\theta - 0.10(s_i - \theta). \quad (12.4)$$

$$B : \pi(s_i, \theta) = 0.60 + (0.20)\theta. \quad (12.5)$$

In Treatment A, there is a private cost to one's action, while there is none in Treatment B. In both treatments, the payoff is increasing in the minimum action taken by the group. A motivating example is a sequential assembly line production process in a firm, whereby work on one

production process can only begin once a previous production process is completed. Hence, final output moves at the rate of the slowest of the production processes. Or consider a team of software developers. A bug in the program introduced by any member may lay to waste the efforts of all team members.

In Treatments A and B, the game was played with $n = 14$ –16 subjects, repeated over 10 rounds. If players could coordinate their actions and beliefs to maximize their payoffs, then in Treatments A and B, they should coordinate on the payoff dominant action $s_i = 7$ for all i . In addition, there are two further treatments. In Treatment C, the payoffs are as in (12.4) but $n = 2$. In Treatment A', three treatments are played in sequence: $A \rightarrow B \rightarrow A$.

Consider Treatment A first. Since players choose their actions simultaneously, they must form beliefs about the actions of others. If player i expects the minimum action to be $E_i[\theta]$, then, using (12.4), the optimal choice is $s_i = E_i[\theta]$. The game is symmetric with respect to the players, and players are identical, hence, it is plausible that they have identical expectations, so $E_i[\theta] = E[\theta]$ for all i . This implies that, in a symmetric Nash equilibrium, the optimal action is $s_i^* = E[\theta] = a \in \{1, 2, \dots, 7\}$ for all i . Thus, the equilibrium strategy profile is given by $s^* = (a, a, \dots, a)$, so $\theta^* = a$, and the following condition for a Nash equilibrium is satisfied,

$$\pi(s_i, a) < \pi(a, a) \text{ for all } s_i \neq a. \quad (12.6)$$

Using (12.4) and (12.6) we get $a = 7$.

In Treatment B, on the other hand, players are not penalized for the deviation of their actions from the median action. So choosing the payoff dominant action $s_i = 7$ is a dominant strategy. Hence, this is a cognitively less challenging game.

The payoff matrix for Treatment A is given in Table 12.14; for any choice of action s_i by player i , the group minimum cannot be higher, hence, the entries below the diagonal are truncated. From Table 12.14, the secure option is for each player to play $s_i = 1$, which guarantees a sure payoff of 0.70 that is greater than the minimum payoff from any other action (maximin outcome). Thus, there is a conflict between the payoff dominant option of $s_i = 7$ and the secure option of $s_i = 1$.

In repeated versions of the game, there was no preplay communication, and at the end of each repetition of the game, the minimum action in that period was announced to all players but individual actions were not revealed.

Table 12.14 Payoff matrix for treatment A.

		Minimum chosen action, θ						
		7	6	5	4	3	2	1
Your choice of action s_i	7	1.30	1.10	0.90	0.70	0.50	0.30	0.10
	6	–	1.20	1.00	0.80	0.60	0.40	0.20
	5	–	–	1.10	0.90	0.70	0.50	0.30
	4	–	–	–	1.00	0.80	0.60	0.40
	3	–	–	–	–	0.90	0.70	0.50
	2	–	–	–	–	–	0.80	0.60
	1	–	–	–	–	–	–	0.70

Source: Van Huyck et al. (1990).

The first period results are of interest because they test for deductive reasoning. In Treatment A, where there is a conflict between the payoff dominant and the secure option, the payoff dominant action was chosen by 31% of the subjects and the secure action by just 2%. Thus, neither of the two deductive principles is very helpful to explain the data. In all seven experiments, the minimum action in period 1 was never greater than 4. Hence, coordination failures are widespread (because coordination on $s_i = 7$ for all i , maximizes payoffs).

As in Van Huyck et al. (1991), the choices of the subjects in experiments were quite dispersed, suggesting lack of common priors. There was a range of optimistic and pessimistic beliefs but actions were consistent with beliefs. Optimistic players (in the sense of having optimistic beliefs) choose high actions and pessimistic players choose lower actions. Repeated play brings inductive reasoning into play but does not achieve any stable outcome in Treatment A. At the end of ten rounds, play converges towards the most inefficient outcome, $s_i = 1$, in all the seven experiments.

In Treatment B, which is cognitively less challenging, and where the prediction is $s_i = 7$, by period 11, 84% of the players conform to the prediction. By period 15, 96% of the players conform to $s_i = 7$. The payoff dominant equilibrium was still not implemented in most cases because at least one player would choose $s_i < 7$.

In Treatment A', which follows Treatment B, by round 16, 25% of the players choose the payoff dominant outcome $s_i = 7$. However, an even larger percentage, 37%, choose the secure option $s_i = 1$.

Are group sizes important for coordination games? This issue is addressed in Treatment C. When subjects are paired with a permanent partner, by round 21, 42% of the players choose their payoff dominant action. However, when subjects were randomly paired then the positive effect on s_i that arises from permanent pairings was reduced. Hence, it is not just group size that is important, but also how groups are organized. Reducing the group size alone does not automatically improve coordination.

12.4.5 *More evidence on equilibrium selection principles*

Straub (1995) considers eight games, which have *risk dominant* and *payoff dominant* equilibria. The main interest is in observing which, if any, of these equilibria does play converge to over ten rounds of a game in which players are matched anonymously with each other. Whenever the Nash product from both strategies is not identical, then the risk dominant equilibrium is the one that is observed. This contrasts with the expectation of Harsanyi and Selten (1988) that the payoff dominant equilibrium is more likely to be observed. When the Nash products are equal, then the risk dominant equilibrium is associated with the play of a mixed strategy. The frequency of play by the players in this case seems consistent with the risk dominant mixed strategy. However, it seems unlikely that the players used a purely deductive process to reason to a mixed strategy risk dominant equilibrium. And Straub (1995, p. 351) writes: "Instead, it appears much more likely that if the players achieved an equilibrium during a session, they were aided by a (possibly evolutionary) process in which their strategies converged to an equilibrium."

12.4.6 *Coordination and the optimization premium in stag-hunt games*

Battalio et al. (2001) considered three *stag-hunt games* that have identical equilibrium outcomes in pure and mixed strategies. Hence, one may expect subjects in experiments to play the same

Game 2R				Game R				Game 0.6R			
		X	Y			X	Y			X	Y
X	45, 45	0, 35		X	45, 45	0, 40		X	45, 45	0, 42	
Y	35, 0	40, 40		Y	40, 0	20, 20		Y	42, 0	12, 12	

Figure 12.17 Game 2R, game R, and game 0.6R.

Source: Battalio et al. (2001) with permission of The Econometric Society.

equilibria across the three games.⁴⁹ The three games, denoted by 2R, R, and 0.6R, are shown in Figure 12.17.

In all three games, (X, X) is the *payoff dominant equilibrium*, and (Y, Y) is the *risk dominant equilibrium*. In the mixed strategy equilibrium, both players play action X with probability $q^* = 0.8$. Strategy X is a best response for any player when the opponent plays X with a probability $q > 0.8$.

The one main difference between the three games is the *optimization premium*, defined as the difference between the best response to an opponent's strategy and an inferior response. Consider, for instance, the equilibrium (X, X). For game R, the optimization premium for either the row or the column player is $45 - 40 = 5$. The optimization premium for game 2R is 10, double that of game R, and for game 0.6R it is 3, which is 0.6 times that of game R. This explains the logic behind the names of the three games.

For a player who is fully rational in the classical sense, the optimization premium should not matter. The best response functions still intersect at the same point. However, the optimization premium alters the slope of the best response. Hence, it might affect the behavior of players who are less than fully rational. Battalio et al. (2001) test this conjecture.

In the experiments, there are eight subjects in each cohort and each cohort played one of the three games, 75 times. Subjects were randomly paired at the end of each round to minimize reputation effects. No preplay communication was allowed and the opponent in each round was anonymous.

The results for period 1 (where choices are purely deductive) and for period 75 (where choices involve deductive and inductive reasoning) are shown in Table 12.15. In period 1, 63% of all subjects play X, the payoff dominant choice. Since $0.63 < q^* = 0.8$, it is not a dominant strategy to play X. Only 37% play the risk dominant strategy. Across the three games, the tendency to play the payoff dominant equilibrium is not statistically different.

Table 12.15 Experimental results.

Game	Period 1 choices (%)		Period 75 choices (%)		Number
	X	Y	X	Y	
0.6R	64	36	44	56	64
R	70	30	25	75	64
2R	53	47	5	95	64
Total	63	37	24	76	192

Source: Battalio et al. (2001).

⁴⁹ In stag-hunt games, unlike a prisoner's dilemma game, there are two Nash equilibria; one equilibrium is payoff dominant and the other is risk dominant.

All this changes by period 75. Now only 5% play the payoff dominant action, X, in game 2R, while 44% play X in game 0.6R. Thus, the payoff dominant action is more likely to be played when the optimization premium is smallest. Analogously, the risk dominant equilibrium is played more often (95%) when the optimization premium is highest, as in game 2R. Hence, which of the two coordinated equilibria gets played depends on the optimization premium, which is predicted to have no effect in classical game theory.

12.4.7 *Forward induction, timing, and coordination failures*

Cooper et al. (1993) consider two players who play the following sequential game that can be represented in an extensive form.

1. In the first stage, player 1 can choose an outside option which gives him a payoff, X, and end the game. Or, he can choose not to exercise the outside option, in which case the game proceeds to the second stage.
2. In the second stage, the two players play the following battle of the sexes (BoS) game (player 1 is the row player and player 2 is the column player):

	1	2
1	0, 0	200, 600
2	600, 200	0, 0

Payoffs are given as points. Dividing the points by 1000 gives the probability of winning a prize. What outcome do we expect? Suppose that the outside option, $X \in (200, 600)$ and this is known to both players. Then players can try to use *forward induction* in order to solve the game.⁵⁰ Player 2 can now reason as follows. If player 1 has forgone his outside option, then he must hope to get a higher payoff. In particular, when $X = 300$, we call the resulting game BoS-300. In this case, forward induction would suggest to player 2 that player 1 intends to play strategy 2. Coordinating expectations in this manner, player 2 would then play strategy 1. Thus, the prediction is that the players will coordinate on the equilibrium (2,1). We summarize the prediction for this game by hypothesis H1.

H1: In the game BoS-300, the equilibrium outcome is expected to be (2,1).

When $X = 100$, rejecting the outside option is not informative about player 1's intentions. Calling this game BoS-100, we get the next hypothesis.

H2: In the game BoS-100, the outcome is predicted to be identical to the game BoS in which the outside option does not exist.

Classical game theory predicts that the choices of players in extensive form games and their equivalent strategic form representations are identical. The evidence examined earlier in this chapter does not support this prediction. Cooper and Van Huyck (2003) reach similar conclusions for the class of generic 2×2 extensive form games of perfect information. They find systematic differences between the choices made in strategically equivalent normal and extensive forms that could not be ascribed to other factors, e.g., other-regarding preferences. Consider the extensive form described above for the case $X = 300$. Denote by 0, the strategy of player 1 to take

⁵⁰ For the concepts of forward induction and stability, the reader can consult any graduate book in game theory; see, for instance, Fudenberg and Tirole (1991) and Maschler et al. (2013).

his outside option, and let player 1 be the row player. Then, the strategically equivalent normal form corresponding to the extensive form game is:

	1	2
0	300, 0	300, 0
1	0, 0	200, 600
2	600, 200	0, 0

Denote this normal form with an outside option of $X = 300$ by BoS-300-NF.

H3: The outcomes under BoS-300 and BoS-300-NF should be identical.

Recall that in the BoS game, both players simultaneously choose their strategies. Now consider a new game, BoS-SEQ. The payoff matrix is as in the BoS game, but player 1 chooses his strategy prior to player 2's strategy choice. However, player 2 is not informed about the strategy chosen by player 1. Under classical game theory, such *timing issues* should not matter.

H4: The outcomes in the games BoS and BoS-SEQ should be identical.

The results are shown in Table 12.16. Results seem consistent with hypothesis H1, i.e., forward induction has sufficient bite; 90% of the equilibrium outcomes in BoS-300 are (2,1), as predicted. However, all the other hypotheses fail.

In the baseline game, BoS, nearly 60% of the outcomes were disequilibrium outcomes characterized by coordination failures. This result was subsequently replicated by Straub (1995) who found that for the BOS game, the non-coordinated outcome accounted for 60% of the cases, coordinated outcomes at (1,2) for 26% of the cases, and coordinated outcomes at (2,1) for 14% of the cases.

Despite the success of hypothesis H1, 20% of the subjects took the outside option in BoS-300. The results in the BoS and BoS-300 games are quite different. In BoS, nearly an equal number of players coordinated on the two equilibria. However, in BoS-300, 90% coordinated on (2,1), which is player 1's preferred equilibrium. The results in BoS-300 and BoS-300-NF are quite different, a finding that rejects hypothesis H3. In the latter game, in more than half the cases, subjects were not able to coordinate on any of the two equilibria, while in BoS-300 they did so in 90% of the cases.

Hypothesis H2 can also be rejected. The presence of the irrelevant outside option $X = 100$ increases the incidence of the (2,1) equilibrium from 19% in BoS to 63% in BoS-100. The overall

Table 12.16 Experimental results.

Game	Outcome (in percentage)			
	(2, 1)	(1, 2)	(1, 1) or (2, 2)	Outside option
BoS	19	22	59	–
BoS-100	63	3	34	1.8
BoS-300	90	0	10	20
BoS-SEQ	62	4	34	–
BoS-300-NF	32	1	52	15

Source: Cooper et al. (1993).

coordination rate jumps up from 41% to 66%. It would seem that the mere presence of the outside option, even when predicted not to matter, acts as a focal point for expectations about how a player will play.

Hypothesis H4 is also rejected. Timing matters, even when it is predicted not to. The overall coordination rate in BoS-SEQ is 66%, as compared to 41% in BoS.

In a nutshell, different outside options, timing issues, and framing effects (whether the game is presented in normal or in extensive form) have an important bearing on the level of coordination.

12.4.8 *Preplay communication and coordination failures*

Suppose that prior to a two-player BoS or a coordination game being played, players can send each other costless messages (*cheap talk*) about their intentions to play particular strategies. If only one of the players is allowed to send a message, we call it one-way communication (1W) and if both players are allowed to send messages, then we call it two-way communication (2W). In classical game theory, cheap talk is believed by one's opponents if it is in the interests of a player to honor their promises. The role of 1W and 2W is explored in Cooper et al. (1989, 1993, 1994). We report below the results of these experiments, summarized in Cooper et al. (1994). Consider the following battle of sexes game.

	1	2
1	0, 0	200, 600
2	600, 200	0, 0

When this game is played, respectively, under one-way and two-way communication, then denote this game as BoS-1W and BoS-2W.

We already know from Section 12.4.7 that in BoS, the coordinated outcomes, either (2, 1) or (1, 2), were achieved in 41% of cases. In BoS-1W, the coordination rate is 96%, while in BoS-2W, it is 58%. Furthermore, when the row player sends out a cheap talk message, in BoS-1W, 95% of the coordination is on (2, 1), which is player 1's preferred equilibrium. In BoS-2W, both coordination equilibria are played with similar frequencies.

These results suggest that 1W works brilliantly to create almost perfect coordination but 2W is relatively less successful. The reason is that, in 2W, players announced their most preferred strategies (both players announced strategy 2). This was less successful in fostering cooperation. However, in 1W, the row player, by announcing his most preferred strategy, most likely induced an *anchoring of beliefs*, so the equilibrium (2, 1) tended to be played 95% of the time.

Cooper et al. (1992) consider a simple coordination game (SCG); the strategic form is shown in Figure 12.18. The two pure strategy Nash equilibria are (1, 1) and (2, 2). The strategy profile (1, 1) is also the risk dominant equilibrium.

The results for SCG are as follows. In the absence of any communication, players mainly play the risk dominant equilibrium (1, 1). In the presence of one-way communication (1W), the

	1	2
1	800, 800	800, 0
2	0, 800	1000, 1000

Figure 12.18 A simple coordination game (SCG).

Source: Cooper, R., DeJong, D. V., Forsythe, R., and Ross, T. W., "Communication in coordination games," *Quarterly Journal of Economic* (1992) 107(2): 739–71, by permission of Oxford University Press.

equilibrium (2, 2) arises more frequently (53% of the time). However, two-way communication (2W) works the best. The Pareto dominant equilibrium (2, 2) arises 97% of the time, presumably because two-way communication serves to mutually increase the confidence of players about the intent of the other player, overriding the risk dominant equilibrium (1, 1).

12.4.9 Historical accidents and choice among coordination equilibria

In the presence of multiple equilibria, the chosen equilibrium may depend on the history of the game. Van Huyck et al. (1997) explore this idea in the context of the *continental divide game*, a static game of full information in which players simultaneously choose their actions.⁵¹

Suppose that there are n players and each chooses some action s_i , $i = 1, 2, \dots, n$. Let the profile of actions chosen by the n players be $s = (s_1, \dots, s_n)$. Denote the median action by M . The payoff of each player $i = 1, 2, \dots, n$ is given by

$$\pi(s_i, M) = aM - b[s_i - (M)^{1/3}]^2 + c; \quad a > 0, b > 0. \quad (12.7)$$

Consider the special case where $a = 26$, $b = 40$, $c = 86$, and players $i = 1, 2, \dots, n$ choose an action s_i from the finite set $\{1, 2, \dots, 14\}$. The resulting payoff matrix is shown in Table 12.17; the payoff of a player in any cell corresponds to a particular action chosen by the player and the median action chosen.

Focus on the two shaded parts of the payoff matrix. When the median, M , lies between 1 and 7, choosing an action s_i between 1 and 7 (rather than 8 to 14) gives any player a higher payoff. Conversely, when the median lies between 8 and 14, any player receives a higher payoff from choosing an action between 8 and 14 (rather than 1 and 7).

Players choose their actions simultaneously, so they must form beliefs about the actions of others in order to guess the likely median. Since players are ex-ante identical and have identical strategy sets and payoff functions, we look for a symmetric Nash equilibrium in which $s_i^* = d \in \{s_1, \dots, s_n\}$ for $i = 1, 2, \dots, n$, so $M = d$. Thus, we are interested in finding a value of d such that $\pi(d, d) \geq \pi(s_i, d)$ for all i and s_i , where π is defined in (12.7). The solution can simply be read off from Table 12.17. There are two possible solutions, shown by the action–median pairs (3, 3) and (12, 12) with respective payoffs equal to 60 and 112 (check that each solution is a symmetric Nash equilibrium).

Suppose that the continental divide game is played several times. Let the initial median, M , be between 1 and 7. In this case, it is a best response for all players to play an action between 1 and 7. Hence, over time, play is expected to converge to the action–median pair (3, 3). Analogously if the first round median, M , is between 8 and 14, play in subsequent rounds is expected to converge to the action–median pair (12, 12). Figure 12.19 bears out this intuition. All play that begins above the dotted line (*separatix*), in Figure 12.19 converges to a median of 12 over the 15 periods of the experiments conducted in groups of 7 each. Conversely, all play that begins below the *separatix* converges to a median of 3 over the 15 periods.


These results make a powerful case for the role of history in choosing among multiple coordination equilibria. Equilibrium analysis in traditional game theory assumes that play instantly converges to one of the two equilibria. Van Huyck et al. (1997) show that their results are

⁵¹ The term *continental divide* refers to a “divide” (typically a raised partition along the terrain) such that if a drop of water falls on either side of the divide it travels in opposite directions, eventually reaching different bodies of water such as rivers and oceans.

Table 12.17 Payoff matrix for the continental divide game.

		Median choice, M													
		1	2	3	4	5	6	7	8	9	10	11	12	13	14
Your choice, s_i	1	45	49	52	55	56	55	46	−59	−88	−105	−117	−127	−135	−142
	2	48	53	58	62	65	66	61	−27	−52	−67	−77	−86	−92	−98
	3	48	54	60	66	70	74	72	1	−20	−32	−41	−48	−53	−58
	4	43	51	58	65	71	77	80	26	8	−2	−9	−14	−19	−22
	5	35	44	52	60	69	77	83	46	32	25	19	15	12	10
	6	23	33	42	52	62	72	82	62	53	47	43	41	39	38
	7	7	18	28	40	51	64	78	75	69	66	64	63	62	62
	8	−13	−1	11	23	37	51	69	83	81	80	80	80	81	82
	9	−37	−24	−11	3	18	35	57	88	89	91	92	94	96	98
	10	−65	−51	−37	−21	−4	15	40	89	94	98	101	104	107	110
	11	−97	−82	−66	−49	−31	−9	20	85	94	100	105	110	114	119
	12	−133	−117	−100	−82	−61	−37	−5	78	91	99	106	112	118	123
	13	173	−156	−137	−118	−96	−69	−33	67	83	94	103	110	117	123
	14	217	−198	−179	−158	−134	−105	−65	52	72	85	95	104	112	120

Source: Van Huyck et al. (1997).

**Figure 12.19** Median choice in sessions 1–10 over 15 periods in the continental divide game.Source: Reprinted from *Journal of Economic Behavior and Organisation*, 32: 488–503, Van Huyck, J., Cook, J., and Battalio, R., "Adaptive behavior and coordination failure." Copyright © 1997, with permission from Elsevier.

consistent with two kinds of behavioral models of learning. (i) Models of *myopic best response dynamics* in which players best respond to the observed median in the previous period. (ii) Models of *fictitious play* in which players best respond to the observed frequency of the historical medians. These models are considered in Part 5.

12.4.10 Financial incentives and coordination

Suppose that we have a history of coordination failures in an organization. Brandts and Cooper (2006) consider how financial incentives could be used to overcome such coordination failures. They use bonuses as a particular example of financial incentives in a *corporate turnaround game*. There are four employees and one manager in the game, representing a corporation that has a history of coordination failures. In the experiment, the experimenter plays the role of the manager and the subjects play the role of the employees.

The manager starts off by announcing a fixed wage, W , and a credible performance-contingent bonus, $B \in \{6, 8, 10, 14\}$. Given W and B , each employee simultaneously chooses effort $s_i \in \{0, 10, 20, 30, 40\}$. The production technology is described by a *weak link minimum effort game*, in which the minimum effort level of a team member is decisive for the team outcome. Denote the minimum effort level of the employees by $\theta = \min \{s_1, s_2, s_3, s_4\}$. The payoffs of the firm, π_F , and the workers, π_i , $i = 1, 2, 3, 4$, are respectively given by the following expressions.⁵²

$$\pi_F = 100 + \theta(60 - 4B). \quad (12.8)$$

$$\pi_i = 200 - 5s_i + \theta B. \quad (12.9)$$

An increase in the minimum effort level benefits workers and firms. Each worker derives disutility from effort. Suppose that initially there are coordination failures, for instance, all employees choose zero effort level, $s_i = 0$. Let $B = 6$ and suppose that employee i contemplates an increase in own effort level from $s_i = 0$ to $s_i = 10$. Denote by p the probability that the minimum effort level of the remaining three players is also at least 10. Then, using (12.9), employee i prefers to increase the effort level from $s_i = 0$ to $s_i = 10$ if $200 < 200 - 50 + p \times 60 + (1 - p) \times 0 \Leftrightarrow p \geq 0.83$. Thus, the barrier to eliciting coordination is very high and, in the absence of some sort of “big push,” coordination is unlikely. The barrier is somewhat reduced with a higher bonus rate. For instance, for $B = 14$ we require $p \geq 0.36$. So we expect effort to be increasing in the bonus rate.

In the experiments, subjects played 30 rounds with the same set of players. For the first 10 rounds, $B = 6$. The expectation, given the results from other weak link games, is that coordination failures will arise. The bonus rate is then varied over the next two blocks of 10 rounds each. From rounds 11–20, the bonus rate was increased, while from rounds 21–30 it was cut back. In each block of 10 rounds, subjects are told of the bonus rate that applies to that block of 10 rounds, but not to future blocks. As expected, the minimum effort is low for the first ten rounds; $s_i = 0$ for 71% of the observations and the average θ over the 10 rounds is 5.71. Figure 12.20 shows the minimum effort over rounds 11–20 for various bonus rates.

The findings for rounds 10–20 are as follows.

⁵² The payoffs are given in an experimental currency, pesetas, such that 500 pesetas equalled either 1 dollar or 1 euro.




Figure 12.20 Average minimum effort levels over rounds 11–20 for different treatments.

Source: Brandts and Cooper (2006) with permission from the American Economic Association.

1. An increase in B leads to an increase in the minimum effort level. This increase cannot be ascribed to simply starting the second block of 10 rounds because with an unchanged bonus $B = 6$, the minimum effort does not change. By rounds 16–20 the minimum effort level stabilizes with the new bonus rate but at a higher level.
2. There is no monotonic relation between minimum effort and the magnitude of the bonus rate. Over rounds 16–20, the bonus levels $B = 8$ and $B = 10$ elicit very similar minimum effort levels, while $B = 14$ elicits the lowest minimum effort level. Thus, what seems to be important is the presence of the bonus, rather than its rate.
3. Examining data at the level of individual employees enriches the picture. In the experiments, subjects received feedback on their effort levels, the minimum effort level, and the effort levels of others in the past. Can sterling efforts on the parts of some players rub-off on others in subsequent rounds? Initially, there is only a modest increase in effort after an increase in B . Gradually, however, the groups bifurcate into two types. Define an individual as a strong leader if he/she jumps up two levels of effort over rounds 10 to 11. There is strong evidence that groups that had a larger number of strong leaders had much higher minimum effort levels.

Finally, over rounds 21–30, the increase in bonus that took place over rounds 11–20 is withdrawn. Whilst this reduces effort levels, interestingly, on average, the effort level does not reduce to the level of rounds 1–10. Indeed, in some firms there is hardly any change in the effort level but in others there is a drop.

12.4.11 *Loss aversion and coordination*

Cachon and Camerer (1996) explore the role of loss aversion as a selection principle in coordination games. They use the median action games of Van Huyck et al. (1991, 1993) that we have

Table 12.18 A median action game (payoffs in cents).

		Median action M						
		1	2	3	4	5	6	7
Your choice of action s_i	1	140	150	140	110	60	−10	−100
	2	130	160	170	160	130	80	10
	3	100	150	180	190	180	150	100
	4	50	120	170	200	210	200	170
	5	−20	70	140	190	220	230	220
	6	−110	0	90	160	210	240	250
	7	−220	90	20	110	180	230	260

Source: Cachon and Camerer (1996).

already considered above. In particular, they consider the payoff matrix (with payoffs denominated in cents), shown in Table 12.18.

The payoffs corresponding to any combination of own action and the median action are given in the appropriate cells. Actions of players are allowed to be integer values between 1 and 7. Using arguments analogous to those in Section 12.4.3, all of the elements along the diagonal constitute distinct symmetric pure strategy Nash equilibria. The Pareto optimal outcome is for each player to choose the highest action, $s_i = 7$ for all i . Recall from Section 12.4.3 that we observed coordination failures (the first round median is 4–5) and there is history dependence (after repeated play, the actions converge to the first round median). Indeed, history dependence might itself be considered to be a selection principle.

Recall the intuition behind the effect of *outside options* based on forward induction in Section 12.4.7 above. Suppose that in the spirit of forward induction, subjects are charged a fee of 225 cents to play the game. Subtracting 225 from each entry in Table 12.18, the play of any of the actions 1 through 4 results in a sure loss, irrespective of the median. How should players be expected to respond? Cachon and Camerer (1996) propose *loss aversion* as a selection principle.⁵³ On p. 167 they describe it as follows: “The loss-avoidance principle is that players do not pick strategies that result in certain losses for themselves, if other ‘equilibrium’ strategies are available. People only pick, and expect others to pick strategies that might result in a gain.” The loss avoidance principle would then imply that players neither play strategies 1–4, nor do they expect others to play them. Thus, only strategies 5–7 should be observed in equilibrium.

Assume now that players must pay a fixed fee to play the game with payoffs prior to paying the fee given in Table 12.18. Consider two conditions.

1. Opt-out condition: In this condition, players can either pay the fee of 225 and play the game, or opt-out by not paying the fee, in which case they get nothing.
2. Must-play condition: In this condition, subjects cannot opt out and must play the game after paying the fee.

Forward induction has no bite in the must-play condition. Once an individual has paid a fee in this condition, it signals nothing about the individual’s subsequent intentions. However, forward

⁵³ See Part 1 for a detailed discussion of loss aversion. Essentially it implies that losses bite more than equivalent gains. In the experiments of Kahneman and Tversky (1979), on average, losses bite approximately 2.25 times more than equivalent gains. However, the estimates of loss aversion vary, depending on the context.

	R	S		R	S		R	S
R	7, 7	1, 5	R	5, 5	-1, 3	R	1, 1	-5, -1
S	5, 1	5, 5	S	3, -1	3, 3	S	-1, -5	-1, -1
High payoff game (H)			Medium payoff game (M)			Low payoff game (L)		

Figure 12.21 Three equivalent stag-hunt games.

Source: "Payoff levels, loss avoidance and equilibrium selection in games with multiple equilibria: an experimental study," Feltovich, Nick, Atsushi Iwasaki, and Sobei H. Oda, *Economic Inquiry* 50: 932–52. Copyright © 2011 Western Economic Association.

induction has bite in the opt-out condition. A player who pays a fee of 225 cents in the opt-out condition must then expect to receive a payoff greater than 225 cents, which coordinates actions to the subset of actions 5–7. Loss aversion, however, has bite in both conditions (assuming no changes in reference points). According to the loss-avoidance principle, subjects will not choose or expect others to choose actions that result in losses. The two main results are as follows.

1. In the opt-out condition, coordination improved once the opt-out was introduced. This could be due to both loss aversion and forward induction.
2. In the must-play condition, coordination also did improve. Since forward induction has no bite in this condition, the improvement must be due to loss aversion.

In conjunction, these results imply that some of the improved coordination that has been ascribed to forward induction is due to loss aversion. The authors conclude that loss aversion is an effective selection principle in this class of games.

Feltovich et al. (2012) consider three classically equivalent versions of the stag-hunt game, shown in Figure 12.21. The original game (H) is shown in the left panel. The medium payoff game (M) in the middle panel is derived by subtracting 2 from each payoff in the game H. The low payoff game (L) is derived by subtracting 6 from each payoff in the game H.

Each of the games has two pure strategy Nash equilibria (R, R) and (S, S). There is also a mixed strategy Nash equilibrium in which both players choose R with probability $2/3$. The equilibrium (S, S) is risk dominant and involves the play of the secure option, R. The equilibrium (R, R), on the other hand, is payoff dominant. Classic game theory does not predict any difference between the outcomes of these games.

Feltovich et al. (2012) distinguish between *certain-loss avoidance* and *possible-loss avoidance*. Possible-loss avoidance is invoked by the play of strategy R in game M, while certain-loss avoidance is invoked by the play of strategy S in the game L. Thus, certain-loss avoidance arises when one receives a negative payoff, irrespective of the strategy of the opponent, while possible-loss avoidance arises when not all payoffs corresponding to a strategy are negative. *Loss avoidance* may then be considered to be an overarching criterion that includes these two sub-cases.

If loss avoidance has bite, then we have the following two predictions. (1) Relative to game H, we should observe a lower incidence of the play of strategy R in game M (possible-loss avoidance). (2) Relative to game H, we should observe a lower incidence of the play of strategy S in game L (certain-loss avoidance).

In all treatments, each of the three games was played in a within-subjects design. Treatment O is the only treatment in which each game was played once. Subjects played 40 rounds of each of the three games in the remaining three treatments. In treatment C, there was complete information about the games. In two further, limited information treatments, F and R, the payoff matrix was not revealed to the subjects but subject pairs were fixed (treatment F) or randomly matched (treatment R) in each round. In treatments F, R, at the end of each round, subjects were told of

their payoffs and the strategy chosen by the opponent; over repeated rounds, subjects could figure out the payoff matrix.

The two main findings in the experiments are as follows. First, both predictions are verified. In particular, in all treatments, there is a higher incidence of the play of strategy R in game H relative to game L. In game L, and in the two limited information treatments, respectively, treatments R and F, the risky action R is played 74.6% and 49.9% of the time. The corresponding figures for game H are 20.3% and 44.6%. In treatment C, the strategy R is played 88.9%, 84.0%, and 79.0% of the time, respectively, in games L, H, and M. Overall, certain-loss avoidance is more prevalent in the data as compared to possible-loss avoidance. Second, loss avoidance has increasing bite as the game is played repeatedly over several rounds, revealing a role for subject experience.

12.4.12 Achieving coordination with gradual growth in group sizes

There is a view that coordination seems harder to achieve in larger groups relative to smaller groups. A nice summary of the evidence for the *minimum effort weak link game*, for fifth round choices of the players, is given in Table 12.19 from Weber (2006). Consider the game in Section 12.4.4 where the action of players is a choice $s_i \in \{1, 2, \dots, 7\}$ and the payoffs depend on the chosen action and the minimum action in the group. The coordinated play of 7, the highest action, is Pareto optimal, while for a given s_i , the worst outcome for player i occurs when the group minimum is the lowest action, 1.

For group sizes of 2, Van Huyck et al. (1990) find a high degree of coordination (86% play 7). However, as group sizes increase (even to size 3) there is a rapid drop in the coordination rate and for group sizes of 8 or more, the minimum for the fifth period, averaged across players/sessions, is the worst possible outcome, i.e., the play of action 1.

Weber (2006) conjectures that in real life, there must be a lot of coordination in large groups such as firms and communities. So how do large groups achieve coordination? The intuition in this paper is that successful groups grow gradually from a small enough size where cooperation might already be existent. The norm of cooperation is then learnt by successive entrants, who also cooperate.

In order to test this intuition, he runs three kinds of sessions. In the control session, 12 subjects played the weak link game for 12 periods. In the remaining two sessions, the *history* and the *no-history* sessions, groups were grown gradually from size 2 to size 12 across 22 periods. These two

Table 12.19 Distribution of fifth period group minimum in various minimum action games.

Minimum choice in fifth period							Group size	Number of groups	Source
7	6	5	4	3	2	1			
86%	3%	3%	3%	0%	0%	5%	2	37	VHBB, CK
18%	4%	0%	11%	15%	15%	37%	3	27	KC, CK
0%	0%	0%	0%	10%	10%	80%	6	10	KC
0%	0%	0%	0%	0%	0%	100%	8	5	CSS
0%	0%	0%	0%	0%	0%	100%	9	2	CC
0%	0%	0%	0%	0%	0%	100%	14–16	7	VHBB

Source: Van Huyck et al., 1990 (VHBB); Camerer and Knez, 2000 (CK); Knez and Camerer, 1994 (KC); Gerard P. Cachon and Camerer, 1996 (CC); Chaudhuri et al., 2001 (CSS). Weber (2006).

growth sessions started with only 2 players who played the weak link game initially for 7 periods. Then, additional individuals were added to the group, in a pre-announced and commonly known manner, until there were 12 individuals in the group by the last period. In the history session, entrants could observe the history of the group, while in the no-history session, they could not.

The results are as follows. First, the minimum action fell to zero in all five control groups with a group size of 12; this replicates earlier findings. Second, in the history and no-history sessions, small groups, of size 2, were able to coordinate successfully. Successful coordination is defined as the following condition: Prior to the growth of the group size, in all but one session, groups must have achieved the Pareto optimal outcome, of a group minimum of 7, for at least two consecutive periods. Third, it was extremely difficult to produce successful coordination even with good histories, i.e., those designed to produce coordination. For three of the history groups, by the time the groups achieved full size, the group minimum was the worst possible action, i.e., 1. However, in three history groups, a high level of coordination was achieved and the minimum stayed at the Pareto optimal level of 7 as the group expanded. Fourth, growing groups never produce coordination in the no-history condition. Thus, the success of “growing coordination” in the lab is mixed and depends on very specific conditions.

One might wish to see if people achieve greater coordination in one game, having successfully achieved coordination in another game, where, possibly, it was easier to achieve coordination. These portability issues are considered by Cason et al. (2012) who consider two games.⁵⁴ A median action game, in which coordination is easier to achieve, and a minimum action game in which coordination failures are more widespread. The authors define the concept of a *behavioral spillover*, which occurs if observed coordination in a game is higher when played in a sequence with another game rather than played on its own. Such spillovers can be positive or negative.

A possible application of these ideas is that firms may wish to allocate sequential tasks in order to maximize the degree of cooperation among the team members. The authors also wish to examine the effects on cooperation of stable groups, relative to groups whose members change randomly. For instance, Duffy and Ochs (2009) show that in contrast to random matching of opponents, the cooperative outcome in a prisoner’s dilemma game is more likely to emerge with a stable partner, which facilitates the development of a cooperative norm.

Cason et al. (2012) run four treatments. In SeqMin, the minimum action game is played for 10 periods, followed by 10 periods of the median action game. The treatment SeqMed reverses this order of play. In the treatment SeqMedDiff, the players are rematched after 10 periods, using a *perfect strangers matching protocol*.⁵⁵ In the treatment Sim, both games were played simultaneously. Table 12.20 shows the frequency of cooperation on the Pareto optimal allocation in the last round, round 10, in which each game is played.

The degree of cooperation reported in the median action game in these experiments (last column of Table 12.20) is unusually high relative to, say, Van Huyck et al. (1991) that was reviewed above. The most important result is evidence of behavioral spillovers. When the median action game precedes the minimum action game (SeqMed), there is much greater coordination in the minimum action game, relative to (i) playing the minimum action game on its own, or (2) simultaneously playing the two games (Sim).

⁵⁴ Other experiments have also noted such spillover effects of coordination; see, for instance, Van Huyck et al. (1991) and Camerer and Knez (2000).

⁵⁵ This protocol requires that no two players can be rematched together after having been previously matched once.

Table 12.20 Frequency of Pareto optimal coordination in period 10.

Treatment	Minimum action game	Median action game
SeqMin	10%	80%
SeqMed	80%	90%
SeqMedDiff	60%	100%
Sim	27%	73%

Source: Cason et al. (2012).

These results are in the spirit of Weber's (2006) findings. Successful experience of coordination (in the median action game) gives rise to behavioral spillovers that lead to greater coordination in the minimum action game. Indeed, even in the perfect strangers treatment (SeqMedDiff), there is greater cooperation in the minimum action game relative to the case where the minimum action game is played first (SeqMin). The transmission mechanism for behavioral spillovers is less clear. The authors construct measures of entropy that relate the cognitive ease of playing various games and advance this as a possible factor. Future research can fruitfully explore these issues further.

12.4.13 Intergenerational advice and coordination

We have seen above that in weak link games, unless groups are very small (say, size 2), the common finding is that coordination is hard to achieve. Chaudhuri et al. (2009) test an intergenerational version of the Van Huyck et al. (1990) model in which each generation can pass on some advice to the succeeding generation. The hope was that the cooperative outcome could arise with sufficient intergenerational advice, i.e., generations would, so to speak, *talk themselves to efficiency*.

Groups of 8 subjects play the minimum effort game for 10 periods. They are then replaced by their laboratory descendants who play for another 10 periods and so on. Several variants of these games were considered.⁵⁶ Game 1, the *replicator or no-advice game*, is as in Van Huyck et al. (1990), except that there are 8 subjects instead of 14–16 and no intergenerational advice is allowed; this is referred to as the VBB game (after the names of the authors). Game 2, the *progenitor game*, is a game played by a group of 8 subjects for the first time, with no advice. In Game 3, the *advice plus history game*, six generations played the game; each generation played the game for 10 periods. In this game, each subject could (1) view the history of all predecessors in all of his/her generations, and (2) receive *private* advice from the immediately preceding generation. Advice took the form of free-form messages. The incentive to provide advice arose from making the payoffs of each generation depend on the payoffs of the next generation. Game 4, the *advice only game*, is identical to Game 3, except that subjects cannot observe the history of the game. Game 5, the *public advice game*, is the most important for the results. This game was run for nine generations. The first generation in Game 5 received advice from the progenitor generation of Game 2. The first five generations of Game 5 played the *public advice with almost common knowledge condition*.⁵⁷ This took the form of advice from previous generations printed on pieces of paper that were available to all

⁵⁶ We mainly describe the results from Block I games. In Block II games, the authors also controlled for the quality of advice given.

⁵⁷ The authors use the *almost common knowledge* concept based on the idea of a common p -belief in Monderer and Samet (1989).




Figure 12.22 Minimum action chosen under different treatments in Block-I games.

Source: "Talking ourselves to efficiency: coordination in inter-generational minimum effort games with private, almost common and common knowledge of advice," Chaudhuri, Ananish, Schotter, Andrew, and Sopher, Barry, *Economic Journal*, 119: 91–122. Copyright © 2008 The Authors. Journal compilation © Royal Economic Society 2009.

subjects. All subjects were aware of this, but did not know how carefully, if at all, other subjects read the pieces of paper. The last four generations in Game 5 played the *public advice with common knowledge condition*. This was achieved by reading aloud the advice from previous generations.

The results are shown in Figure 12.22 and can be summarized as follows. First, the results of Game 1 are similar to the results of Van Huyck et al. (1990), except that there was an even lower level of coordination. The introduction of purely private advice, in Game 4, made matters even worse relative to Game 1, and relative to the Van Huyck et al. (1990) results. Game 3, the advice plus history game, elicits slightly more coordination. In the public advice game with almost common knowledge, there is some evidence of higher coordination in the initial generations, however, the minimum drops to a low level. Coordination levels improve significantly when the public advice game with common knowledge is introduced in generation 6. This immediately increases the group minimum action to 7 for all 10 rounds. In none of the subsequent generations/rounds did the minimum fall below 5 and was at the Pareto optimal level in 30 out of 40 rounds.

In Block II games, the same games, 1 through 5, are run, however, the quality of advice is controlled. In these games, the authors find a quality threshold. If advice is sufficiently exhortative, public, and exceeds the quality threshold, unanimously for all subjects, then the efficient outcome (a minimum of 7) is highly likely to be achieved.

The role of intergenerational advice has also been examined by Schotter and Sopher (2003, 2007). Schotter and Sopher (2003) consider the role of intergenerational advice in the following BoS game.

	1	2
1	150, 50	0, 0
2	0, 0	50, 150

The row and the column players have identical strategies, 1 and 2. There are two pure strategy Nash equilibria (1, 1) and (2, 2). There is also a mixed strategy equilibrium, which we ignore. A

baseline, *advice plus history*, treatment is run for 81 generations. Each generation plays the BoS game once and passes on advice, both in terms of the choice of strategy and free-form advice, to the succeeding generation. The incentive for providing advice arises because payoffs of one generation are made conditional on the payoffs of the next generation. Subjects in this treatment can also examine the history of play of all preceding generations.

When the baseline treatment had run for 52 periods, three further continuation games were run. In one, the baseline treatment was allowed to run for 29 more periods (till period 81). The remaining two continuation games were called Treatment I and Treatment II. Treatment I, the *advice only treatment*, begins with a history of 52 periods of play of the baseline treatment and then continues for 81 more periods. In this treatment, advice is offered to successive generations, but subjects are not allowed to view the history of the game. In Treatment II, the *history only treatment*, played for 66 periods (after a history of 52 periods of the baseline game), subjects could not seek any advice from the previous generation but could view the history of the game.

The results are shown in Figure 12.23 for the time periods following the 52 periods for which the baseline treatment had already run. In each treatment, we get *punctuated equilibria*. Punctuated equilibria are considered in Part 5. For the moment, these are equilibria in which a convention seems to get established for some time, only to disappear and then be re-established again. Interestingly, when a convention gets established, some generational advice comes along, suggesting a departure from the convention. At other times, once a convention is established, advice comes along to adhere to it, but such advice is often not followed.

The greatest success in achieving an equilibrium outcome arose in the baseline treatment (advice plus history); coordination is observed 47 times out of 81. History-only (Treatment II) is the least successful in eliciting a coordinated equilibrium; 19 out of 66 times and a continuation probability of the equilibrium of 0.315. Advice-only (Treatment I) did better than Treatment




Figure 12.23 Time paths of equilibria.

Source: Schotter and Sopher (2000).

Game A				Game B				Game C			
	1	2	3		1	2	3		1	2	3
1	5, 5	0, 0	0, 0	1	9, 9	0, 0	0, 0	1	7, 3	0, 0	0, 0
2	0, 0	5, 5	0, 0	2	0, 0	5, 5	0, 0	2	0, 0	5, 5	0, 0
3	0, 0	0, 0	5, 5	3	0, 0	0, 0	1, 1	3	0, 0	0, 0	3, 7

Figure 12.24 Three coordination games.

Source: Reprinted from *Games and Economic Behavior*, 4(4), John B. Van Huyck, Ann B. Gillette, and Raymond C. Battalio, "Credible assignments in coordination games", 606–26. Copyright © 1992, with permission from Elsevier.

II but worse than the baseline treatment. A coordination equilibrium was achieved in Treatment I in 39 out of 80 times with a continuation probability of an equilibrium equal to 0.512.

12.4.14 *External arbiter recommendations*

In each of the Nash equilibria in a coordination game, conditional on the play of the other players, it is a best reply for each player to play their part of the equilibrium strategy profile. Hence, external recommendations of an arbiter to players to play their part of the equilibrium strategy profile should be self-enforcing. Van Huyck et al. (1992) perform experiments to test this conjecture. They consider the three coordination games shown in Figure 12.24.

In game A, all coordination equilibria give identical payoffs. In game B, the coordination equilibrium (1, 1), is payoff dominant. In game C, all three equilibria give the same aggregate payoff but the coordination equilibrium (2, 2) is the most equitable. In game A, whenever the external arbiter suggested the play of any equilibrium, nearly 100% of the subjects complied with the recommendation. However, in games B and C, the arbiter's recommendations can be in conflict with the two focal principles of payoff dominance (game B) or equality (game C).

When the recommendations were in conflict with these two focal principles, they were ignored by a majority of the players. In game B, for instance, only 48% of the players complied with a recommendation to play strategy 2 and an even lower percentage, 38%, with a recommendation to play 3. In game C, in the absence of the recommendation, most subjects (82%) chose the equitable outcome (2, 2). When any of the inequitable outcomes was recommended, the compliance rate was less than 50%. Thus, in the presence of multiple equilibria, recommendations to play a particular Nash equilibrium are self-enforcing only if they do not conflict with some other salient principle, in this case payoff dominance and equity.

12.4.15 *Communication, experience, and coordination*

Brandts and Cooper (2007) and Cooper (2006) examine the role of communication and experience in the *corporate turnaround game* described in Section 12.4.10. This is a *weak link game* in which a manager interacts with four employees and offers them a fixed wage and variable levels of bonus in different treatments. The payoffs of the workers depend on the wage, bonus, the actions that they take, and also on the minimum action taken by any member in the group. Starting from a history of coordination failures, and, typically a low bonus level, can the manager take actions to turn around his company?

Brandts and Cooper (2007) modify the corporate turnaround game to include the possibility that (i) the manager sends messages to the workers (one-way communication), and (ii) the manager and the workers both send messages to each other (two-way communication). They also allowed for the possibility of a computer replacing the human manager in order to check if coordination is particularly enhanced by human to human contact.




Figure 12.25 Average minimum effort in different treatments and rounds.

Source: "It's what you say, not what you pay: an experimental study of manager-employee relationships in overcoming coordination failure," Brandts, Jordi and David J. Cooper, *Journal of the European Economic Association*, 5: 1223–68. Copyright © 2010 John Wiley and Sons.

The results are shown in Figure 12.25 that plots the average minimum effort for different rounds/treatments. All graphs begin from round 10; the minimum effort prior to round 10 is zero, i.e., there is a history of coordination failures. The computer-manager and the no-communication treatments are statistically indistinguishable and lead to low effort levels. Communication enhances effort and two-way communication is more effective in eliciting higher effort than one-way communication. Incentives in the form of bonuses enhance effort; however, communication is even more effective in eliciting effort. Hence, the title of the paper, "It's what you say not what you pay."

The authors conjecture that communication is cognitively more salient relative to bonuses. It may be more cognitively taxing to infer that higher bonuses are paid in the expectation of higher effort and that all other workers do interpret bonuses in the same manner. The exact form of communication that enhances coordination is surprisingly simple: request more effort, point out the mutual benefit of higher effort, and imply that employees are well paid. Insofar as the mutual benefits are financial, the communication strategy seems to stress financial incentives.⁵⁸

In an artefactual experiment, Cooper (2006) replicates the Brandts and Cooper (2007) experiments for experienced managers taking the executive MBA course.⁵⁹ Their performance in a

⁵⁸ The role of communication in this study ties in with results from other work in median action games, e.g., Blume and Ortmann (2007). However, one main difference, other than the application to the corporate turnaround game, is that in Brandts and Cooper (2007), communication is considered in conjunction with financial incentives.

⁵⁹ All managers have at least ten years of work experience of which at least five years were in a supervisory role. The average annual earning of a manager is in excess of \$120,000.




Figure 12.26 Average minimum effort level in different treatments and rounds.

Source: Cooper (2006).

corporate turnaround game, that has a history of coordination failures, is compared with the performance of undergraduate students. Feasible bonus levels, B , are restricted to the set of integers from 6 to 15, and the effort levels s_i are chosen from the set $\{0, 10, 20, 30, 40\}$, as in Section 12.4.10.

The results are shown in Figure 12.26 for the case of firms where, in round 10, there was coordination failure in the sense that the minimum effort level was 0. These results serve two purposes. First, they increase confidence in the external validity of the Brandts and Cooper (2007) experiments. From Figure 12.25, following a history of coordination failures, the managers achieve greater coordination slightly faster than the undergraduate students. However, by approximately round 13, the undergraduates catch up with the managers and the average minimum effort levels of the two subject pools looks very similar and comparable. Second, the method of communication that is effective in the lab experiments is very similar to that in the artefactual field experiments.

12.4.16 *Coordination in repeated games*

Dal Bó and Fréchette (2011) conduct experiments on cooperation in repeated games. The game that they consider is shown in Figure 12.27. The parameter c takes three possible values, 32, 40, 48, and they consider two possible continuation probabilities of the stage game, $p = \frac{1}{2}$ and $p = \frac{3}{4}$. This gives six possible treatments, each corresponds to the pair (c, p) . Subjects are randomly rematched at the end of each round of the stage game in each treatment and there is no discounting.

The treatments are designed to test various hypotheses. For instance, in the treatment $(c, p) = (32, 0.5)$, the only possible equilibrium is (D, D). To see this, consider the grim strategy that if a player observes D once, then he always chooses D. If C is played every period, then, conditional on the opponent also playing C every period, the payoff of the player is $\frac{32}{1-0.5} = 64$. Suppose

	C	D
C	c, c	12, 50
D	50, 12	25, 25

Figure 12.27 A coordination stage game.

Source: Dal Bó and Fréchette (2011) with permission from the American Economic Association.

that one party plays C every period and the other defects to D for one period. The payoff of the defector, given the strategies, is $50 + 0.5 \left(\frac{25}{1-0.5} \right) = 75 > 64$. In all other treatments, cooperation can be sustained in the repeated game as an equilibrium outcome. Furthermore, the following cases also make (C, C) a risk dominant equilibrium of the stage game: $(c, p) = (48, 0.5)$, $(40, 0.75)$, $(48, 0.75)$.

The results are as follows. In the case $(c, p) = (32, 0.5)$, where cooperation cannot be sustained as an equilibrium outcome, the outcome (C, C) converges to very low levels as players gain experience in the game, which is in line with earlier results, e.g., Bereby-Meyer and Roth (2006). However, the other cases in which cooperation can be sustained as a subgame perfect outcome, do not necessarily produce greater cooperation as subjects learn with experience. Thus, a cooperative prediction is necessary but not sufficient to produce cooperation in repeated game experiments. It is also not necessary that cooperation increases with experience when the cooperative outcome is subgame perfect and also risk dominant at the same time.

Comparing relative cooperation rates in finitely and infinitely repeated games, an earlier literature concluded that the gains in terms of the cooperative outcome were small; see Roth (1995) for a summary of the evidence. Some later studies, such as Dal Bó (2005), find larger gains in cooperation. Other findings reveal the beneficial effects for the cooperative outcome of (i) the quality of public signals in infinitely repeated games with imperfect public signals (Aoyagi and Fréchette, 2009), and (ii) fixed matching relative to random matching of opponents (Duffy and Ochs, 2009).

12.5 Bargaining games

In this section, we consider several kinds of bargaining games, their theoretical framework, and the relevant empirical evidence. Section 12.5.1 considers normative bargaining games, inspired by the seminal work of Nash (1950b). Section 12.5.2 considers bargaining models using non-cooperative game theory that are based on sequential offers and counteroffers by players (Stähl, 1972; Rubinstein, 1982). We consider the cases of full information, and one-sided and two-sided asymmetric information. Section 12.5.3 considers non-standard theories that do not fit any of the two frameworks in Section 12.5.1, 12.5.2. Finally, in Section 12.5.4, we explore the pattern of search and lookups of subjects, using MOUSELAB software, when they play bargaining games with a view to inferring their cognitive processes.

12.5.1 Normative bargaining solutions

The Nash (1950b) bargaining solution is the leading normative model of bargaining in economics. We first outline the model and then consider empirical tests of its axioms and predictions. The main focus here is on unstructured bargaining.

Consider a bargaining game in which there are two players, 1 and 2. The set of feasible expected utility payoffs of the players is denoted by $S \subset \mathbb{R}^2$; a typical element of S is denoted by $x = (x_1, x_2)$, where $x_i, i = 1, 2$ is the expected utility of player i . The disagreement payoffs of the players should they fail to conclude an agreement are given by $d = (d_1, d_2) \in S$. The Nash bargaining solution easily generalizes to more than two players.

Definition 12.1 A bargaining problem is denoted by (S, d) such that (i) S is compact, convex, and bounded above, and (ii) there exists at least one element $x \in S$ such that $x_i > d_i, i = 1, 2$. Let B denote the set of all bargaining problems. The solution to any problem $(S, d) \in B$ is given by a function $f(S, d) : B \rightarrow \mathbb{R}^2$, where $f(S, d) \in S$.

In Definition 12.1, S is convex, so all linear combinations of the possible bargaining outcomes are feasible. The compactness of S holds if, for instance, the feasible outcomes consist of a finite number of distinct alternatives and lotteries over them (recall that elements of S are the pairs of expected utilities of the two players). Nash was interested in the solution, f , when the following axioms hold.

1. Pareto optimality: Let $z = (z_1, z_2) \in S$ be the solution to the bargaining problem (S, d) , i.e., $f(S, d) = z$. Then there does not exist an $(x_1, x_2) \in S$ such that $x_i \geq z_i, i = 1, 2$. Thus, the bargaining solution does not lead to an outcome that is less desirable for both players relative to another feasible outcome.
2. Symmetry: (S, d) is a symmetric bargaining problem if (i) $d_1 = d_2$, and (ii) if $(x_1, x_2) \in S$, then $(x_2, x_1) \in S$. For symmetric bargaining problems, if $f(S, d) = (z_1, z_2)$, then $z_1 = z_2$. Thus, if the players are identical for all relevant purposes, then so must be the solution.
3. Independent of irrelevant alternatives: Let $(S, d) \in B$ and $(T, d) \in B$ such that $T \subset S$. If $f(S, d) \in T$, then $f(S, d) = f(T, d)$. Thus, if the best outcome in (S, d) lies in the smaller set T , then the outcome is unchanged if we consider the more restricted bargaining problem (T, d) .
4. Independence of equivalent utility representations: Consider some $(S, d) \in B$ and $f(S, d) = (z_1, z_2)$. Now consider another bargaining problem $(S', d') \in B$ such that the relation between (S, d) and (S', d') is as follows.

$$S' = \{(a_1x_1 + b_1, a_2x_2 + b_2) : (x_1, x_2) \in S\}; d' = (a_1d_1 + b_1, a_2d_2 + b_2), a_1 > 0, a_2 > 0,$$

then $f(S', d') = (a_1z_1 + b_1, a_2z_2 + b_2)$. This is a reasonable property because points in S are pairs of von Neumann–Morgenstern utilities, which are invariant with respect to positive affine transformations.

Remark 12.1 From the fourth property, independence of equivalent utility representations, the knowledge of the utility level of the other player should not affect the bargaining outcome. This is because the utility of the other player is invariant with respect to a positive affine transformation.

Proposition 12.4 (Nash, 1950b): Consider the bargaining problem $(S, d) \in B$. There exists a unique solution, $f^*(S, d)$, that satisfies Properties 1–4, above. This unique solution is given by

$$f^*(S, d) = \arg \max_{x \in S, x > d} (x_1 - d_1)(x_2 - d_2). \quad (12.10)$$

Example 12.3 Consider a Nash bargaining game in which players share a fixed sum of money, M , so that for any $x = (x_1, x_2) \in S, x_1, x_2 \geq 0$, and $x_1 + x_2 = M$. Then, from

Proposition 12.4, we have $f_i^*(S, d) = d_i + \frac{M-d_1-d_2}{2}$, $i = 1, 2$. Each player gets his disagreement payoff plus a share of half the net surplus. Thus, a higher outside option, d_i , ensures a more favorable bargaining position for the player. This solution is often called a “split the difference” solution.

The key properties that determine the Nash bargaining solution are 3 and 4. Pareto optimality (Property 1) is not critical. Suppose that we replaced Property 1 by Property 5, below.

5. Strong individual rationality: Consider the bargaining problem $(S, d) \in B$ and its solution $f(S, d)$. Strong individual rationality requires that $f(S, d) > d$.

Proposition 12.5 (Roth, 1977): Consider the bargaining problem $(S, d) \in B$. There exists a unique solution, $f^*(S, d)$, that satisfies Properties 2–5, above. This unique solution is given by (12.10).

The next proposition shows that symmetry (Property 2) is fairly innocuous too.

Proposition 12.6 (Roth, 1978): Consider the bargaining problem $(S, d) \in B$. Then, any solution, $f^*(S, d)$, that satisfies Properties 3–5 is of the form

$$f^*(S, d) = \arg \max_{x \in S, x > d} (x_1 - d_1)^{\theta_1} (x_2 - d_2)^{\theta_2}; \theta_1 + \theta_2 = 2 \text{ and } \theta_1 > 0, \theta_2 > 0. \quad (12.11)$$

Comparing (12.10), (12.11), symmetry simply imposes the restriction $\theta_1 = \theta_2 = 1$. In the absence of symmetry, we can consider other interesting issues such as differential bargaining powers of the two players; these are captured by the relative sizes of θ_1, θ_2 .

From the discussion above, the Nash solution fundamentally relies on Properties 3 and 4. Hence, these properties have been the main focus of empirical work that we shall review below.

Consider now the implications of relaxing symmetry. Define the maximum possible payoffs of each player, m_i , $i = 1, 2$, while still restricting to the set S as follows.

$$m_1 = \max \{x_1 : x_1 > d, (x_1, x_2) \in S\}; m_2 = \max \{x_2 : x_2 > d, (x_1, x_2) \in S\}.$$

Proposition 12.7 (Raiffa, 1953): Consider the bargaining problem $(S, d) \in B$. Then any solution, $h(S, d) = z$, that satisfies Properties 1, 2, and 4 (but not 3), satisfies

$$\frac{z_1 - d_1}{z_2 - d_2} = \frac{m_1 - d_1}{m_2 - d_2}.$$

In Proposition 12.7, the non-unique solution has the property that the gains to the two players in the optimal solution, relative to the disagreement point, are in the same proportion as the maximum possible gains relative to the disagreement point.⁶⁰

Two early empirical studies of the Nash bargaining solution were conducted by Nydegger and Owen (1975) and Rapoport et al. (1977). Both allowed unstructured face to face bargaining. The monetary payoffs were interpreted as the utilities of the players and players knew their own and the other player's monetary payoff; this is the *full information condition*.⁶¹ Both studies rejected Property 4, independence of equivalent utility representations. The Nydegger and Owen

⁶⁰ Kalai and Smorodinsky (1975) showed that an extra property (individual monotonicity) ensures that the solution in Proposition 12.7 is unique.

⁶¹ This created obvious problems (see Remark 12.1) that were rectified later by Roth and Malouf (1979), see below.

results supported Properties 1–3. By contrast, the Rapoport et al. results rejected Property 1 and observed violations of Property 2, although these were not statistically significant (this study was not designed to test Property 3). Even though there were departures from the Nash bargaining solution towards the equitable solution, the Nash solution was able to predict which player would come out ahead, say, players with a higher disagreement payoff. Although not directly related to the Nash bargaining solution, the early studies by Siegel and Fouraker (1960) and Fouraker and Siegel (1963) also found a tendency for the outcomes of bilateral bargaining outcomes to be in the direction of payoff equality.

Early studies that used the *partial information condition* (players know their own payoffs only), however, confounded between monetary prizes and the utilities of the players, which did not allow for a clean test of the theory. Roth and Malouf (1979) solved this problem through a novel framework that we consider below. Suppose that (S, d) is defined in monetary units. So S is the set of feasible monetary payoffs and $d = (d_1, d_2)$ gives the monetary payoffs if there is disagreement. Consider the larger class of bargaining problems B^* in which S is compact but not convex. Since the elements of the set S are no longer von Neumann–Morgenstern utilities, Property 4 does not apply. Consider, therefore, an alternative property, Property 6.

6. Independence of ordinal transformations, preserving interpersonal comparisons: Consider a bargaining problem $(S, d) \in B^*$ with the solution $f(S, d) = (z_1, z_2)$. Let $(x_1, x_2), (y_1, y_2) \in S$. Consider the transformations t_1 and t_2 that satisfy the following two conditions.

$$\begin{aligned} t_i(x_1, x_2) &\geq t_i(y_1, y_2) \text{ if and only if } x_i \geq y_i, i = 1, 2. \\ t_1(x_1, x_2) - t_1(d_1, d_2) &\geq t_2(x_1, x_2) - t_2(d_1, d_2) \text{ if and only if } x_1 - d_1 \geq x_2 - d_2. \end{aligned}$$

Now consider a second bargaining problem $(S', d') \in B^*$ that is related to (S, d) through the transformations t_1 and t_2 , i.e.,

$$S' = \{(t_1(x_1, x_2), t_2(x_1, x_2)) : (x_1, x_2) \in S\}; d' = (t_1(d_1, d_2), t_2(d_1, d_2)).$$

Then we require the relation between the solutions to problems (S, d) and (S', d') to be

$$f(S', d') = (t_1(z_1, z_2), t_2(z_1, z_2)).$$

Proposition 12.8 (Roth and Malouf, 1979): *Consider a bargaining game $(S, d) \in B^*$ and suppose that Properties 1, 2, 3, 6 are satisfied. Then the unique solution to the problem is the equal gains solution, $f = E$ such that $E(S, d) = z$ where z is a Pareto optimal point in S that satisfies, for all other Pareto optimal points (x_1, x_2) in S the condition*

$$\min\{z_1 - d_1, z_2 - d_2\} > \min\{x_1 - d_1, x_2 - d_2\}.$$

From Proposition 12.8, the solution to the problem under monetary payoffs and Properties 1, 2, 3, 6 is Pareto optimal and it maximizes the minimum gain to any player, while restricting to the set of Pareto optimal outcomes. This is the sense in which it may be termed as an equitable solution. Unlike Property 4 on which the Nash bargaining solution is based, Property 6 allows interpersonal comparisons of payoffs.

The earlier experiments, reviewed briefly above, did not allow players to know the von Neumann–Morgenstern utilities of others. Rather, under conditions of full information, players simply knew each other's monetary payoffs. The experiments of Roth and Malouf (1979)

permitted players to know the utility of the opponents, corresponding to the monetary payoffs. It is worth spelling this out in some detail because it formed the template for work in unstructured bargaining in several important subsequent contributions.

The principle behind this is as follows. Consider three outcomes, x_1, x_2, x_3 and suppose that an individual reveals the following preferences: $x_1 \succeq x_2 \succeq x_3$. If a utility function u represents the preferences $\succeq$, then we must have $u(x_1) \geq u(x_2) \geq u(x_3)$. Since u is defined only up to an interval scale, we may normalize $u(x_1) = 1$ and $u(x_3) = 0$. Consider the lottery $L = (x_1, p; x_3, 1 - p)$. Then the expected utility of this lottery is $EU(L) = p$. Suppose that we could elicit the probability p such that the individual reveals that he is indifferent between L and $(x_2, 1)$. Then, we must have $u(x_2) = p$. Thus, if players can observe the relevant p for other players, they can observe the utility function of the opponent. This key insight was then used to design the following experiment.

Two players bargained over the probability, p , that they receive a (possibly non-identical) prize, in an unstructured manner. This took the form of bargaining over lottery tickets. The proportion of tickets that each player gets in the lottery bargaining, gives the probability of winning a prize. For instance, a player who successfully bargains for 65% of the lottery tickets wins his prize with probability 0.65. If players fail to reach an agreement, both get nothing, so $d = (0, 0)$.

There were two information conditions. In the *full information condition*, players knew their own and the opponent's prize, while in the *partial information condition*, they only knew their own prize. The full information condition satisfies Property 6 and the partial information satisfies Property 4. The two players played four games in sequence, shown in Table 12.21. In Games 2 and 4, player 2 is restricted to get no more than 60% of the lottery tickets, while in the other two games there are no such restrictions.

To apply the method outlined above, let $u(x_1) = 1$ be the utility from receiving own prize and $u(x_3) = 0$ be the utility from not receiving it. Then for any intermediate outcome (winning prize with probability p) the expected utility is p , which is independent of the size of the respective prizes of the players. Games based on such experiments are known as *binary lottery games* (Roth and Malouf, 1979; Roth et al., 1981).

Under each information condition, both players play the four games in Table 12.21. Hence, the feasible utility payoffs, $S \in B$, are identical in each information condition. The sample is quite small (11 pairs in the full information condition and 8 pairs in the partial information condition). However, the value of this exercise is to understand how one properly constructs the relevant hypotheses and predictions of the various bargaining models.

Table 12.21 Description of the four bargaining games. D is the percentage difference between the lottery tickets received by player 2 and player 1. M and SD are the means and standard deviations of D.

Game	Prizes for individual players (\$)		Maximum % of lottery tickets allowed for individual players		D under Full Information		D under Partial Information	
	Player 1	Player 2	Player 1	Player 2	M	SD	M	SD
1	1	1	100	100	0	0	0	0
2	1	1	100	60	-1.909	12.169	1.325	8.335
3	1.25	3.75	100	100	-34.600	19.277	2.500	4.629
4	1.25	3.75	100	60	-21.636	22.482	-2.500	4.106

Source: Roth and Malouf (1979).

The four main predictions are as follows.

1. Properties 1–4 predict an equal 50–50 split in all four games: By definition, Property 1 (Pareto optimality), is consistent with all possible division of the 100 tickets between the two players. Since Game 1 is symmetric (recall that $d = (0, 0)$), Property 2 (symmetry) requires that each player gets 50% of the tickets. Since, a 50–50 split is optimal in Game 1, it should also be optimal in the more restricted Game 2 (Property 3, independence of irrelevant alternatives). Property 4 (independence of equivalent utility representations) requires Games 1 and 3 to have identical outcomes (use $a_1 = 1.25$, $b_1 = 0$, $a_2 = 3.75$, $b_2 = 0$ and recall that disagreement payoffs are zero). By the same reasoning, Games 2 and 4 should have identical outcomes.

Define D as the percentage difference between the lottery tickets received by player 2 and player 1.

2. Given the first prediction, the Nash solution (Proposition 12.4) predicts $D = 0$ in all four games. However, under both information conditions, Raiffa's solution (Proposition 12.7) predicts $D = 0$ for Games 1 and 3 and $D = -25$ for Games 2 and 4. Raiffa's solution does not require Property 3 (independence of irrelevant alternatives), so the outcomes in Games 1 and 2 will be different. The outcome in Game 1 under Raiffa's solution, like the Nash solution, is a 50–50 split. Using Proposition 12.7, in Game 2, $\frac{m_1 - d_1}{m_2 - d_2} = \frac{m_1}{m_2} = \frac{100}{60} = 1.6667$. Since all lottery tickets must be used up, the ratio of prizes must be $\frac{z_1}{z_2} = \frac{62.5}{37.5} = 1.6667$. Thus, the predicted Raiffa solution for Game 2 is a split of the lottery tickets to players 1 and 2 in the respective percentages 62.5 and 37.5. So $D = 37.5 - 62.5 = -25$. Property 4 applies to Raiffa's solution, so Games 1 and 3 should have identical outcomes, and so should Games 2 and 4.
3. In the “partial information condition,” Games 1 and 3 should have the same Nash bargaining solution and so should Games 2 and 4: Property 4 requires that each player should know only their own utility function. Since the utility function of the other player is invariant with respect to positive affine transformations, the bargaining solution should not be sensitive to the knowledge that a player has of the other player's utility functions.
4. In the full information condition, the data should respect Properties 1–3 but not necessarily Property 4: In the full information condition, players share utility and monetary payoff information; the latter enables Property 6 to hold and the solution is as in Proposition 12.8. Thus, we expect the solution to be between the Nash outcome, f^* , and the equal gains solution E .

The results are shown in Table 12.21 (M denotes mean and SD denotes standard deviation of D). There is no variation in the results in Game 1; all 19 pairs come up with 50–50 splits in both information conditions. This is consistent with the Nash and Raiffa solutions.

In the *partial information condition*, based on a comparison of the mean outcomes for D , the outcomes in Games 1 and 3 are statistically indistinguishable and so are Games 2 and 4. This is consistent with the third prediction above; hence, there is strong confirmation for the Nash bargaining solution. Thus, as predicted, in the partial equilibrium condition, players go in the direction of equal splits of the lottery tickets.

In the *full information condition*, the outcomes in Games 1 and 3 are statistically different and so are the outcomes in Games 2 and 4. In the full information condition, players bargain in the direction of splitting lottery tickets in such a way that equalizes the expected monetary payoff from the prizes. Thus, the player with the smaller prize tended to get a larger share of the lottery

tickets. The set of feasible utility payoffs is unaffected by the two information conditions, yet the actual outcomes are very different, hence, empirical studies that do not differentiate between the two cases do not fully capture the complexity of this sort of bargaining.

A total of 99% of the bargaining outcomes were Pareto optimal under both information conditions (this supports Property 1). Property 3, independence of irrelevant outcomes, holds in the full information case but is not supported in the partial information case for Games 3 and 4. This means all games are statistically identical, which supports prediction 1 for the Nash solution.

A series of further experiments were then conducted by Alvin Roth and several co-authors to explore extensions of the framework described above. Roth et al. (1981) looked for sociological factors such as norms or conventions that may govern the bargaining outcomes, in addition to the standard factors considered by theory. They conjectured that these norms, if any, must relate to known objects such as money. However, for unfamiliar objects, such norms may not apply. Their unfamiliar object is an intermediate good, *chips*, that can be converted into money at rates that may or may not be mutually known to the players. The authors find that the results accord with their conjectures.

Roth and Murnighan (1982) find that the bargaining results are sensitive to the information possessed by players with different levels of payoffs. For instance, a critical determinant of their bargaining results is whether the player with the smaller monetary payoffs (there were two prizes, \$5 and \$20) knows the prizes of both players. When the player with the \$5 prize knew both prizes, the bargaining outcomes were mainly equal division of the lottery tickets or equal expected values. However, when the player with the \$5 prize did not know the other player's prize, then the distribution of agreements was unimodal at equal divisions of the lottery tickets. Results are also sensitive to whether information on payoffs is common knowledge or not. Disagreements increased when it was not common knowledge that the \$5 player knew both prizes, possibly for reasons to do with reciprocity and intentions. These results suggest that there are nuances in the data that are not captured by any of the existing models, yet among the ones that are available, the Nash bargaining solution does reasonably well.

Murnighan et al. (1988) find that the *risk attitudes* of players have an influence on the bargaining outcome. In the first stage, they elicit the risk aversion parameters of the players. In the second stage, players play the binary lottery bargaining game, above. When the disagreement payoffs were high, more risk averse players did better as compared to games in which the disagreement payoffs were low. The disagreements in the bargaining game were found to be increasing in the size of the disagreement payoff.

Roth et al. (1988) document a *deadline effect*. Most of the bargaining agreements occur close to the deadline for the conclusion of bargaining. They wished to test experimentally anecdotal and folk wisdom that is encapsulated in popular phrases such as "eleventh hour agreements." Clearly, field data would be very hard to obtain for such testing, however, this question can easily be explored using experimental data. Subjects were allowed to bargain for a maximum of 9–12 minutes; bargainers could also send messages to each other over computer terminals but were prevented from violating anonymity. If subjects could not reach an agreement in the allocated time, they received a payoff of zero. The results are shown in Figure 12.28. Most of the action happens in the last 30 seconds; 41% of the agreements are reached during this time. Breaking down the last 30 seconds into 5 second components, 53% of the agreements occur in the last 5 seconds. The classical axiomatic approaches do not have a role for deadline effects.⁶²

⁶² The authors conjecture possible theoretical directions in section IV of their paper.




Figure 12.28 Frequency of agreements and disagreements over the whole period (left panel) and in the last 30 seconds (right panel).

Source: Roth et al. (1988) with permission from the American Economic Association.

What accounts for the bimodal distribution of outcomes at the equal split and the equal expected levels? Could it be that players are trying to coordinate on focal bargaining outcomes, and can coordination failures be used to explain these results? Initial work by Roth (1985) was supportive of an implication of this hypothesis, namely, that disagreements between players should increase as the relative payoffs become more inequitable. However, a later attempt by Roth et al. (1988) shows that the results are not very supportive of this hypothesis (in fact they do not get a bimodal distribution of outcomes).

There is also a significant theoretical literature on coalitional bargaining that also fits here, but there is not enough experimental evidence yet to enable one to make an assessment of its importance.⁶³

12.5.2 *Positive bargaining solutions*

BARGAINING UNDER FULL INFORMATION

Most experimental tests of positive bargaining models have revolved around the following two-person, sequential, bargaining game in which players have self-regarding preferences. Time is divided into discrete periods $\{1, 2, 3, \dots, T\}$ with generic element t . The problem is to agree on the division of a pie of size c . Delaying agreements is costly, say, due to time discounting, or loss in production as in a labor dispute. The discount factors of players 1 and 2 are, respectively, $\delta_1 \in [0, 1)$ and $\delta_2 \in [0, 1)$; these capture the cost of delay.

Player 1 makes the initial offer in period 1 and offers in every odd period, while player 2 makes counteroffers in even periods. In any odd period, t , an offer by player 1 is a proposed division of the pie $(x_1, x_2) = (x, c - x)$. The offer is either accepted or rejected by player 2. If player 2 accepts, the game ends with respective payoffs $\delta_1^{t-1}x$ and $\delta_2^{t-1}(c - x)$. If player 2 rejects, then the game

⁶³ The interested reader can consult the recent survey by Bandyopadhyay and Chatterjee (2012) and follow up the references therein.

proceeds to the next time period, $t + 1$. Since $t + 1$ is even, player 2 makes a new offer; the game now proceeds as in round t with the roles of the two players reversed. If an offer made at time T is rejected, then both players get their *disagreement payoff* of 0 each. Such a game is often called a T period game; see Ståhl (1972) who initially considered this problem formally. The solution is by backward induction.

Example 12.4 When $T = 1$, we get the ultimatum game that is considered in Chapter 5. In this case, since players are assumed to have self-regarding preferences, player 1 makes the smallest possible offer to player 2, which is immediately accepted by player 2; if c is denominated in cents, then $(x_1, x_2) = (c - 1, 1)$.

Most experiments consider the cases $T = 1, 2$, and 3. Let us consider the cases $T = 2, 3$ first and assume that money is infinitely divisible.

Example 12.5 Let $T = 2$. At $t = 2$ player 2 knows that disagreement payoffs are zero, so player 1 will accept an offer of zero, thus, he makes an offer $(x_1, x_2) = (0, c)$. At $t = 1$, when player 1 makes an offer, he must offer at least $\delta_2 c$, thus, $(x_1, x_2) = (c - \delta_2 c, \delta_2 c)$, which is immediately accepted by player 2. This is the unique backward induction outcome. Notice that when $T = 2$, the bargaining outcome is predicted to be independent of δ_1 , which is a testable prediction; we report the empirical tests, below.

Example 12.6 Let $T = 3$. Using the same logic as in Example 12.5, at $t = 3$, player 1 offers $(x_1, x_2) = (c, 0)$, which is accepted by player 2. At date $t = 2$, player 2 offers $(x_1, x_2) = (\delta_1 c, c - \delta_1 c)$, which is accepted by player 1. At $t = 1$, player 1 offers $(x_1, x_2) = (c - \delta_2(c - \delta_1 c), \delta_2(c - \delta_1 c))$, which is immediately accepted by player 2. In addition to testing for the unique backward induction outcome, there are two further implications that can be tested when $T = 3$. (a) The unique outcome is influenced by the discount factors of both players, δ_1, δ_2 . (b) Comparing the outcomes for player 1 in the two cases, $T = 2$ and $T = 3$, player 1 receives an extra payoff of $\delta_1 \delta_2 c > 0$ when $T = 3$.

In each of the examples above, the subgame perfect outcome is Pareto optimal, which in the current context requires that no money is wasted. The fraction of the shared pie depends only on the delay costs.

Proposition 12.9 (Rubinstein, 1982): As $T \rightarrow \infty$, in the alternative offers bargaining game (described above), player 1 offers to player 2 the following split of the pie that is immediately accepted by player 2,

$$(x_1^*, x_2^*) = \left(\frac{1 - \delta_2}{1 - \delta_1 \delta_2} c, \frac{\delta_2(1 - \delta_2)}{1 - \delta_1 \delta_2} c \right).$$

From Proposition 12.9, when the discount rates are identical, $\delta_1 = \delta_2 = \delta$, the bargaining outcome is $(x_1^*, x_2^*) = \left(\frac{c}{1 + \delta}, \frac{\delta c}{1 + \delta} \right)$. In particular, when players are sufficiently patient, so that $\delta = 1$, then we get $x_1^* = x_2^* = \frac{c}{2}$. Thus, the Nash bargaining or cooperative solution and the alternating offers bargaining solution, or the non-cooperative solution, coincide in this case.

We have already described the results of the ultimatum bargaining game ($T = 1$) in Chapter 5. In particular, Güth et al. (1982) who conducted the first such experiments, found that the modal proposal of players who played the game for the first time was a 50–50 share.

These experiments led to a paper by Binmore et al. (1985) that generated some debate. They begin with the observation (p. 1178) that “The work of Güth et al. seems to preclude a predictive

role for game theory insofar as bargaining behavior is concerned. Our purpose in this note is to report briefly on an experiment that shows that this conclusion is unwarranted.” They studied a two-player game with $c = 100$ pence and $T = 2$. The size of the pie shrunk from 100 to 25 from period 1 to period 2, so, the implied discount factors are $\delta_1 = \delta_2 = 0.25$. Thus, the prediction (see Example 12.5) is that in period 1, player 1 offers (0.75, 0.25) in dollar terms, which is immediately accepted. Example 12.5 relied on money being infinitely divisible. If the smallest unit is a cent, then the prediction is that player 1 makes an opening demand of 0.74–0.76 dollars and player 2 accepts any demand by player 1 that is lower than 0.74 dollars.

There were two features in Binmore et al. (1985) that created subsequent controversy. First, they played two bargaining games that reversed the order of the players in the following manner. In game 1, subjects in their role as player 2 switched to being player 1s in game 2. However, in game 2, player 1s did not know that, in fact, they were not playing against anybody; the experimenters simply wished to observe their choices subsequent to the role reversal. Second, the authors gave explicit instructions to the players (in capitals) that “You would be doing us a favour if you simply set out to maximize your earnings.” Insofar as actual bargaining considerations may involve other factors, e.g., fairness concerns, this potentially reduces the validity of the results.

The results were as follows. In game 1, the results of Güth et al. (1982) were replicated; modal settlements were 50–50 and the predictions in Example 12.5 are rejected. In game 2, after the role reversals, the modal demand by player 1 was just below 0.75 dollars and in 30 out of 81 cases, the demands were in the interval [0.74, 0.76]. The authors took this as support for classical bargaining theory and concluded (p. 1179): “The results (shown in Table 1) indicate little support for the view that a substantial proportion of the population are ‘fairmen’ as opposed to ‘gamesmen.’”

Güth and Tietz (1988) responded by conducting experiments on two games, similar to those in Binmore et al. (1985) with (i) $\delta_1 = \delta_2 = 0.9$ and, (ii) $\delta_1 = \delta_2 = 0.1$. They allowed for role reversals among players 1 and 2, exactly as in Binmore et al. (1985). However, they omitted the instruction: “You would be doing us a favour . . .” Using Example 12.5, the prediction in the first game is that player 1 offers $(x_1^*, x_2^*) = (0.10c, 0.90c)$, which is immediately accepted by player 2. In the second game, the prediction is $(x_1^*, x_2^*) = (0.90c, 0.10c)$.⁶⁴ The average demand by the first mover in the two games is, respectively, 0.76c and 0.67c. Thus, not only were they unable to replicate the earlier results, the results did not support the theoretical predictions of the classical model. The authors conclude “Our main result is that contrary to Binmore, Shaked and Sutton, ‘gamesmanship’ is clearly rejected, that is, the game theoretic solution has nearly no predictive power.”⁶⁵

Ochs and Roth (1989) then conducted a series of careful experiments using a 4×2 design shown in Table 12.22. They varied the basic alternating offers bargaining game along two dimensions: heterogeneity in discount factors, and the number of time periods, $T = 2, 3$. The games are played over 10 rounds with different opponents. In the first column in Table 12.22 are the combinations of discount factors, δ_1, δ_2 , used in the four games.⁶⁶

In period 1, each chip is worth \$0.30 to each bargainer. In period 2, each chip is worth δ_1 (\$0.30) to player 1 and δ_2 (\$0.30) to player 2. In period 3, these respective values are δ_1^2 (\$0.30) and

⁶⁴ One difference from the strict sequential game experiments was that if player 2 rejected the initial offer from player 1, yet in the counter offer claimed less than the initial offer, then this was treated as a disagreement.

⁶⁵ Another follow-up study by Neelin et al. (1988) imposed identical discount factors and considered two-player alternating offers bargaining games lasting over two, three, and five rounds. The results rejected the Stahl–Rubinstein model and also, equal splits.

⁶⁶ In order to implement different discount factors, the binary lottery procedure of Roth and Malouf (1979) is used but the lottery tickets or chips are redeemed for money at different rates for the two players.

Table 12.22 Description of various bargaining games in Ochs and Roth (1989).

	Two-period		Three-period	
	Chips	Money	Chips	Money
$\delta_1 = 0.4, \delta_2 = 0.4$	Cell 1:		Cell 5:	
	(59,41)	(\$17.70,\$12.30)		
	to	to	(76,24)	(\$22.80,\$7.20)
$\delta_1 = 0.6, \delta_2 = 0.4$	(61,39)	(\$18.30,\$11.70)		
	Cell 2:		Cell 6:	
	(59,41)	(\$17.70,\$12.30)		
$\delta_1 = 0.6, \delta_2 = 0.6$	to	to	(84,16)	(\$25.20,\$4.80)
	(61,39)	(\$18.30,\$11.70)		
	Cell 3:		Cell 7:	
$\delta_1 = 0.4, \delta_2 = 0.6$	(39,61)	(\$11.70,\$18.30)	(77,23)	(\$23.10,\$6.90)
	to	to	to	to
	(41,59)	(\$12.30,\$17.70)	(76,24)	(\$22.80,\$7.20)
$\delta_1 = 0.4, \delta_2 = 0.6$	Cell 4:		Cell 8:	
	(39,61)	(\$11.70,\$18.30)		
	to	to	(65,35)	(\$19.50,\$10.50)
	(41,59)	(\$12.30,\$17.70)		

Source: Ochs and Roth (1989).

δ_2^2 (\$0.30). Using Examples 12.5, 12.6, the subgame perfect equilibrium predictions for each game can now be readily calculated. These predictions are noted in cells 1–8 in Table 12.22 and are expressed in terms of chips and money (notice that the predictions are made over a range of values to reflect that chips can take only integer values).

One can now directly test if experimental data supports the predictions in Table 12.22. However, there are also two secondary *qualitative predictions*.

1. From Example 12.5, we know that when $T = 2$, the solution should only depend on δ_2 . Thus, the outcomes for cells 1 and 2 on the one hand, and cells 3 and 4 on the other, are predicted to be identical. Since player 2 is relatively more patient in cells 3, 4, he is predicted to receive a higher payoff as compared to cells 1, 2 where he is less patient.
2. From Examples 12.5, 12.6, we also know that, keeping discount factors constant, player 2 is predicted to receive a lower payoff when $T = 3$ rather than $T = 2$ (the opposite holds for player 1).

The results are shown in Figures 12.29, 12.30. The legend shows the information plotted for each cell, which includes the maximum and minimum offers. The dotted horizontal line shows equal division, the continuous horizontal line shows the perfect equilibrium offer and the hatched line shows the interval that contains the perfect equilibrium, taking account of integer values of chips.

In general, the subgame perfect prediction of the alternating offers bargaining game is a poor predictor of the experimental data. Other than cell 1, the estimated population mean plus minus two standard deviations does not include the theoretical prediction. Both qualitative predictions are also violated by the data. The first qualitative prediction fails because the difference in sample means between cells 1 and 2 is statistically different at the 95% confidence level; a similar result




Figure 12.29 Opening offers to player 2 for cells 1, 2, 5, and 6 in Table 12.22.

Source: Ochs and Roth (1989) with permission from the American Economic Association.

also holds for cells 3 and 4. In contrast to the second qualitative prediction, player 2 has a higher payoff when $T = 3$ rather than $T = 2$, in two out of the four possible comparisons. A second set of predictions refers to the acceptance decision of player 2 to the opening round offer. Theory predicts that the equilibrium offers should be immediately accepted. However, in 16% of the cases, the opening offers were rejected; this is a pervasive feature in bargaining experiments.




Figure 12.30 Opening offers to player 2 for cells 3, 4, 7, and 8 in Table 12.22.

Source: Ochs and Roth (1989) with permission from the American Economic Association.

Surprisingly, following a rejection of the opening offer by player 2, 101 of the 125 counterproposals by player 2 are on disadvantageous terms. Kennan and Wilson (1990) observe that in many strikes, the terms on which parties eventually settle are worse than the terms they rejected in earlier rounds of negotiations. Collating various studies, Roth (1995) shows that the incidence of bargaining disagreements in experimental and field studies is not very different. Despite the different methodologies, there is a striking similarity in the rejection behavior of subjects in Binmore et al. (1985), Neelin et al. (1988), and Ochs and Roth (1989): The respective percentage of first round rejections in the three studies is 15, 14, and 16, followed by disadvantageous counterproposals in

75, 65, and 81% of the cases, respectively. This is difficult to reconcile with the view that players are rational and have self-regarding preferences.

The preferred explanation of Ochs and Roth (1989) for such disadvantageous counteroffers is that (p. 365) “the unobserved element in the bargainer’s utility function may have a component related to the perceived ‘fairness’ of a proposal.” Thus, players may reject insulting offers even at the cost of a reduction in their payoffs. But this is not the only explanation. Güth and Tietz (1990) reject this explanation and favor an alternative one in which players switch between the heuristic of *strategic* and *equitable* considerations.

An important insight into the potential reason for disadvantageous counterproposals was given by Bolton (1991). He suggested that the utility of players depends on two arguments: own payoffs and relative payoffs. Thus, while players may make disadvantageous counterproposals, their relative payoffs from such proposals might be sufficiently greater to compensate for the lower absolute amounts. Indeed, this model formed the basis of the ERC model of Bolton and Ockenfels (2000) that we have considered in Chapter 6.

As a test of his proposal, Bolton (1991) considered two treatments. In the first treatment, players are paid tournament style incentives (winner takes all). In the second treatment, players are paid classical incentives, i.e., a share in the pie. The conjecture is that in treatment 1, since no pie is being shared, fairness considerations will not be as important, hence, we would get results more in accord with the classical theoretical prediction. There is partial support for the predictions. Inexperienced subjects show no noticeable effect of the two treatments. However, the behavior of experienced subjects, for some configurations of the discount factors, is consistent with the predictions. Clearly, a full account of the relevant data is an open problem.

De Bruyn and Bolton (2008) first use ultimatum games data reported in Roth (1995) to estimate a parameter of fairness. Then, they use this estimate to test the “out of sample” predictions of a new model, based on the original idea in Bolton (1991). This is a worthwhile attempt at discovering which kinds of models may explain results from bargaining games in a more satisfactory manner. They consider a utility function similar to that in Bolton and Ockenfels (2000), which is given by

$$U(r) = \begin{cases} \left(r - \frac{b}{2} \left(r - \frac{1}{2}\right)^2\right) c & \text{if } r < 0.5 \\ rc & \text{if } r \geq 0.5 \end{cases}, \quad (12.12)$$

where c is the size of the pie that is to be distributed in an ultimatum game, and r is the relative share of the player (we omit player subscripts). The utility function in (12.12) assumes that the reference point of the player is a share of one half of the pie. When the player gets less than half the pie, he suffers a reduction in utility, as his payoff falls further below the equitable share. The single parameter, $b > 0$, controls the extent of this reduction in utility. The utility function is unbounded below ($U \rightarrow -\infty$ as $b \rightarrow \infty$ and $r < 0.5$) but it is bounded above by c .

Players are assumed to choose their strategies with some noise, as in models of *quantal response equilibrium* (QRE); see Chapter 13 for the details. In a QRE, for each observed share r offered by the proposer, the responder plays a mixed strategy over the two actions: accept or reject. Taking account of the probability with which the responder chooses his two actions, the proposer also plays all his actions, various possible choices of r , with strictly positive probability. Each player plays the action that gives a higher expected payoff, with a higher probability; under QRE, players play *better responses* rather than *best responses*. A QRE is a fixed point in beliefs such that the strategy of each player is optimal, given the equilibrium beliefs and the strategies, in turn, generate the equilibrium beliefs.




Figure 12.31 Predictions of the model (bold line) versus the actual data in Roth et al. (1991) (dotted line) for offers and rejections in ultimatum games.

Source: De Bruyn and Bolton (2008). Reproduced with permission. Copyright, INFORMS, <http://www.informs.org>

The authors then use this model to predict out of sample data from the experiments reported in Roth et al. (1991). The fit is quite good, as can be seen from Figure 12.31. The authors show that their model predicts better in out of sample data as compared to the Fehr and Schmidt model (see Chapter 6).⁶⁷ The simple QRE also predicts the data quite well, but when it is combined with the ERC model, it predicts even better. For the results from the ultimatum game, this is an alternative explanation to a reciprocity-based explanation. The imposition of a reference point of $r = 0.5$ in (12.12) is reasonably close to its estimate in the data ($\hat{r} = 0.521$). Obviously, in other sorts of experimental games, such a reference point need not be the most suitable one, hence, the authors recommend that (p. 1786) “the reference point should not be assumed but rather freely estimated.”

There is evidence of greater efficiency of outcomes when in experiments subjects engage in face to face bargaining, relative to anonymous bargaining in which individuals make choices over computer terminals. Face to face bargaining may open up many other channels, such as social identity, body language, and facial cues, that would be blocked in anonymous bargaining. To the extent that these channels are important in real-life bargaining, anonymous bargaining misses out on some important factors, even as it allows extensive communication through computer terminals.

An early empirical study that attempts to address issues of communication is reported in Roth (1995, Table 4.6). He ran ultimatum game experiments ($T = 1$) with three conditions. The baseline condition is the standard anonymous game condition. A second condition, *unrestricted face to face communication*, allows unrestricted communication for 2 minutes, prior to the actual game. A third condition, *social face to face communication*, restricts communication to social conversation only (e.g., name and year in school) but not to conversation related to the bargaining game. The disagreement rate is 33% in the baseline condition but in the next two conditions, the disagreement rates, respectively, 4% and 6%, are not statistically distinguishable. These results suggest that it is not the content of the communication that is important in enhancing agreements. Rather, as

⁶⁷ Perhaps a fairer comparison should instead have been with a non-linear Fehr–Schmidt model augmented by the QRE framework.

Roth (1995, p.295) puts it: “powerful social motivations that may have little to do with bargaining may be responsible for the observed behavior.”

A related issue is that of experimenter scrutiny. Are the results affected because the subjects in an experiment perceive that they are being watched over by the experimenter? Early results by Hoffman et al. (1994) suggested that they do. But, Roth (1995, pp. 299–300) questions several aspects of these experiments.⁶⁸ However, in properly controlled experiments, Bolton and Zwick (1995) do not find evidence of experimenter scrutiny on the bargaining results. There is also evidence of the effect of cultural factors in bargaining when $T = 1$.⁶⁹

Binmore et al. (1989) examine the role of outside options. In Example 12.3, above, we considered a Nash bargaining game in which players share a fixed sum of money, M , and for any $x \in S$, $x_1, x_2 \geq 0$ and $x_1 + x_2 = M$. The solution is given by the *split the difference formula*. Continue to assume that the disagreement payoffs of both players are zero. Now introduce outside options for each player. In a sequential bargaining game, an outside option for any player is a number $b_i \geq 0$ such that if the player unilaterally exits the game, he receives a payoff of b_i . Assume that the outside options of the two players are given by the tuple (b_1, b_2) ; these payoffs satisfy $b_1 \geq 0, b_2 \geq 0$ and $b_1 + b_2 \leq M$. This alters the nature of the bargaining problem. At any date, the player making the offer must ensure that the other player gets at least the outside option. The set S is now defined as $\{(x_1, x_2) : b_1 \leq x_1, b_2 \leq x_2, x_1 + x_2 \leq M\}$.

Since $d_1 = d_2 = 0$, the solution to a Nash bargaining game, using the *split the difference formula* is $\frac{M}{2}$. In the presence of the outside option, if $b_i \leq \frac{M}{2}$ for $i = 1, 2$, then this remains the solution. However, if for any player, the constraint $b_i \leq \frac{M}{2}$ is violated, then this player gets $x_i = b_i$ and the other player gets $M - b_i$ (the constraint $b_i \leq \frac{M}{2}$ cannot be violated for both players because $b_1 + b_2 \leq M$). Such a solution is known as the *deal me out solution*.⁷⁰

Consider a pie of initial size £7 to be divided among two players, using a sequential alternating offers bargaining game in which player 1 makes offers in odd periods. In three different groups, the outside option of player 2, b_2 , is £0, £2, and £4. The outside option of player 1, $b_1 = 0$. At each stage, the payoffs are reduced by a factor of $\delta = 0.9$. The game terminates either if an offer is accepted, or a player takes his outside option, or the payoffs shrink to become negligible in some experimentally defined sense.

To explain the solution, suppose, for simplicity that $\delta = 1$, so that in the absence of the outside option, the *split the difference outcome* is (3.5, 3.5). Consider now the *deal me out solution*. When the outside option of player 2 is £0 or £2, we have the case $b_i \leq \frac{M}{2}$, so the 50% split, (3.5, 3.5), is still the equilibrium outcome. However, as soon as the outside option becomes greater than £3.5, as in group 3, where it is £4, then player 2 should get £4 (57.14% of the pie) and player 1 should get $7 - 4 = £3$ (42.86%) of the pie.

Now consider $\delta = 0.9$. In the *deal me out solution*, when $\delta = 0.9$, player 2 gets 47.3% of the pie in groups 1 and 2 and 57.1% of the pie in group 3. Split the difference (Example 12.3) predicts the shares of player 2 in the three groups to be, respectively, 47%, 64%, and 76%.

The results are shown in Figure 12.32. The horizontal axis is the fraction of pie offered by player 1 to player 2. The height of the histograms captures the frequency of each offer for each of the three groups. The deal me out solution is shown by the solid line connecting the predicted equilibrium

⁶⁸ For instance, variations in the degree of anonymity were confounded with other factors and some of the results are reported for some games in the paper but not for others.

⁶⁹ See, for instance, Roth et al. (1991) and other results reported in Chapter 5.

⁷⁰ The alternating offers bargaining game with an outside option that corresponds to this solution is described in Appendix 1 in Binmore et al. (1989).




Figure 12.32 Player 1's offer of the pie to player 2.

Source: Binmore, K. G., Shaked, A., and Sutton, J., "An outside option experiments," *Quarterly Journal of Economics* (1989) 104: 753–70, by permission of Oxford University Press.

points (47.3% in groups 1 and 2 and 57.1% in group 3). The split the difference solution is also similarly shown as the dashed line.

The main contribution of Binmore et al. (1989) is to show that the deal me out solution describes the data from experimental bargaining games better, relative to the split the difference solution. There are a large number of bargaining disagreements (see the hollow squares and the squares with a dot in Figure 12.32) that are inconsistent with the predictions of any of the main sequential bargaining theories. However, focusing just on those bargaining outcomes where there are agreements, there are spikes in the data that are consistent with the predictions of the deal me out solution. In group 2 games, in 11 out of 25 games, the share of player 2 was between 50% to 51%, while in group 3 games, in 7 out of 25 games, the share was between 57% to 58%. The results for group 3 are the most interesting. In 4 out of 25 cases, player 2 received a fraction of the pie that was well below 57.1%; this is inconsistent with both deal me out and split the difference solutions. Of the remaining 21 cases, in 18 cases the agreed splits are closer to the deal me out solution as compared to the split the difference solution. The authors also note that their results are consistent with fairness considerations in bargaining, conditional on players not settling below their outside option.

Zwick et al. (1992) test for the predictions in Proposition 12.9 by implementing an experimental design in which there is a probability, p , in each round, that the game will terminate. Thus, in any round, if a player rejects an offer, then the game ends with probability p . Otherwise, with probability $1 - p$, the player can make a counter offer. They chose three values of $p = \frac{1}{10}$, $\frac{1}{3}$, and $\frac{5}{6}$. When $p = \frac{5}{6}$, this game is similar to an ultimatum game (which corresponds to $p = 1$), and when $p = \frac{1}{10}$ and $\frac{1}{3}$, respectively, the game is expected to last for 10 and 3 rounds. The results are as follows. First period demands are neither consistent with the predictions in Proposition 12.9,

nor with equal splits. However, focusing attention on the final accepted offers, mean demands were closer to equal splits rather than to the predictions of Proposition 12.9.

BARGAINING UNDER ASYMMETRIC INFORMATION

One-sided asymmetric information

We explore first *one-sided asymmetric information* in the bilateral bargaining game experiments of Rapoport et al. (1995). Consider a single indivisible object that a seller, S, wishes to sell to a buyer, B, at some price p_t at time $t = 0, 1, 2, \dots$. The seller's cost of the object, c , is known to both parties, but the buyer's valuation, v , is private information. The seller believes that v is drawn from some commonly known density of prior beliefs, $\phi(v)$. At any time, t , the seller sets a price, p_t , that the buyer can either accept or reject. If the buyer rejects, then the game proceeds to the next round where the seller makes a new price offer, p_{t+1} . If the buyer accepts at time t , then the game ends and the payoffs of the seller and the buyer are, respectively, $\delta^t(p_t - c)$ and $\delta^t(v - p_t)$, where $0 \leq \delta \leq 1$ is the common discount factor.

We are interested in a sequential equilibrium of this bargaining game. This is a price sequence $(p_t)_{t=0}^{\infty}$ and the binary decision of the buyer $d_t \in \{0, 1\}$ $t = 0, 1, 2, \dots$ to accept ($d_t = 1$) or reject ($d_t = 0$) such that the following conditions hold. Conditional on any history of the game: (i) The price p_t , at any time, maximizes the profit of the seller, conditional on his belief about the type of the buyer at time t and the subsequent strategies of the players. (ii) The buyer's decision at any time t , d_t , is optimal, given his beliefs and the subsequent strategies. (iii) Beliefs should be consistent with Bayes' rule wherever possible.

Proposition 12.10 (Fudenberg et al., 1985; Gul et al., 1986) *Suppose that $c = 0$, and ϕ is a uniform distribution.*

(i) *If ϕ is uniform over $(0, 1]$, then there is a unique equilibrium sequence of prices given by*

$$p_t = \gamma^t p_0, \quad t = 1, 2, 3, \dots \quad (12.13)$$

where,

$$p_0 = \frac{\gamma(1-\delta)}{1-\gamma\delta}, \quad \gamma = \frac{1}{\delta} \left(1 - \sqrt{1-\delta}\right). \quad (12.14)$$

The strategy of a buyer with valuation v is to accept whenever $p_t \leq \frac{v}{\gamma} p_0$.

(ii) *If ϕ is uniform over $[0, 1]$, then the equilibrium is not unique.*

The intuition behind Proposition 12.10 is not difficult to understand. Both sides have market power. The seller's market power comes from being a monopolist, hence, he can control the sequence of prices. The buyer's market power arises from his superior information. The seller would like to sell his good early on account of time discounting, yet he does not want to set "too low" a price because the buyer's valuation might be high.

When $0 < \delta < 1$, it follows from (12.14) that $0 < \gamma < 1$, hence, the optimal price sequence in (12.13) is a strictly decreasing one. The seller sets a high price initially, and gradually decreases it over time at an exponential rate. The idea is to engage in *intertemporal price discrimination*, so that he can sell initially to the high valuation buyers and later to the low valuation buyers. The reason that high valuation buyers do not wait long enough for the price to drop even further is that the rate of price drop slows down sufficiently over time, so that the extra benefit from a lower price is swamped by the delay costs on account of time discounting.

When $\delta = 1$, the buyers are infinitely patient and are willing to wait forever for the price to drop. In this case, the seller will want to sell at the earliest date by offering the lowest possible price, $p_0 = 0$ (because $c = 0$ in Proposition 12.10), thus, the seller's profit is zero. At the other extreme ($\delta \rightarrow 0$), using L'Hopital's rule, $\lim_{\delta \rightarrow 0} p_0 = \lim_{\delta \rightarrow 0} \gamma = 0.5$. A simple calculation shows that p_0 is decreasing in δ . Thus, $p_0 \leq 0.5$.

This discussion outlines the predictions of the model that can be experimentally tested. Three values of δ are used: a high value $\delta = \frac{9}{10}$ (condition H), a medium value $\delta = \frac{2}{3}$ (condition M), and a low value $\delta = \frac{1}{3}$ (condition L).⁷¹ The authors chose the distribution of valuations ϕ to be uniform over $[0, 1]$; from Proposition 12.10, this allows a multiplicity of equilibria that also include the equilibrium described in Proposition 12.10(i). A discrete counterpart of the distribution ϕ was implemented by letting the buyer's valuations take equally probable integer values \$0, \$1, . . . , \$100.

Consider first the prediction $p_0 \leq 0.5$ and that prices are declining in δ . For each value of δ , mean p_0 falls as δ increases, exactly the opposite of the theoretical prediction. When evaluated at the level of individuals, the theoretical prediction holds only for 2 out of 48 subjects. Mean p_0 is significantly higher than 0.5 in condition H for all iterations, statistically not different from 0.5 in condition M and lower than 0.5 in condition L.

A second prediction is the exponentially declining prices. Prices do fall over time in all sessions, however, they do not fall exponentially, which is critical to the incentive compatibility of the mechanism. A linear trend with an extra quadratic term fits as well as an exponentially declining trend. Experienced sellers choose linear descending prices.

The seller's profit should fall as δ increases and in the limit, as $\delta \rightarrow 1$, the profit should be zero. The predictions for the buyer's profits in this zero-sum game are just the reverse. However, the results show that the seller's profits are close to \$25 for all values of δ and in all sessions. The expected valuation of the buyer is \$50, and the seller's expected profits are just half of this, so there could be some bargaining norms in operation such as equal expected splits. Even when high valuation buyers have had experience of the game, they do not attempt to free ride on low valuation buyers. They typically accept the first or second instance where they find that the offered price is lower than their valuation. Thus, they accept the seller's offers "too soon" relative to the predictions of theory. This feature contributes to the high observed efficiency of the outcomes.

Overall, the authors take the position that their results constitute a rejection for the sequential equilibrium predictions of theory. They write (p. 391): "The equilibrium predictions capture very little of the qualitative features of the data . . . we find that subjects follow simple rules of thumb in choosing strategies, reflected in the behavioral consistencies revealed in the study . . . Rather, the decision rules adopted by buyers and sellers violate any possible sequential equilibrium of the mechanism under investigation, or described by fictitious strategies."

Forsythe et al. (1991) consider two-player, pie-splitting games in which an informed player, I, knows the size of the pie. The uninformed player, U, believes that the state of the world is *good*, pie size π_g , with probability p , and the state of the world is *bad*, pie size $\pi_b < \pi_g$, with probability $1 - p$. Both players are risk neutral, expected utility maximizers. Players are allowed to bargain for a fixed length of time and if they do not successfully conclude an agreement, both players get zero. The players are allowed to send unconstrained messages to each other, hence, the structure of the experiments makes it difficult to write down the corresponding theoretical model.

Following Myerson (1979), any Nash equilibrium of this game may be implemented by a direct revelation mechanism of the following type. Player I moves first and announces the state of the

⁷¹ Each of these conditions was iterated three times in each session. In session 1, all 48 subjects participated, and in sessions 2 and 3, a subset of 16 subjects participated. In each session, each subject played 18 two-player bargaining games; 9 as seller and 9 as buyer.

world: good, g , or bad, b . Player U can then shrink the pie by a fraction $1 - \gamma_j$, $\gamma_j \geq 0$ in state $j = g, b$; this corresponds to a *strike* by the player. Thus, in the good and bad states, respectively, the pie available for redistribution reduces to $\gamma_g \pi_g$ and $\gamma_b \pi_b$ (so $\pi_j - (1 - \gamma_j) \pi_j = \gamma_j \pi_j$). Letting $x_j \geq 0$ be the share of player 2 in state $j = g, b$, the state-dependent payoffs of players I, U are, respectively, $(\gamma_g \pi_g - x_g, x_g)$ and $(\gamma_b \pi_b - x_b, x_b)$. A mechanism is denoted by $\mu = (\gamma_g, x_g, \gamma_b, x_b)$. An incentive compatible (IC) mechanism must ensure that in each state, player I reveals the true state of the world, thus

$$\gamma_g \pi_g - x_g \geq \gamma_b \pi_g - x_b \text{ and } \gamma_b \pi_b - x_b \geq \gamma_g \pi_b - x_g. \quad (12.15)$$

The inequalities in (12.15) imply that an incentive compatible mechanism must satisfy

$$(\gamma_g - \gamma_b) \pi_b \leq x_g - x_b \leq (\gamma_g - \gamma_b) \pi_g. \quad (12.16)$$

A mechanism is *efficient* if there exists no allocation of outcomes such that at least one of the following three players (each type of player I is a distinct player) can be made better off, without making anyone worse off: Player U, player I who observes the good state, and player I who observes the bad state.

Proposition 12.11

- (i) *Strikes are not efficient in the good state, i.e., $\gamma_g = 1$.*
- (ii) *In an efficient mechanism, the right-most inequality in (12.16) must bind.*

Proof: (i) First notice that since $\pi_b < \pi_g$, (12.16) requires that $\gamma_g \geq \gamma_b$, thus, the incidence of strikes in the good state is always lower. Suppose that we have a mechanism $\mu = (\gamma_g, x_g, \gamma_b, x_b)$ such that $\gamma_g < 1 - \varepsilon < 1$. Now consider an alternative mechanism $\mu^* = (\gamma_g^*, x_g, \gamma_b^*, x_b) = (\gamma_g + \varepsilon, x_g, \gamma_b + \varepsilon, x_b)$. Player U's payoff is unaffected because x_b, x_g are unchanged. However, in each state of the world, a smaller fraction of the pie is destroyed due to strikes, hence, the payoff of player I must increase. Thus, the original mechanism $\mu = (\gamma_g, x_g, \gamma_b, x_b)$ in which $\gamma_g < 1$ is not efficient. It follows that in any efficient mechanism, $\gamma_g = 1$, and no strikes take place.

(ii) Suppose that $x_g - x_b < (\gamma_g - \gamma_b) \pi_g$. Then, a small increase in γ_b that does not disturb the inequality, does not affect the payoff of player U. But this increases the payoff of player I, because less of the pie is now shrunk, thus, such a move is efficient. Continue to increase γ_b until $x_g - x_b = (\gamma_g - \gamma_b) \pi_g$. ■

Consider a non-cooperative game in which player U makes a non-negotiable (take it or leave it) demand to player I. A well-known result in bargaining theory shows that the risk-neutral player U must find it optimal to demand π_g whenever $p\pi_g > \pi_b$, where p is the probability of the bad state (otherwise player U demands π_g). This is known as the *strike condition*. The same condition holds in a cooperative game (the optimal mechanism) as we now show.

Proposition 12.12 (Forsythe et al., 1991):

- (i) *Strikes are not efficient when $p\pi_g < \pi_b$.*
- (ii) *However, in an efficient mechanism strikes can exist if $p\pi_g > \pi_b$.*

In order to predict the shares of the players (determined entirely by x_g, x_b), the authors use the *random dictator axiom* in Myerson (1984). The outcomes in this case are as if each of the two players chooses the allocation as a dictator in a dictator game, with equal probabilities. Thus, we get the mechanism

$$\mu = \begin{cases} (1, \frac{1}{2}\pi_g, \frac{1}{2}, 0) & \text{if } p\pi_g > \pi_b \\ (1, \frac{1}{2}\pi_b, 1, \frac{1}{2}\pi_b) & \text{if } p\pi_g < \pi_b \end{cases} \quad (12.17)$$

When player I proposes in the dictator game, he imposes the mechanism $\mu = (1, 0, 1, 0)$ in which strikes never occur ($\gamma_g = \gamma_b = 1$) and player U gets a zero payoff ($x_g = x_b = 0$). When player U proposes in the dictator game, there are two possibilities. If $p\pi_g > \pi_b$, the player should always demand a pie of size π_g , so he imposes the mechanism $\mu = (1, \pi_g, 0, 0)$; thus, strikes occur if the state turns out to be bad. If $p\pi_g < \pi_b$, then strikes are never optimal, so player U imposes the mechanism $\mu = (1, \pi_g, 1, \pi_b)$ in which he gets the entire surplus in each state. Noting that the proposals made by the two players are equiprobable gives us (12.17).

Four games were considered in the experiments. The strike condition is “on” in games I and II ($p\pi_g > \pi_b$) and “off” in games III and IV ($p\pi_g < \pi_b$). If the predictions of the theory are found to hold, then one may be able to conclude that strikes are perhaps an efficient response to asymmetric information.

The results are as follows. Incentive compatible efficiency (Proposition 12.11) enables one to sort out individual-level data in a large number of cases. For instance, strikes happened more frequently in games I and II as compared to games III and IV. However, the predictions of the theory are not perfect. Strikes do happen in good states in games I and II and also in the bad state in games III and IV (where they are not predicted to happen at all). However, these violations happen in about 8% of the cases.

One does obtain *bargaining impasse* (strikes) even under full information (see Section 12.5.3). However, this paper makes the sharper prediction that the difference in the incidence of strikes in games I and II, relative to games III and IV, is significant. Thus, to a large extent, these results indicate that strikes are a response to information asymmetries between players. However, this does not solve all the puzzles. The incidence of strikes should be similar across games I and II, yet the incidence is significantly different (39.2% in game I and 17.2% in game II; the prediction in both games is 50% because $p = 0.5$). This result holds true not only for the bargaining games but also for the games where the players choose as dictators in order to test Myerson’s random dictator axiom (Proposition 12.12).

Myerson’s random dictator axiom is not very useful in organizing the data. The authors conjecture that the reason for this failure was a violation of the basic assumptions of the model such as risk neutrality and homogeneity of preferences by the subjects in the experiment.

Two-sided asymmetric information

Radner and Schotter (1989) consider bargaining experiments with *two-sided asymmetric information*. The focus is on a particular mechanism, the *sealed-bid mechanism*. Consider a single indivisible object that a seller, S, wishes to sell to a buyer, B, in a static game. The seller’s valuation (or cost), v_s , and the buyer’s valuation, v_b , are private information to the respective players. The joint distribution of v_s and v_b is given by Φ , and it is commonly known. The privately known values, v_s and v_b , are also sometimes known as *reservation values* because the seller will not accept a price below v_s and the buyer will not pay a price above v_b .

In the sealed-bid mechanism, the seller and the buyer simultaneously submit bids. The seller’s bid, p_s , is an asking price and the buyer’s bid, p_b , is an offer price. If $p_b \geq p_s$, then the two parties trade the object at a price $p = \frac{p_b + p_s}{2}$. If $p_b < p_s$, then there is no trade and each party gets a payoff of zero. In the event of an agreement, the respective payoffs of the seller and the buyer are, respectively, $\pi_s = p - v_s$ and $\pi_b = v_b - p$.

This is a static Bayesian game in which each party has a type: the seller’s type is v_s , drawn from some set of “cost” types C and the buyer’s type is v_b , drawn from some set of “valuation” types V .

In equilibrium, each player chooses a type-contingent strategy. When both v_b, v_s are drawn from uniform distributions over $[0, 100]$, then Chatterjee and Samuelson (1983) showed that the following type-contingent linear bidding strategies constitutes an equilibrium (an exercise asks you to prove it).

$$p_b = \begin{cases} v_b & \text{if } v_b < 25 \\ \frac{25}{3} + \frac{2}{3}v_b & \text{if } v_b \geq 25 \end{cases} \quad (12.18)$$

$$p_s = \begin{cases} v_s & \text{if } v_s > 75 \\ 25 + \frac{2}{3}v_s & \text{if } v_s \leq 75 \end{cases} \quad (12.19)$$

Both optimal bidding strategies are stepwise linear. The buyer's bidding strategy equals the true valuation if $v_b < 25$ but thereafter, the buyer *shades his bid* (because the bid increases less than one for one with v_b). By contrast, the seller initially shades his cost until $v_s = 75$ and thereafter, bids his true cost. Figure 12.33 shows a plot of (12.18) and (12.19). The diagonal line shows *truthful bids*. The seller's optimal bid is shown by the dashed line that is either above or on the diagonal, so he typically tries to overstate his true cost. The buyer's optimal bid, shown by the continuous line, reveals that the buyer typically understates his valuation.

To test the predictions of this model, one can test if the behavior of subjects is sensitive to the cutoff points, as shown in (12.18) and (12.19). The equilibrium in linear strategies is not the only equilibrium. There are a multitude of other equilibria in this model (Leining et al., 1989). However, the equilibrium in linear strategies, when players are not allowed preplay communication, is the most efficient across all possible equilibria; this holds even when preplay communication is allowed prior to bargaining. It is most efficient in the sense that it maximizes the ex-ante expected total profits $E(\pi_s + \pi_b)$, where the expectation is taken with respect to Φ (Myerson and Satterthwaite, 1983). An important motivation behind the experiments is to see if subjects choose the equilibrium in linear strategies, and if such an equilibrium is efficient.

Radner and Schotter conducted eight different experiments. The optimal linear strategies in (12.18) and (12.19) come from the first set of experiments. We only describe their experiments




Figure 12.33 Optimal linear bidding functions for sellers and buyers.

Source: Reprinted from *Journal of Mathematical Psychology*, 42: 133–66, Daniel, T. E., Scale, D. A., and Rapoport, A., "Strategic play and adaptive learning in sealed bid bargaining mechanism." Copyright © 1998, with permission from Elsevier.

1, 2, and 5 that give the clearest test of theory. In the second set of experiments, at the end of each round, the bids of all participants were publicly revealed (without compromising subject identity). In the fifth experiment, the distributions for v_b and v_s were non-uniform but symmetric.

The two main findings of the experiments, based on data pooled across all buyers on the one hand and across sellers on the other, are as follows. First, an eyeball test suggests that both sellers and buyers seem to employ linear bidding strategies. For the seller, this is confirmed by running regressions of p_s on an intercept term, v_s , v_s^2 and v_s^3 . An F-test did not reject the hypothesis that the coefficients of v_s^2 and v_s^3 are jointly zero. Likewise, when one runs a regression of p_b on an intercept term, v_b , v_b^2 and v_b^3 , an F-test does not reject the hypothesis that the coefficients of v_b^2 and v_b^3 are jointly zero.

Second, it can be confirmed that there are cutoff points that are consistent with (12.18) and (12.19) such that the behavior of buyers and sellers is different on each side of the cutoffs. For the sellers, the actual bidding strategies on both sides of the cutoff are consistent with the predictions. But buyers did not shave their bids as much as predicted and they change their bids almost one for one with their valuations. This was a factor in the high actual efficiency of this mechanism in the experiments.

One feature of the Radner and Schotter (1989) study is that they maintain fixed pairings of subjects, so one may wonder if their results are confounded by reputational factors. Daniel et al. (1998) address this concern in an experimental design that has three differences from the Radner–Schotter study. First, they conduct their experiments over 50 rounds to allow for learning possibilities. Second, they randomly rematch partners after every round to reduce reputational confounds. Third, the experiment is conducted on computer terminals, rather than manually, and each subject is given the complete history of his bids so far.⁷² In experiment 1, the authors attempt to replicate the results of Radner and Schotter, and in experiment 2, they test some additional features of the solution.

In experiment 1, the seller's cost is distributed uniformly over the interval $[0, 100]$ and the buyer's valuation is distributed uniformly over the interval $[0, 200]$. The optimal linear bidding strategies in this case are given by

$$p_b = \begin{cases} v_b & \text{if } v_b \leq 50 \\ \frac{50}{3} + \frac{2}{3}v_b & \text{if } 50 < v_b \leq 150. \\ 116.7 & \text{if } v_b > 150 \end{cases} \quad (12.20)$$

$$p_s = 50 + \frac{2}{3}v_s \text{ for all } v_s \in [0, 100]. \quad (12.21)$$

The optimal bidding strategies are shown in the left panel of Figure 12.35. Relative to the experiments in Radner and Schotter, the buyer's optimal strategy has two kinks, hence, this is an even more demanding test of theory.

In experiment 2, the buyer's valuation is distributed uniformly over the interval $[0, 200]$ as in experiment 1. However, the seller's cost is distributed uniformly over a very narrow interval $[0, 20]$. The optimal linear bidding strategies in this case (see right panel of Figure 12.34 for a pictorial representation) are given by

⁷² This experimental design (with the three newer features) was initially used by Rapoport and Fuller (1995). They replicated the findings of Radner and Schotter. In the main, they found that subjects either behaved in accordance with the optimal linear bidding strategies or they simply stated their reservation values truthfully.




Figure 12.34 Optimal bidding strategies for buyers and sellers in experiment 1 (left panel) and in experiment 2 (right panel).

Source: Reprinted from *Journal of Mathematical Psychology*, 42: 133–66, Daniel, T. E., Scale, D. A., and Rapoport, A., “Strategic play and adaptive learning in seated bid bargaining mechanism.” Copyright © 1998, with permission from Elsevier.

$$p_b = \begin{cases} v_b & \text{if } v_b \leq 50 \\ \frac{50}{3} + \frac{2}{3}v_b & \text{if } 50 < v_b \leq 75. \\ 63.3 & \text{if } v_b > 75 \end{cases} \quad (12.22)$$

$$p_s = 50 + \frac{2}{3}v_s \text{ for all } v_s \in [0, 20]. \quad (12.23)$$

The focus of this study is on individual behavior (as compared to average behavior), which provides a more accurate test of the theory. Ten buyers and ten sellers are randomly matched with each other over 50 rounds for each of the two experiments. The results are interpreted using *spline regressions*.⁷³

For experiment 1, three of the ten sellers are truth-tellers (so they bid their costs), while the behavior of the remaining seven sellers is consistent with their predicted bidding strategy in (12.21). Except for one buyer, the regression fitted well either the linear optimal bidding strategies of buyers (six out of ten buyers), or buyers are truth-tellers. For all ten buyers, in the middle range, $50 < v_b \leq 150$, the optimal linear strategy, $\frac{50}{3} + \frac{2}{3}v_b$, fits well the observed play. Thus, the more satisfactory experimental design successfully replicates the findings of Radner and Schotter.

Consider now the results of experiment 2. For all ten sellers, the bidding behavior falls between the optimal linear bidding strategy and truth-telling. The R^2 values are quite low. It is as if sellers bid independently of their reservation values. In the last 25 rounds, sellers adjust their bids downwards, relative to the first 25 rounds, suggesting evidence of learning. Several other bargaining studies do not find evidence of learning; see, for instance, Neelin et al. (1988) and Weg et al. (1990). It is likely that subject feedback, as in Daniel et al. (1998), plays an important role in the efficacy of learning. In contrast to sellers, the bidding behavior of buyers is more consistent with linear bidding strategies, although buyers typically bid below their optimal linear strategies. Sellers earned significantly less, while buyers earned significantly more as compared to the predicted optimal linear strategies. It is as if the sellers struggled to settle on their optimal bids in this case, and buyers took advantage of suboptimal bidding by sellers.

⁷³ In other words, regressions are run for each of the regions that separates the cutoff points in the optimal strategies and the ends of the best fit lines are smooth-pasted together.




Figure 12.35 The trade and no-trade regions in a two-sided asymmetric information bargaining game.

Source: Reprinted from *Games and Economic Behavior* 38: 127–55, Valley, K., Thompson, L., Gibson, R., and Bazerman, M., “How communication improves efficiency in bargaining games.” Copyright © 2002, with permission from Elsevier.

Radner and Schotter (1989) found that when face to face communication and bargaining occur simultaneously, the bargaining outcomes achieve nearly full efficiency. Valley et al. (2002) design sealed-bid experiments to isolate the source of this efficiency in a two-player bargaining game with unstructured bargaining and varying degrees of pre-bargaining communication. Suppose that the valuation of buyers and sellers is distributed uniformly over $[0, 1]$. The optimal linear bids of the two players are then given by

$$p_b = \frac{1}{12} + \frac{2}{3}v_b \text{ and } p_s = \frac{1}{4} + \frac{2}{3}v_s. \quad (12.24)$$

Trade only takes place if $p_b \geq p_s$, or, using (12.24), if $v_b - v_s \geq \frac{1}{4}$; the relevant region in which trade takes place is shown in Figure 12.35 in the top left hand corner. In the experiments, the authors assume that the valuations of the two players are drawn uniformly from the interval $[0, 50]$, and trade takes place if $p_b \geq p_s$, or if $v_b - v_s \geq 12.5$.

The authors use a 3×2 design; there are three prebid communication conditions (no/written/face to face⁷⁴) and two feedback conditions (no feedback/feedback⁷⁵). Thus, the sealed-bid mechanism had four stages. In stage I, players received information on the realization of their private values, v_b, v_s . In stage II, one of the three communication treatments were run. In stage III, bargaining took place via the sealed-bid mechanism, and in Stage IV, one of the two feedback treatments was implemented. Trade only took place when $v_b > v_s$, which occurred for 192 out of the 343 observations; we shall mainly focus on these results below.

⁷⁴ In the written communication, players were given a 13 minute period in which to exchange unrestricted messages delivered to each other by courier. In the face to face treatment, players could communicate for a 6 minute period with each other.

⁷⁵ In the no feedback treatment, players were told of the outcome once all the rounds of play were completed. In the feedback treatment, players were informed at the end of each round, if a trade had occurred or not, the price at which a trade had occurred, and profits, etc.

In Figure 12.36, the observations that represent trade are shown as filled-in diamonds and those in which no trade took place are shown as hollow squares. The two diagonal lines shown in Figure 12.36 are the same as those in Figure 12.35 in order to separate the trade and the no-trade regions. The sealed-bid mechanism under no communication predicts that no trade should occur within these lines (this area is also known as the Chatterjee–Samuelson no-trade region). This prediction is borne out in panel A in Figure 12.36, confirming the earlier results of Radner and Schotter (1989). However, once we allow for communication, in panels B and C, this result no longer holds. Now one observes a lot of trades within the diagonal lines, more so under face to face communication as compared to written communication.

Define *high value trades* as those where $v_b > v_s + 12.5$ and *low value trades* where $v_s < v_b < v_s + 12.5$. In the no communication treatment, the proportion of high and low value trades was, respectively, 89% and 11%. In the written communication treatment, these figures were 83% and 44%, while in the face to face treatment, these figures were 94% and 74%. Thus, communication, especially face to face communication, was particularly effective in realizing not only high value trades but also low value trades. Thus, face to face communication captured 94.1% of the total available surplus from trade. This is significantly better than the corresponding figure of 82.3% under no communication. However, there is little evidence of any learning over successive rounds.

What accounts for greater efficiency under communication? The authors find that the actual strategy of players is significantly different from the predicted strategy in which bids depend on own valuation alone. A strategy in which one's bid depends on the valuation and bids of others, fits much better the bargaining data. For instance, fitting the classical linear strategies for the buyer as in (12.24), the regression line is $p_b = 3.24 + 0.65v_b$ and $R^2 = 0.57$. However, adding further explanatory variables about the seller's valuation and bids gives the regression $p_b = -0.52 + 0.25v_b + 0.10v_s + 0.57p_s$ and $R^2 = 0.85$; both newly added variables were individually significant.

The authors term bidding strategies where bids depend on the types and actions of others as *dyadic strategies*. In fact, in 73 of the 129 cases where there are positive gains from trade, in the presence of communication both players name an identical price. It need not be the case that both players reveal their true valuations to each other (although 8 out of 129 did), however, even if sellers overstate their costs and buyers understate their valuations, they are mostly able to negotiate on an identical price that allows trades to take place. Furthermore, there is a lower inaccuracy in stated valuations when communication is face to face rather than written. Hence, coordination achieved by interpersonal communication allows the observed efficiency level to exceed the Myerson–Satterthwaite upper bound.

12.5.3 *Self-serving bias and focal points in bargaining*

In this section, we consider two theories of bargaining that are motivated by neither the normative framework nor by the framework of the sequential alternating offers bargaining game. Güth (2012) makes a particularly strong criticism of the approach described in Sections 12.5.1 and 12.5.2, and takes up important issues that we do not consider here.⁷⁶

⁷⁶ We encourage the interested reader to browse through the edited book by Bolton and Croson (2012) for a rich variety of alternative models of bargaining that we have not discussed here. These models incorporate the roles of overconfidence, beliefs, trust, procedural fairness, reference points, gender, and culture, etc.




Figure 12.36 Trade outcomes for the no-communication, written, and face to face treatments.

Source: Reprinted from *Games and Economic Behavior* 38(1), Kathleen Valley, Leigh Thompson, Robert Gibbons, and Max H. Bazerman, "How communication improves efficiency in bargaining games," 127-55. Copyright (2002), with permission from Elsevier.

Section 12.5.2 highlighted one possible factor for bargaining disagreements and impasse: the presence of asymmetric information. We now consider the role of *self-serving bias* in creating such *bargaining impasse* under full information. Insofar as such explanations are supported by the evidence, at the very minimum, they show that asymmetric information may be sufficient, but not necessary for observed bargaining disagreements. A self-serving bias arises when people attempt to rationalize their actions by either selectively drawing only supporting evidence from memory, or using selectively self-serving aspects of the evidence from a situation.

Loewenstein et al. (1993) created a hypothetical situation in which a plaintiff sues a defendant for damages worth \$ x . Subjects are then randomly allocated to the roles of the plaintiff and the defendant, provided with the relevant case material, and asked to guess the damages awarded by a judge. The differences in their estimates are statistically significant and are partial to their own roles in the situation. Yet, both parties have access to identical information, so issues of asymmetric information are absent.

Babcock et al. (1995) test this hypothesis further by considering two treatments that differed in the timing of the assignment of roles as plaintiff and defendant. In treatment 1, the roles were assigned prior to handing over of the case material, while in treatment 2, the roles were assigned afterwards. Consistent with the predictions of a self-serving bias, players gave more biased judgments in treatment 1. The authors also tried several treatments to debias the self-serving bias of subjects. There are two findings. First, it is quite hard to debias players. Second, while the debiasing procedure may make players more aware of the self-serving bias of others, they nevertheless show little awareness of their own bias.

In another debiasing experiment, Babcock et al. (1996) made subjects aware of the self-serving bias and told them that listing the weaknesses in their own reasoning may reduce the bias. The subjects were then asked to list the weakness in their own case. This did significantly reduce the bias. Such a debiasing procedure may also have policy implications.

Babcock and Loewenstein (1997) outline a simple theory of bargaining impasse. Parties in a bargaining situation may be subject to a self-serving bias in that they confound between what really serves their own interests and what is a fair bargaining outcome. This hampers bargaining negotiations in the following ways. First, by focusing on self-serving agreements, the bargaining parties narrow down the set of possible agreements. Second, if parties believe that their self-serving bargaining position is also fair, then they interpret the intransigence of the other party in rejecting their offer as unfair and opportunistic. Insofar as issues of reciprocity may be involved in the bargaining process, this hardens the bargaining position of both parties. Third, the bargaining parties are typically unwilling to settle below their perceived fair distribution of the pie (Loewenstein et al. 1989).

The authors make a case for reinterpreting the existing findings of bargaining under asymmetric information using their framework. While the authors make plausible suggestions (pp. 119–121) on how this could be done, one would need sharp, controlled, experiments to test these issues more formally. For some further potential pitfalls in testing theories of asymmetric information in the lab, see Roth (1995, Section 2a).

Isoni et al. (2013) test for the insights of Schelling (1960) on *tacit bargaining*, using experimental games. Schelling argued that payoff-irrelevant cues may serve to coordinate the actions of players on some mutually beneficial bargaining outcome. It might well be that the players benefit asymmetrically from such coordination but being rational, they realize that if all players were to follow the relevant cues, they would be better off. Thus, such cues may serve as an *equilibrium selection device*. There are many possible payoff-irrelevant cues, such as spatial cues, geographical cues, and precedent based cues.




Figure 12.37 Games G1 and G2.

Source: Isoni, A., Poulsen, A., Sugden, R., and Tsutsui, K. (2013). "Focal points in tacit bargaining problems: experimental evidence." *European Economic Review* 59: 167–88.

Example 12.7 (Schelling, 1960, pp. 58–67): *Two opposing army commanders look to maximize the territory that they can gain, yet not come into conflict with each other. How far should they go? When shown a map, the respondents in Schelling's study choose a river between the two territories as the limit of advance of each commander. The cue here is spatial and it is payoff irrelevant. Both commanders gain asymmetrically unless the river divides their advance exactly evenly.*

We shall consider the evidence from coordination games as applied to level- k models in Chapter 13 (Crawford et al., 2008). The empirical evidence in Crawford et al. (2008) shows that the salience of some cues that successfully induce coordination in symmetric games may be lost in asymmetric games. One of the aims of Isoni et al. (2013) is to show that the salience of cues can still be a useful coordinating device in asymmetric games. The authors conjecture that the difference in results can perhaps be reconciled on two grounds. (i) In Crawford et al. (2008), the games were framed as matching games, while they framed them as bargaining problems. (ii) There are differences in the strategic complexity of the games in the two sets of papers.

Figure 12.37, also known as a bargaining table, shows 2 games out of the 24 games used in the study. Each of two players is assigned a color, light gray or dark gray. The left square in each game is colored light gray and is called the light gray player's base while the right square is colored dark gray and it is termed, the dark gray player's base. The numbers on discs in each bargaining table are amounts of money in UK pounds. Each player must choose one or both discs. An agreement is reached only if each player claims the disc that the other does not; in this case, players get the monetary payoff shown on their chosen disc. If the choice of discs overlaps, then both players get zero.

Each bargaining table in Figure 12.37 has nine columns. Taking out the middle column in either of games G1 or G2, the four columns on the left and the right sides constitute the *left and right sides* of the bargaining table. It is conjectured that these sides may have spatial salience for the choices of the two players. Players whose base is on the left side are called left players and players whose base is on the right are called right players.

Which disc should each player choose? Perhaps they could use some spatial-cues and choose the disc closest to their base. Using this reasoning, in the left hand side bargaining table of Figure 12.37, the left side player may choose the disc with number 6 and the right side player may choose the disc with number 5. If both players followed such spatial cues, then both may successfully coordinate.




Figure 12.38 Description of a bargaining table for one of the 24 games in Isoni et al. (2013).

Source: Isoni, A., Poulsen, A., Sugden, R., and Tsutsui, K. (2013). "Focal points in tacit bargaining problems: experimental evidence." *European Economic Review* 59: 167–88.

Each player has the following choices: choose none of the discs, the near disc, the far disc, or both discs. Writing out the normal form for the game in the left panel of Figure 12.37, there are two strict Nash equilibria, (near, near) and (far, far) with respective payoffs (6, 5) and (5, 6). In the game shown in the right panel of Figure 12.37, the payoffs from the two strict Nash equilibria are both (5, 5), hence, this game is termed as an *equality-compatible game*. Can players use some tacit agreements to play one of the two strict Nash equilibria?

The results in Figure 12.37 are compared with those from the game in Figure 12.38, where spatial cues based on proximity are absent, because both discs are equidistant from the light gray and dark gray bases. If it turns out that there is relatively more coordination in the games in Figure 12.37, then it provides evidence for spatial cues.

The results are as follows. Spatial cues provide a strong coordinating device. In particular, there is an *own side effect*; 69–82% of the left players choose a disc from the left side of the table. By contrast, only 13–21% of the right players choose a disc from the left side; 64–81% of right players choose a disc from their own side. Only 7–16% of left players choose a disc from the right side. The differences in choices of the players are confirmed to be statistically significant. Thus, the central column appears to have been perceived as a boundary by the players, akin to the river in Example 12.7. The spatial focal points retain their significance in both equality compatible and equality incompatible games, suggesting a role of such cues even in asymmetric games. However, the class of games with interesting economic insights that also have spatial focal points might not be very large.

12.5.4 Inferring cognition through search and lookups in a bargaining game

One indirect test of game theoretic predictions is to see if the human cognitive process is consistent with traditional game theoretic reasoning. Of course, the solution concepts in game theory, such as Nash equilibrium, do not make any direct assumptions about the cognitive process. Indeed, a static Nash equilibrium does not even make any assumptions about the beliefs of players about other players. Hence, observations on the cognitive process of humans do not necessarily constitute a direct refutation.

Some economists may argue that if the predictions of game theory are accurate, then even incorrect assumptions about the cognitive process do not constitute a refutation.⁷⁷ There are two main responses to these arguments; see the introductory chapter for more details. First, as shown in this chapter, several predictions of classical game theory fail the “direct test” of not being in accord with the evidence. Second, if a model with refuted assumptions makes correct predictions, it could be purely due to chance or due to the role of assumptions that are not properly understood within the model.

This leaves open the question of how one may observe the cognitive processes used by humans when they play games. Suppose that we have a static, full information, 2×2 normal form game in which the row and the column players have two strategies each. Suppose also that we can observe the pattern of the *lookups* and *searches* of the players as they try to compute their strategies. Consider three kinds of searches by the row player. (1) Up and down over one’s own payoffs. (2) Left or right over the opponent’s payoffs. (3) Transiting frequently from own to the opponent’s payoffs. The first two patterns of searches are consistent with a player using game theoretic solution concepts such as Nash equilibrium. The third pattern of search seems to be suggestive of a concern for one’s payoffs relative to those of other players (i.e., social preferences).

One important method of identifying the pattern of lookups and searches is offered by the use of the MOUSELAB software.⁷⁸ This method is increasingly used in behavioral economics. We explain the method, using a three-round bargaining game from Johnson et al. (2002).⁷⁹

In the spirit of the alternative offers sequential bargaining game of Rubinstein (1982), suppose that two players, 1 and 2, bargain over the division of a pie over three rounds ($T = 3$; see Example 12.6). The players make offers sequentially, and Player 1 makes the first move. When it is the turn of any player to make a move, the player offers a split of the pie. The other player either agrees (in which case, the agreed upon split is made) or rejects and makes a counteroffer. No counteroffers can be made subsequent to round 3. The size of the original pie is \$5, but it shrinks to \$2.50 in the second round, and \$1.25 in the third round, to capture time discounting; otherwise players do not discount future payoffs. The outside option of each player is zero.

Suppose that we take as given the typical game theoretic assumptions, such as common knowledge of rationality, and selfish concern for own-payoffs. The subgame perfect equilibrium can now be computed by starting from the optimal offers in round 3 and folding the game tree backwards. In round 3, when it is the turn of player 1 to move, and the size of the pie is \$1.25, he offers \$0.01 to the opponent and proposes to keep \$1.24 for himself. It is individually rational for player 2 to accept this offer because the outside option is zero. In round 2, when it is the turn of player 2 to make the offer, he knows that any offer to player 1 that is lower than or equal to \$1.24 will be rejected, so he makes an offer of \$1.25 to player 1 and keeps $\$2.50 - \$1.25 = \$1.25$ for himself. In round 1, when player 1 makes the offer, he knows that any offer to player 2, lower than or equal to \$1.25, will be rejected by player 2. Thus, he offers \$1.26 to player 2 and keeps $\$5.00 - \$1.26 = \$3.74$ for himself. Thus, in the subgame perfect equilibrium, player 1 offers \$1.26 to player 2 (at the very first node), keeping \$3.74 for himself and this is immediately accepted

⁷⁷ This methodological view goes back to the “as if” reasoning of Milton Friedman (see the introductory chapter). Gul and Pesendorfer (2007) use a similar argument to criticize the research agenda in neuroeconomics that we consider elsewhere.

⁷⁸ See, for instance, Payne et al. (1993) for an overview of studies using the MOSUSELAB software in the domain of individual choice.

⁷⁹ Several important results that we refer to in this section were produced earlier in the pioneering work of Camerer et al. (1994).



Figure 12.39 The MOUSELAB computer display for the three-round bargaining game.

Source: Reprinted from *Journal of Economic Theory* 104: 16–47, Johnson, E. J., Camerer, C. F., Sen, S., and Rymon, T., “Detecting failures of backward induction: monitoring information search in sequential bargaining.” Copyright © 2002, with permission from Elsevier.

by player 2. The objective of Johnson et al. (2002) is to use MOUSELAB software to test if the cognitive process of the players is consistent with this outcome.

It is worth understanding the basic structure of the MOUSELAB interface, as well as the arguments for and against it. We shall also encounter the use of the MOUSELAB interface in the context of tests for the level- k model, below. The MOUSELAB computer display in the three-round bargaining game of Johnson et al. (2002), described above, is shown in Figure 12.39.

The display has two columns of three boxes each. The boxes in the first column show the size of the pie in each round, while the boxes in the second column show the role of the player in each round. The contents of each box appear as blanks. In order to examine the contents of any box, subjects can use the mouse to move the cursor to that box, which immediately opens and reveals the contents. In Figure 12.39, the subject has moved the cursor to the first box in column 1, which reveals the contents to be \$5. The box immediately closes when the cursor is moved outside the box. In order to make an offer, players move the cursor over the dollar scale which is present horizontally at the bottom of the box. Subjects can then click any amount to make an offer. The software records the activity of the subjects every 16 milliseconds, which provides a detailed account of which boxes are opened and the order in which they are opened. This pattern of information acquisition reveals the lookups and searches of each player.

An important assumption in interpreting the results of the MOUSELAB experiments is that the pattern of information acquisition mirrors the structure of the cognitive process or information usage. Critics, steeped in traditional game theory, might object to the MOUSELAB method in the bargaining game by giving the following two arguments.

1. It is possible that subjects in experiments are memorizing the information relevant for the game by searching the game forward, as revealed in their information acquisition process. However, once memorized, they could be using the information to solve the game

backwards. Hence, information search during the play of the game may not be particularly informative of the cognitive process. Johnson et al. (2002) give the following response to this potential criticism. First, memorizing the information content of boxes is more cognitively taxing than opening boxes and searching for information in MOUSELAB. Second, players opened the same box more than once, suggesting that they were not successfully memorizing. Third, from one game to the next, the size of the pies was altered, making memorization even more difficult, while keeping the equilibrium outcome unchanged.⁸⁰

2. Game theory makes no assumptions about information search and information usage. In order to test this, Johnson et al. (2002) look at a subset of the subjects who are trained in the use of backward induction. The search and lookups of these subjects can then be compared to non-trained subjects to detect any differences.

If subjects behaved in the manner assumed by classical game theory, we should expect that they play the equilibrium outcome. Call such subjects “equilibrium players.” If players are equilibrium players, then player 1 should offer \$1.26 to player 2, keeping \$3.74 for himself, which should be immediately accepted by player 2.

How might boundedly rational subjects play the bargaining game? Johnson et al. (2002) make the following conjecture.⁸¹ Suppose that players are of two types, naive and partially sophisticated. A player 1 who is naive, interprets any round of play as an ultimatum game and ignores future rounds of play. Such players make an offer of 40% of the pie (the average offer in ultimatum game experiments), i.e., \$2.0 to player 2 in round 1. A player 1 who is partially sophisticated, believes that all other players are naive. Such players look one step ahead, and ask what would be the counteroffer made by a naive player in the next round. For instance, in round 1, a partially sophisticated player 1 thinks that in round 2, the counteroffer made by a naive player 2 will be 40% of the pie of size \$2.50 (so player 1 will be offered \$1 and player 2 will keep \$1.50). Thus, looking ahead, in round 1, a partially sophisticated player 1 offers \$1.51 to player 2, which should be immediately accepted if the beliefs of player 1 are accurate.

Proposition 12.13 *The first round offer made by player 1 to player 2 is as follows: Equilibrium players offer \$1.26, naive players offer \$2.0, and partially sophisticated players offer \$1.51.*⁸²

In the trained subjects treatment, subjects are trained in backward induction. The findings are as follows.

1. Baseline treatment: In the baseline treatment, untrained human subjects play against untrained human opponents. The average first round offer by player 1 is \$2.11. Offers below \$1.80 are rejected half the time. This is in line with earlier results from similar bargaining experiments. The information searches using MOUSELAB reveal that players spend less time searching backward among the boxes, relative to searching forward, which is contrary to the expected lookup patterns of equilibrium players. In 19% and 10% of the trials, respectively, players do not even look once at the second and the third round pie sizes. On average, subjects in the role of player 1 spend 62% of their time looking at the first round pie size. About 31% of the time is spent looking at the size of the second round pie, leaving

⁸⁰ The size of the pies was \$5 + a, \$2.50 + b, \$1.25 + c, such that $a - b + c = 0$. The reader can check that the equilibrium outcome is identical in each case to the game in the text with $a = b = c = 0$.

⁸¹ The conjecture is based on the supposition that all players are *level-k* types. We give a formal treatment of the level- k model in Chapter 13.

⁸² Depending on the tie breaking rules, these numbers can respectively be \$1.25, \$2.0, and \$1.50.

only about 7% of the time for looking at the third round pie. Furthermore, a strong correlation is found between searching the boxes in a backward direction and the closeness of one's offer to the equilibrium offer. Similar results were found on the pattern of searches for subjects in their roles as player 2.

2. Robot treatment: In the robot treatment, subjects play against a robot in order to switch-off their social preferences, if any. Robot players are programmed to behave selfishly and play the subgame perfect outcome. In their roles as player 1, human subjects, on average, offer \$1.84 to robot opponents (as compared to \$2.11 in the baseline treatment). Thus, concern for others, alters the equilibrium offer by only \$0.27. In their roles as player 2, human subjects frequently reject the robot's opening offer of \$1.25, which is inconsistent with the behavior of equilibrium players. If subjects are trained in backward induction, then they learn quickly to make near-equilibrium offers to robot subjects (\$1.22).
3. Mixture treatment: In this treatment, subjects who had previously been trained in backward induction in the robot treatment, play against untrained human subjects. Trained subjects start by offering relatively low offers that the untrained subjects reject; over time, the trained subjects learn to raise their offers. Conversely, the untrained subjects start by making high offers but learn, over time, that lower offers are also accepted, so they reduce their offers. The average offers in later rounds are the outcome of this "tug of war" between the two kinds of subjects. The average first round offer is \$1.60, which is lower than the robot treatment but higher than the baseline treatment. Hence, the presence of some equilibrium types moves aggregate behavior towards, but not all the way to, the equilibrium prediction. Across all treatments, the departure of actual behavior from equilibrium behavior seems to arise from a combination of cognitive limitations (untrained players) and social preferences (human or robot opponents).

12.6 Asymmetric information, signaling, and cheap talk

In this section we consider signaling games that are more cognitively challenging as compared to the games considered so far. We assume that the reader is familiar with the basics of signaling games.⁸³

12.6.1 *Evidence from a generic signaling game*

The signaling game shown in Figure 12.40 is experimentally tested in Brandts and Holt (1992); it is based on Spence's (1973) job market signaling model. The two panels show the same game but different equilibria (see below). There are two players, a sender (player 1) and a receiver (player 2). Nature moves first at the oval circle in the middle and chooses one of two types for player 1, H or L, with respective probabilities, $2/3$ and $1/3$.⁸⁴ The chosen type is revealed only to player 1 whose pure strategies are given by A and B. Player 2 has two pure strategies, C and D, but does not know the type of player 1. This is indicated by the two vertical dotted lines connecting the

⁸³ For the basics, see the appendix on game theory.

⁸⁴ In the signaling experiments that we report, players have common priors about the probability with which nature initially moves. When such priors are not told to the subjects they might have to learn about them from repeated play of the experiment. The long-term outcomes in the experiments in the two cases of known and unknown initial probability distribution of moves by nature are quite different; see Drouvelis et al. (2012).




Figure 12.40 The signaling game in Brandts and Holt (1992).

Source: Brandts and Holt (1992) with permission from the American Economic Association.

two information sets of player 2; for instance, when player 1 chooses the action A, player 2 does not know whether the game has proceeded to the upper or the lower node on the right hand side information set. The payoffs are shown at each terminal node.

For a detailed consideration of this game and the associated equilibria, the reader can consult the appendix on game theory. There are two equilibria and both are pooling equilibria, shown respectively in the left and right hand side panels of Figure 12.40; the equilibrium strategies are shown by the bold lines. For instance, in the panel on the left hand side, labeled as Intuitive Equilibrium, both types of player 1 choose the action B. On observing action B, player 2 does not update prior beliefs about the type of player 1, and his best reply is the action C. The restriction on “out of equilibrium beliefs” that suffices to support this pooling equilibrium is that if player 2 observes the action A, then player 1 is believed to be of type L; clearly, action C is optimal for player 2.

The second pooling equilibrium is shown on the right hand side of Figure 12.40. It can be ruled out using the Cho and Kreps (1987) equilibrium dominance criterion, hence, it is labeled as the Unintuitive Equilibrium. Refinements that are stronger than the equilibrium dominance criterion, such as divinity and strategic stability, can also be used to rule out the pooling equilibrium on the right. Hence, if the unintuitive equilibrium is not played in experiments, then this result is also consistent with these other theories.

The experiment had three parts, a, b, and c; each part ran over four rounds (or matchings). In *part a*, 8–12 subjects in each two-hour session, played different rounds of the signaling game. Subjects were split into two equal groups of senders and receivers and placed in separate rooms. Subjects chose their actions simultaneously and then pairs were randomly matched in the two rooms. In *part b* of the experiment, subjects played a slightly different game in which an extra action I (standing for irrelevant) is added to the set of actions of the receiver. Action I is dominated by the other two actions (C and D), so it can be pruned to generate the game in Figure 12.40. However, if it is not pruned, then the equilibrium dominance arguments cannot be used. Part b is played over matchings 5–8 of the two sets of players in the two rooms. This is followed by *part c* of the experiment with matchings 9–12, played with the original game (without action I). The results are shown in Table 12.23.

To understand Table 12.23, consider the first row. This is the data from the first set of matches in part a of the experiments in which there are a total of four matches. There were 8 type-H senders, all of whom choose B. Of the 8 type-L senders, only 1 chose B, all others chose A. For the receivers, 9 out of 9 receivers responded to the message B by choosing the action C, and 6 out of 7 responded to the message A with action D. The other rows can be read off analogously.

Table 12.23 Type-contingent messages sent by senders and message-contingent actions of the receivers for the matchings in each of the parts, a, b, and c.

Part/matching	Signals given type		Responses given signals	
	B/H	B/L	C/B	D/A
a/1	8/8	1/8	9/9	6/7
a/2	8/8	2/8	9/10	2/6
a/3	12/12	1/4	13/13	3/3
a/4	12/12	1/4	13/13	3/3
b/5	12/12	–	12/12	–
b/6	4/4	4/8	8/8	4/4
b/7	12/12	–	12/12	–
b/8	8/8	3/4	11/11	1/1
c/9	16/16	–	16/16	–
c/10	12/12	3/4	15/15	0/1
c/11	4/4	9/12	12/13	2/3
c/12	12/12	3/4	15/15	1/1
Percentage for part a:	100	21	98	74
Percentage for part b:	100	58	100	100
Percentage for part c:	100	75	98	60

Source: Brandts and Holt (1992).

Consider, as our theoretical prediction, the intuitive equilibrium shown in the left panel in Figure 12.40. In this equilibrium, type H plays B, which is consistent with the results in Table 12.23. However, type L senders in the experiments mostly choose A in part a of the experiments, which is inconsistent with the pooling equilibrium.

In part a, the response of the receivers to message B is almost always C (and to A it is mainly D). This gives type L senders a payoff of 60 from choosing A and 100 from choosing B. Type L senders seem to learn this lesson as the experiment progresses through various rounds. Whereas the proportion of type L senders who send message B is only 21% in part a, this percentage rises to 58 in part b and 75 in part c of the experiments. The unintuitive equilibrium in which the sender sends message A and the receiver responds with C is played in only 7 decision pairs, mostly in the initial rounds of the experiment when subjects are less experienced. The choices of 112 decision pairs in parts b and c give no evidence for the unintuitive equilibrium.

One may interpret these results as good evidence that subjects play a sequential equilibrium. This comes as a bit of a surprise given the widespread failure of subjects to play equilibrium strategies in less cognitively demanding games (see above). Furthermore, Bayes' rule plays a critical role in signaling equilibria, yet we know from Part 7 that humans do not follow this rule very well. In the initial rounds, the play is not consistent with sequential equilibrium. However, with repeated rounds of the experiment, it does become closer to the equilibrium. Another possibility is that in repeated rounds of experiments, subjects could just be playing simple learning strategies such as adaptive play and reinforcement learning. For instance, in the experimental results of Brandts and Holt (1992), the reason that type-L senders eventually switch to playing strategy A could be that once played, they realize that it has a higher payoff, so they continue to play it. Learning dynamics are examined in a follow-up paper by Brandts and Holt (1993) in which

they show that plausible learning dynamics that converge to an equilibrium might contradict the intuitive criterion.

The intuitive criterion will not work when both players can potentially earn higher payoffs by defecting from the equilibrium action to an out of equilibrium action. A range of other refinements have been given to deal with this case. It is beyond the scope of this book to provide the details, but many excellent sources are available.⁸⁵ The *divinity* and *universal divinity* criteria proposed by Banks and Sobel (1987) are refinements that can be used when the intuitive criterion fails. Although both types of senders might have an incentive to play an out of equilibrium action, the incentive to play such an action may be stronger for one of the types, say type- t . The word “stronger” means here that for a relatively larger range of posterior beliefs of the receiver, it is more attractive for type- t senders to play an out of equilibrium action. Universal divinity requires that receivers should then place a posterior belief of 1 that the sender type is type- t when they observe this out of equilibrium action. When universal divinity does not help in distinguishing between the sender types who could have sent an out of equilibrium message, further theoretical criteria are available. However, empirical support for these is weak; such criteria include Kohlberg and Mertens (1986) *never-a-weak-best-response* (NWBR) criterion and the most ambitious of these refinements, i.e., Kohlberg and Mertens (1986) notion of *stability*.

Banks et al. (1994) test if the refinements of sequential equilibrium have empirical support. They use the following nesting of equilibria: *stable* $\subseteq$ *NWBR* $\subseteq$ *Universally Divine* $\subseteq$ *Divine* $\subseteq$ *Intuitive* $\subseteq$ *Sequential* $\subseteq$ *Bayesian Nash*. Kohlberg and Mertens (1986) show that every game contains at least one stable equilibrium. Thus, the existence of all the various kinds of equilibria in the empirical study is guaranteed. Four subject pools played 7 different games in 13 experiments. The main results are as follows. (1) Refinements that, roughly, require greater cognitive faculties relative to the intuitive equilibrium (i.e., divinity, NWBR and stability) do not do very well. By and large, the play by subjects in experiments is not supportive of these refinements. (2) All of the games, except one, had a pooling equilibrium. But subjects do not play the pooling equilibrium. (3) The authors do not find consistent support for any of the various alternative decision rules that subjects might have been using, such as minimax and the principle of insufficient reasoning.

12.6.2 An application to a limit entry pricing game

Cooper et al. (1997b) conduct experiments using the *limit entry pricing game* of Milgrom and Roberts (1982). In this two-period model, an incumbent monopolist faces a downward sloping demand curve and has private information on the cost. The monopolist can either be a high cost type, M_H , or a low-cost type, M_L ; the prior probability of each type is 0.5. In the first period, the monopolist chooses to produce a quantity, q , at the price, p . The quantity choice is observed by a potential entrant in the market, who then decides either to enter (IN) or stay out (OUT).

In the second period, if the entry does not occur, then, the monopolist produces the monopoly output, q_m , at the price p_m . If entry occurs, then both firms play the Cournot outcome, producing the Cournot quantities at the price p_c . The profit of the incumbent of type M_i , $i = H, L$ over the two periods, denoted by $\Pi_I(M_i)$, is given by

$$\Pi_I(M_i) = \pi_{II}(p) + \delta\pi_{II}(p_j); j = m, c.$$

⁸⁵ The interested reader can consult one of many graduate level game theory texts, e.g., Fudenberg and Tirole (1991). The paper by Banks et al. (1994) has a self-contained and rigorous discussion of the various equilibrium refinements. Camerer (2003, pp. 414–18) gives a more intuitive discussion.

$\delta > 0$ is the discount factor and π_{Ii} is the profit of the incumbent of type M_i , $i = H, L$ in a single period. First period profit, $\pi_{Ii}(p)$, depends on first period price, p . The second period profit, $\pi_{Ii}(p_j)$, $j = m, c$, depends on the second period price p_j that is determined by the entrant's decision, IN (price p_c) or OUT (price p_m). The entrant's payoff, which only materializes in the second period, depends on the type of the incumbent, M_i , and is given by $\pi_E(M_i)$. It is $\pi_E(p_c)$ from the IN decision and some exogenously specified outside option $\bar{\pi}_E \equiv 0$ from the OUT decision.

In their experiments, Cooper et al. (1997b) allow the incumbent monopolist to choose discrete output levels $q = 1, 2, \dots, 7$. The incumbent's payoff over the two periods, $\Pi_I(M_i)$, for each level of first period output and the entrant's decision, is shown in Table 12.24. For instance, when $q = 1$, then $\Pi_I(M_H) = 150$ if the potential entrant enters and $\Pi_I(M_H) = 426$ if the potential entrant stays out. Two treatments are run for the entrant. These are differentiated in Table 12.25 by the background color. In Treatment I (the lighter background), the expected payoff of the entrant from IN is lower than from OUT. By contrast, in Treatment II (the darker background), the expected payoff from IN is relatively higher.

The next Lemma follows immediately from Table 12.25.

Table 12.24 Type-contingent payoffs of the incumbent in the market entry game.

Quantity choice	High cost type, M_H		Low cost type, M_L	
	IN	OUT	IN	OUT
1	150	426	250	542
2	168	444	276	568
3	150	426	330	606
4	132	408	352	628
5	56	182	334	610
6	-188	-38	316	592
7	-292	-126	213	486

Source: David J. Copper, Susan Garvin, and John H. Kagel, "Signalling and adaptive learning in an entry limit pricing game." *The RAND Journal of Economics* 28(4) (Winter, 1997): 662-83. Copyright © 1997, RAND.

Table 12.25 Payoffs of the entrant in the two treatments.

Entrant's choice of action	Type of the incumbent		Expected payoff
	M_H	M_L	
IN	300	74	187
OUT	250	250	250
IN	500	200	350
OUT	250	250	250

Source: David J. Cooper, Susan Garvin, and John H. Kagel, "Signalling and adaptive learning in an entry limit pricing game." *The RAND Journal of Economics* 28(4) (Winter, 1997): 662-83. Copyright © 1997, RAND.

Lemma 12.1 Suppose that the entrant has full information about the incumbent's type. In Treatments I and II, the entrant enters if the incumbent's type is M_H and stays out otherwise.

Proposition 12.14 Suppose that the entrant has full information about the incumbent's type.

- (i) The high cost incumbent, M_H , chooses a quantity $q = 2$ in the first period for an aggregate payoff of 168.
- (ii) The low-cost incumbent, M_L , chooses a quantity $q = 4$ in the first period for an aggregate payoff of 628.

Proof: Using the entry decisions of the entrant, given in Lemma 12.1, and the data in Table 12.24, the highest payoffs for types M_H and M_L , respectively, correspond to $q = 2$ and $q = 4$. ■

The traditional rationale behind limit pricing has been that a high quantity (thus, a low price) chosen by the incumbent may signal to a potential entrant that the profits are low, hence, making entry more costly. We state this formally.

Definition 12.2 (Limit Pricing): A monopolist is said to engage in limit pricing if the type-contingent quantity chosen in the first period, q , exceeds the quantity chosen under full information.

Using Proposition 12.14 and Definition 12.2, type M_H engages in limit pricing if $q > 2$ and type M_L engages in limit pricing if $q > 4$. The next proposition summarizes the main results that are checked against the experimental data.⁸⁶

Proposition 12.15 Consider Treatment I.

- (a) There exist pooling equilibria for output levels 1–5. The out of equilibrium beliefs are that if output differs from the equilibrium level, then the incumbent is of type M_H . In all these pooling equilibria, the entrant enters. If the Cho–Kreps intuitive criterion is invoked, then the only pooling equilibria that survive are output levels 4 and 5. Of these two equilibria, pooling at $q = 4$ satisfies “efficiency” in the sense of Milgrom and Roberts (1982).
- (b) There are two separating equilibria. In both these equilibria, type M_H chooses $q = 2$, and conditional on observing $q = 2$, the entrant enters. In the first separating equilibrium, type M_L chooses $q = 6$, and on observing $q = 6$, the entrant does not enter. In the second separating equilibrium, type M_L chooses $q = 7$, and on observing $q = 7$, the entrant does not enter. These equilibria are supported with out of equilibrium beliefs, which require that if output differs from the equilibrium level, then the incumbent is of type M_H . Of these equilibria, only the first separating equilibrium is consistent with the intuitive criterion.

Proof: (a) Let us consider a particular case of the pooling equilibrium at $q = 3$ (the other cases are analogous). The incentive compatibility condition requires that none of the types find it profitable to choose any $q \neq 3$. In a pooling equilibrium, the prior beliefs of the entrant are not updated, hence, the entrant must make his decision based on his expected payoff (last column of Table 12.25). The expected payoff of the entrant in Treatment I is relatively greater from choosing OUT rather than IN, hence, the entrant does not enter. Thus, from Table 12.24, the payoffs of

⁸⁶ Readers who are rusty with the basics of signaling games may wish to have a quick look at the Appendix on game theory.

types M_H and M_L in the pooling equilibrium are, respectively, 426 and 606. If any of the types of the incumbent decides to choose an out of equilibrium action $q \neq 3$, then the out of equilibrium beliefs require the entrant to believe that the incumbent is of type M_H , so the entrant enters ($\because 300 > 250$). From Table 12.24, there exists no $q \neq 3$ such that either of the incumbent types M_H and M_L get a higher payoff relative to the respective equilibrium payoffs. Each incumbent type gets a strictly positive payoff that is greater than the outside option, which is normalized to zero, hence, the participation constraint is satisfied. The other pooling equilibria can be checked to hold in a similar manner.

The plausibility of the out of equilibrium beliefs can be checked against the intuitive criterion. Continue with the case of the pooling equilibrium at $q = 3$. If type M_L chooses another action, he must not be able to get more than the equilibrium payoff of 606. However, by choosing $q = 4$, there is a possibility that he can get more. Hence, the receiver's belief that for $q \neq 3$ the type of the sender is M_H is not intuitive in the Cho–Kreps sense. Proceeding in this manner one can show that pooling at $q = 4$ and $q = 5$ satisfies the intuitive criterion.

Using Milgrom and Roberts (1982), an equilibrium satisfies “economic efficiency” if there exists no other equilibrium at which one type is strictly better off and no type is strictly worse off. Pooling at $q = 5$ does not satisfy this definition because compared to pooling at $q = 4$, both types are worse off. However, pooling at $q = 4$ satisfies the Milgrom–Roberts efficiency criterion.

(b) Consider now the separating equilibrium in which type M_H chooses $q = 2$ and type M_L chooses $q = 6$. In a separating equilibrium, the actions of the incumbent fully reveal the type to the entrant. Thus, on observing $q = 2$, the entrant believes, correctly, that the incumbent is type M_H and, thus, enters, giving the type M_H incumbent a payoff of 168. On observing $q = 6$ the entrant believes (consistent with Bayesian updating) that the type of the incumbent is M_L , so, using Table 12.25 for Treatment I, the entrant stays out, giving the type M_L incumbent a payoff of 592. The participation constraint holds for each type of incumbent.

Let us check for incentive compatibility, i.e., none of the incumbent types should benefit by mimicking the quantity chosen by the other type. If type M_H chooses $q = 6$, then the entrant believes the incumbent to be type M_L and, thus, stays out. This gives type M_H a payoff of $-38 < 168$. So type M_H does not find it worthwhile to mimic the quantity choice of type M_L . Analogously, if type M_L mimics the quantity chosen by type M_H , it is identified (mistakenly) by the entrant as a type M_H , so the entrant enters. The payoff to type M_L in this case is $276 < 592$. Hence, the proposed separating equilibrium is incentive compatible.

The second separating equilibrium in which type M_H chooses $q = 2$ and M_L chooses $q = 7$ can be checked to hold in an analogous manner. Along the lines of the proof of part (a), it is straightforward to check that the first separating equilibrium satisfies the intuitive criterion, while the second does not. ■

Proposition 12.16 *Consider Treatment II.*

- (1) *No pooling equilibrium exists.*
- (2) *The two separating equilibria identified in Treatment I also exist in this case.*
- (3) *There also exist hybrid equilibria in Treatment II. In particular, the following is a hybrid equilibrium. Type M_L chooses $q = 5$. Type M_H plays a mixed strategy, choosing $q = 5$ with probability 0.20 and $q = 2$ with probability 0.80. The decision of the entrant is as follows. On observing $q \neq 5$ the entrant always enters. But when the entrant observes $q = 5$, he enters with probability 0.11 and stays out with probability 0.89.*

Proof: Similar to the proof of Proposition 12.15, thus, it is omitted. The reader is asked to prove this as an exercise. ■

Cooper et al. (1997b) also experimentally test for other equilibrium refinements. The next proposition summarizes the relevant results on refinements that they test; the interested reader can look up Garvin (1993) for the proof.

Proposition 12.17 *Consider Treatment I. Both pooling equilibria in Proposition 12.15 also satisfy the Banks and Sobel (1987) equilibrium refinement of “divinity.” The “perfect sequential equilibrium” of Grossman and Perry (1986) has a unique outcome, namely, pooling at $q = 4$. The “undefeated equilibrium” of Mailath et al. (1993) predicts pooling at $q = 2, 3$, and 4. Noldeke and Samuelson’s (1994) evolutionary model predicts pooling at $q = 4$ as the unique long-run steady state.*

The experimental results for Treatment I are shown in Figure 12.41 for three sessions. Each session was played by 12–16 subjects allocated to the roles of the incumbent and the entrant over




Figure 12.41 Results for Treatment I.

Source: David J. Cooper, Susan Garvin, and John H. Kagel, “Signalling and adaptive learning in an entry limit pricing game.” *The RAND Journal of Economics* 28(4) (Winter, 1997): 662–83. Copyright © 1997, RAND.

three sets of 12 rounds, each set being called a cycle. Subjects are inexperienced in Sessions 1 and 2. Session 3 is conducted with experienced subjects who have prior experience of Treatment I. Each of the histograms shows the percentage of incumbents choosing various levels of quantities. The percentage of entrants who choose to enter is shown as a row of numbers below the histogram beginning with the prefix IN (%).

Inexperienced subjects initially play the full information (or *myopic maximum*) equilibrium. In sessions 1 and 2, in cycle 1, 69% of the choices made by type M_L incumbents coincide with $q = 2$ and 77% of choices made by type M_H incumbents coincide with $q = 4$. As subjects gain experience, in session 3, both types of incumbent subjects gravitate towards the pooling equilibrium in which they choose $q = 4$. However, there is still a substantial fraction of subjects who do not play the pooling equilibrium. Recall from Proposition 12.15 that $q = 4$ corresponds to Milgrom and Robert's (1982) notion of economic efficiency. In cycle 3 of session 3 (constituting the most experienced subjects), the quantity choice of $q = 4$ is made by about 70% of incumbents of type M_H and 92% of the incumbents of type M_L . However, at this pooling equilibrium, only about 6% of the entrants choose IN, the expected equilibrium choice (see Proposition 12.15a).

The separating equilibrium (Proposition 12.15b) is not observed and neither is the pooling equilibrium at $q = 5$ (Proposition 12.15a). Cooper et al. (1997b, p. 670) conclude that: "Indeed, for all of our experiments it is striking how slowly anything resembling a full-blown equilibrium emerges (typically only late in experienced-subject sessions). Even in the best cases, there is a substantial minority of choices (5–10%) that deviate from any given equilibrium."

The experimental results from Treatment II are shown in Figure 12.42. As in Treatment I, in the early periods, each of the incumbent types disproportionately play their full information strategies. Data in the first six periods of play in Treatments I and II are statistically indistinguishable. There are no pooling equilibria in Treatment II (Proposition 12.16). However, there is some tendency for incumbents of type M_H to try to pool with incumbents of type M_L , prior to the emergence of a separating equilibrium. In session 5, the modal choice for all types of incumbents is $q = 4$. By cycle 3 in session 6 (in which subjects have the greatest experience), about half of the incumbents of type M_H and about 80% of incumbents of type M_L play the separating equilibrium (respectively, $q = 2$ and $q = 6$). In response to $q = 2$, 100% of entrants enter, while in response to $q = 6$, about 12% of entrants enter, which is consistent with equilibrium predictions.

The authors explain the data on individual choices by using an adaptive learning model of fictitious play. It does take many rounds of repeated play to converge to equilibrium. Cooper et al. (1997a) consider Treatment II only. They show that the intuitive criterion is clearly violated. An adaptive learning model explains when the intuitive criterion is likely to be satisfied or violated.

The results on equilibrium refinements (Proposition 12.17) are not positive. Cooper et al. (1997a, p. 555) write: "Game theory is at its foundation a hypothesis about how people behave. In particular, it posits that individuals will attempt to anticipate the behavior of others and respond accordingly. This is the soul of game theory, and the experiments indicate that it is alive and well. What may not be so healthy are the legion of assumptions which have been added to the theory in order to get tractable results. In the rush to get theories which give sensible outcomes and are easily employed by theorists, the reality of how people actually behave may have been left behind. We do not suggest that game theory be abandoned, but rather as a descriptive model that it needs to incorporate more fully how people actually behave."

12.6.3 *Finitely repeated games of asymmetric information*

Consider finitely repeated games of one-sided asymmetric information. Suppose, for simplicity, that a player has private information about his two possible types, *strong* and *weak*, in a sense to




Figure 12.42 Results for Treatment II.

Source: David J. Cooper, Susan Garvin, and John H. Kagel, "Signalling and adaptive learning in an entry limit pricing game." *The RAND Journal of Economics* 28(4) (Winter, 1997): 662–83. Copyright © 1997, RAND.

be made precise below. Since the game is finitely repeated, the solution can be found by backward induction. Strategies for each type of each player specify actions for all possible histories of the game. One also needs to specify beliefs of a player on observing out of equilibrium actions.

The equilibrium for such games, due to Kreps and Wilson (1982) and Milgrom and Roberts (1982) (henceforth, KWMR), has generated much interest in economics. The weak type successfully mimics the strong type, initially, by engaging in a costly signal, building up a reputation for being strong. In later periods, the weak type milks its reputation that it has built up in previous periods and reduces the cost of sending a costly signal by randomizing between the costly and a less costly signal.

Consider, for instance, the model of *entry deterrence* that is examined experimentally in Jung et al. (1994). Eight rounds of an entry deterrence game are played by a long-lived incumbent monopolist and a series of short-lived entrants, one for each period. An incumbent monopolist

Table 12.26 Payoff matrix in an entry deterrence game.

Entrant's decision	Monopolist's decision	Payoffs		
		Entrant	Weak monopolist	Strong monopolist
In	Fight	80	70	160
	Accommodate	150	160	70
Out	No choice	95	300	300

Source: Jung et al. (1994).

Table 12.27 Predictions of the equilibrium actions of the weak monopolist and the entrant.

Period	Weak monopolist	Entrant's probability of entering	
	Probability of accommodating	Conditional on entry and fight in previous period	Conditional on no previous entry
1	0.000	0.000	0.000
2	0.000	0.000	0.000
3	0.000	0.000	0.000
4	0.188	0.000	0.000
5	0.346	0.357	1.000
6	0.416	0.357	1.000
7	0.560	0.357	1.000
8	1.000	0.357	1.000

has private information about his two possible types, *weak* (high marginal cost) and *strong* (low marginal cost). The payoffs are shown in Table 12.26. In a one-shot game, a weak monopolist gets a higher payoff from accommodating entry and a strong monopolist gets a higher payoff from fighting entry (in each case because $160 > 70$). Relative to the decision to stay out and get a payoff of 95, it only pays the entrant to enter if the incumbent accommodates ($150 > 95$). Otherwise if the incumbent fights, it is best to stay out ($80 < 95$). The entrant does not know the type of the incumbent but assigns a commonly known prior probability of 0.67 that the type is weak and 0.33 that the type is strong. The predictions are given in Proposition 12.18; one of the exercises asks you to prove it.

Proposition 12.18 *Consider the eight period chain store model of Jung et al. (1994). The predictions of the model, conditional on no accommodation having taken place in any previous round, are given in Table 12.27 (the strong monopolist always prefers to fight, so Table 12.27 only reports the actions of the weak monopolist and the actions of the entrant, conditional on the history of play).*

The weak monopolist fights entry, if it takes place in any of the first three periods; this action prevents the entrant's entry for a while. Entrants play a mixed strategy, and in equilibrium, begin to enter in period 4. From period 4 onwards, the weak monopolist plays a mixed strategy in

which the probability of accommodation increases gradually until it reaches 1 in the last round. The entrant's entry decision is history dependent. Following periods in which there has been entry and fights, the entrants enter with a stationary probability of 0.357. If no previous entry has taken place, then the incumbent monopolist has not had a chance to build his reputation, so entry takes place with probability 1 in periods 5–8. If the weak monopolist does not fight in the first three rounds of play, then the out of equilibrium beliefs of the entrant are that he is a weak type, in which case, all subsequent entrants enter (for the proof of Proposition 12.18, you will need to use these out of equilibrium beliefs).

From Proposition 12.18, standard theory predicts that monopolists will engage in *predatory pricing* (pricing and other practices that deter entry and reduce welfare). The Chicago school has traditionally expressed skepticism about the possibility of predatory pricing. The main source of the skepticism is that the incumbent monopolist cannot forever make losses, thus, entry cannot be stopped in the long run. Unravelling backwards, it should not pay to deter entry in the short run. The possibility of asymmetric information, however, alters the results as shown in Proposition 12.18. In practice, evidence of predatory pricing has been hard to find. Thus, the experimental results of Jung et al. (1994) are important.

Subjects in experiments played between 45–90 rounds of the game. The observed play of the subjects up to round 31 is not very supportive of the equilibrium predictions. For instance, in the initial rounds, weak monopolists accommodate about half the times in the first three periods, although theory predicts that they should be fighting all the time. From rounds 31 onwards, as subjects gain experience with the game, accommodation in the first three rounds was done in only 15% of the plays, while for very experienced subjects it was almost eliminated in these rounds. From round 31 onwards, predatory pricing was found in 80% of the rounds.

There are, however, important departures from the Kreps–Wilson predictions in two important respects. First, the entry rates are not stationary, as predicted. Entry rates increase steadily over the various experimental rounds. Second, while the predicted numbers in the last column of Table 12.27 are higher than those in the penultimate column, the experimental findings are just the opposite. In an alternative approach, Selten (1978) argued that the backward induction argument is more compelling in the last few rounds of the game. So the monopolist might accept the backward induction argument for the last few rounds, while choosing to be aggressive in the initial periods. While this theory makes less precise predictions, it can be a useful device to make sense of the data arising from experiments.

While the data in Jung et al. (1994) violates important aspects of the KWMR predictions, the findings of Camerer and Weigelt (1988) are more supportive of these predictions. Camerer and Weigelt (1988) consider a borrower–lender game in which there is asymmetric information about the type of the borrower. The borrower can be one of two possible types: *normal* (who prefers to default on a loan) and *nice* (who prefers to repay the loan). The structure of the game is similar to the one in Jung et al. (1994) and is also played over eight rounds. The predictions of the model are similar to those in Proposition 12.18 and can be derived using backward induction.

Using the notion of *homemade priors*,⁸⁷ Camerer and Weigelt (1988) conclude that their data broadly supports the Kreps–Wilson predictions. Jung et al. (1997, p. 91) write that: “We give the Kreps and Wilson sequential equilibrium model lower marks than Camerer and Weigelt (1988). . . . Reexamining their data, there are some tendencies to deviate from the Kreps and Wilson sequential equilibrium along the lines found here. In particular, the bankers in the

⁸⁷ In these priors, the lenders think that some proportion of the normal type borrowers will always pay back the loans (i.e., behave just like the nice types).

borrower–lender game, who correspond to the entrants in the chain-store game, lent less as the last period approached, rather than following the stationary strategies required under mixed-strategy play.” It seems that the predictions of the Kreps and Wilson equilibrium have a greater chance of holding under the parameter values employed in Camerer and Weigelt’s borrower–lender game relative to the chainstore game in Jung et al. (1994). Neral and Ochs (1992) show that the results of Camerer and Weigelt are not very robust to changes in parameter values, particularly those related to their homemade priors.

12.6.4 *Signaling in corporate finance*

In actual practice, individual investors in a firm are likely to be less informed about the profitability of the firm as compared to the managers of the firm. Managers of firms with good prospects might then send a costly signal on profitability in order to separate themselves from firms with poorer prospects (Myers and Majluf, 1984).

Suppose that we are interested in explaining excessive levels of debt to equity ratios (leverage) of companies and managers are more informed about the firm’s future profit prospects as compared to investors. We may, heuristically, construct a separating equilibrium in which managers signal their private information by using debt as an instrument, as follows. Only a firm with good prospects may hope to repay a higher level of debt in the future. Hence, a firm with good prospects can, by deciding on a high leverage, separate itself from a type of firm that has bad prospects. Ex-post, when investors observe high leverage, they become informed of the firm’s prospects.

Dividends are sometimes not easy to explain because they just represent a movement from one account (the investors’ share in the firm) to another (investors’ personal bank account), which is taxed. However, under asymmetric information, firms with better prospects may offer greater dividends to signal to investors that current cash flows are healthy.

These sorts of issues are examined by Cadsby et al. (1990, 1998). Consider a firm and its potential investors. The firm has two possible types, indexed by $j = H, L$; types are private information. The firm has initial level of assets, $A_j > 0$, and a project that has a net present value, $B_j > 0$. The cost of the project, $c > 0$, is identical for both types. A type H firm has a relatively higher initial level of assets, so $A_H > A_L$, and a relatively more profitable project, so $B_H > B_L$. We consider Cadsby et al. (1990) first, where signaling is costless. The firm has one of two possible actions. If the firm invests in the project, the end of period value of the firm is $V_j = A_j + B_j - c$, and if it chooses to forgo the project, the end of period value of the firm is $V_j = A_j$.

There is a large pool of risk-neutral investors who are potentially interested in investing in the project. Investors do not know the type of the firm but have prior beliefs that each type is equally likely. Denote by EV_j , the expected value of the firm and by s_j , the share in equity of a type $j = H, L$ firm, that is demanded by an investor in return for the investment. The firm moves first and chooses one of its type-contingent actions, forgo the project, or invest in the project. If the firm forgoes, then the game ends and the firm gets $V_j = A_j$, while investors get nothing. If the firm chooses to invest, it needs to raise the money to pay for the cost c of investing. In this case, a descending value first price auction is conducted among the investors to determine the sole winner who finances the cost of the project, c . Since there is a large pool of risk-neutral and competitive investors, competitive bidding reduces the profits of the investors to zero. Thus, the winning bid, s_j^* , satisfies

$$s_j^* EV_j = c, \quad (12.25)$$

Table 12.28 Type-contingent assets and end of period value of firms.

Experiment	Type L		Type H	
	A_L (Forgo)	V_L (Invest)	A_H (Forgo)	V_H (Invest)
A	250	625	500	1250
B	160	500	620	1000
C	300	750	625	1250
D	80	400	270	600
E	50	375	200	625
F	400	750	800	1250
G	150	480	450	810
H	150	480	450	1050
I	375	750	1875	2250
J	240	600	880	1400

Source: Cadsby et al. (1990).

where the expectation is taken over the type of the firm and the value of the firm is $V_j = A_j + B_j - c$, $j = H, L$. Since the firm knows its type, its profits equal $(1 - s_j^*) V_j$, while the investors get $s_j^* V_j$. A firm of type $j = H, L$ invests if $(1 - s_j^*) V_j \geq A_j$, otherwise it forgoes the investment.

Depending on the parameter values, separating, pooling, and hybrid equilibria can arise in this model. Consider one choice of type-contingent parameter values for A_j , V_j shown in Table 12.28. The first column alphabetically lists the ten sets of experiments, A through J. The first two experiments, A, B, with shaded backgrounds, are the full information benchmarks in which investors know the type of the firm. In all cases, $c = 300$.

The different equilibria that arise for each of the cases in Table 12.28 are given in Table 12.29. The predicted equilibria are listed in the second column and are of the following types: Full information (FI), separating (S), pooling (P), and hybrid or semi-separating (H). The firm has two possible investment actions, forgo or invest. The column labeled Pr(Invest) gives the probability of investment (possibly a mixed strategy). The predicted share of the firm demanded by the investor in equilibrium, s_j^* , for each of the types, $j = H, L$, is also shown.

In a pooling equilibrium, both types of firms either forgo the investment or invest in the project. Investors do not update their beliefs in a pooling equilibrium, thus, using (12.25), $s_L^* = s_H^*$. In a separating equilibrium, one type invests, while the other does not, thus, $s_L^* \neq s_H^*$. For some parameter values, there is a unique equilibrium, while for others, such as those in the experiments E, F, H, and J, there are multiple equilibria.

In the full information equilibrium, in experiments A, B, the type of the firm is known to the investors, hence, using (12.25), $s_j^* = c/V_j$. Using $c = 300$, and the information given in Table 12.28, one can readily compute s_L^*, s_H^* for the first two rows in Table 12.29. For these values of s_j^* , and the values given in Table 12.28, a simple calculation shows that $(1 - s_j^*) V_j > A_j$, $j = H, L$, so both types of firms invest with probability 1 (Pr(Invest) = 1).

Now consider the pooling equilibrium in experiment C. In a pooling equilibrium, prior beliefs are not updated, hence, using (12.25) and Table 12.28, $s_j^* = c/EV_j = \frac{300}{0.5(750) + 0.5(1250)} = 0.3$, $j = H, L$. A firm of type L undertakes the project with probability 1 because $(1 - s_L^*) V_L = 0.7(750) > A_L = 300$. Similarly, a firm of type H invests because $0.7(1250) > 625$.

Table 12.29 Equilibrium predictions of the signaling model in Cadsby et al. (1990).

Experiment	Equilibrium	Type L firm		Type H firm	
		Pr(Invest)	S_L^*	Pr(Invest)	S_H^*
A	FI	1	.480	1	.240
B	FI	1	.600	1	.300
C	P	1	.300	1	.300
D	S	1	.750	0	.000
E	S	1	.800	0	.000
F	P	1	.600	1	.600
	H	1	.680	.36	.680
	S	1	.400	0	.000
	P	1	.300	1	.300
G	H	1	.360	.20	.360
	S	1	.625	0	.000
	P	1	.392	1	.392
H	S	1	.625	0	.000
	H	1	.571	.09	.571
I	S	1	.400	0	.000
J	S	1	.500	0	.000
	P	1	.300	1	.300
	H	1	.371	.35	.371

Source: Cadsby et al. (1990).

Consider experiment E in Table 12.28. From Table 12.29, we know that in this case there are three possible equilibria, S, P, and H. Consider, for instance, the separating equilibrium. Let the prior beliefs of investors be such that if the firm invests in the project, it is type L, otherwise, if it forgoes, it is type H. Suppose that a firm of type L chooses to invest and a type H firm chooses to forgo. These choices reveal the types of the firm. Hence, using (12.25), for a type L firm, investors are willing to accept a share in the equity equal to $s_L^* = c/V_L = 300/375 = 0.8$. Hence, the profit of a firm of type L is $(1 - s_L^*)V_L = 0.2(375) = 75.0 > 50$, thus, type L prefers to invest rather than forgo. If firm H were to try to mimic a type L firm, then its profits are $(1 - s_L^*)V_H = 0.2(625) = 125.0 < 200$, hence, incentive compatibility holds, giving rise to a separating equilibrium.⁸⁸

The reader can proceed analogously to construct all other entries in Table 12.29.

We show the results of the experiments in Table 12.30. In the full information cases A and B, there was quick convergence to the predicted equilibrium, hence, the results for these cases are

⁸⁸ By contrast, for experiment E, there is no separating equilibrium in which type H invests and type L forgoes. To see this, let the beliefs of investors be that if the firm invests in the project, it is type H, otherwise, if it forgoes, it is type L. Suppose that a firm of type H choose to invest and a type L firm chooses to forgo. These choices reveal the types of the firm. Hence, using (12.25), for a type H firm, investors are willing to pay a share in the equity equal to $s_H^* = c/V_H = 300/625 = 0.48$. Hence, the profit of a firm of type H is $(1 - s_H^*)V_H = 0.52(625) = 325 > 200$, hence, type H prefers to invest rather than forgo. If firm L were to try to mimic a type H firm, then its profits are $(1 - s_H^*)V_L = 0.52(375) = 195.0 > 50$, so the putative separating equilibrium breaks down.

omitted. For cases C through J, the mean level of s_L and the predicted levels s_L^* are quite close. The other main results can be summarized as follows.

1. When there is a unique predicted equilibrium for an experiment, convergence to that equilibrium takes place in even fewer number of rounds. Convergence was almost instantaneous when the unique equilibrium was pooling (as in C). Convergence took longer when there was a unique separating equilibrium (as in cases G, I).
2. When there is a unique separating equilibrium (as in cases G, I), subjects initially try to play some pooling strategies. However, as they discover in successive rounds that separating is more profitable, actual play converges to a separating equilibrium (in round 6 for G and in round 4 for I).
3. When there are multiple equilibria (as in cases E, F, H), the experimental outcomes converge to the pooling equilibrium quite rapidly (round 2 in E and almost instantaneously in F, H). This occurred irrespective of whether subjects had prior experience of pooling or separating equilibrium.⁸⁹

These results provide good support for the theory. The authors do not report data on “individual” levels of bids, s_i . Strictly speaking, the predictions of theory only apply to these individual decisions, however, the reported results are for the mean levels (see Table 12.30). On the whole, it is surprising that in the light of people failing to match the theoretical predictions in much simpler games, they actually do so remarkably well in these quite complex games. Furthermore, people often violate Bayes’ rule in experiments (see Part 7), which plays a critical role in a Perfect Bayesian equilibrium of signaling games. For these reasons, replicating these experiments, checking for the robustness of the results, looking at data at the disaggregated level, and having a richer action space for firms, is important.

Cadsby et al. (1998) extend their earlier work to introduce the possibility of costly signaling (e.g., advertising). Each of the types of the firm can buy a costly signal. Recall that the profit of

Table 12.30 Actual and predicted results of the signaling model in Cadsby et al. (1990).

Experiment	Type L		Type H		Mean s_L from the data	s_L^* predicted by the model
	Number of subjects	Number of investors	Number of subjects	Number of investors		
C	6	6	6	6	0.283	0.300
D	8	8	4	0	0.575	0.750
E	4	4	8	8	0.583	0.600
F	5	5	7	7	0.293	0.300
G	6	6	6	0	0.617	0.625
H	6	6	6	6	0.378	0.392
I	4	4	8	0	0.385	0.400
J	5	5	7	7	0.288	0.300

Source: Cadsby et al. (1990).

⁸⁹ The authors offer one possible conjecture for the greater prevalence and salience of the pooling equilibrium. If any of the investors is aggressive enough to anticipate pooling, he will drive the auction price (s_L) down, creating an incentive for firms to pool.

each firm is given by $(1 - s_i^*) V_i$. On the one hand, buying a costly signal reduces the value of the firm, V_i . However, because the costly signal might reveal the type of the firm to the investors, a type H firm might benefit on account of investors bidding a lower share s_H^* . On balance, it might be optimal for a firm to send out a costly signal.

The main aim of the paper is to test for the Cho–Kreps equilibrium dominance criterion (or the intuitive criterion) but also to test if in this more complicated environment signaling equilibria still arise. The results are more negative for the theory in this paper relative to their earlier paper. It is no longer the case that the signaling outcome was obtained for all experiments. Even when subjects reached an equilibrium prediction, the equilibrium was not maintained, i.e., play drifted away from the equilibrium in subsequent periods. The evidence for the intuitive criterion is mixed. Indeed, the data cannot be explained by equilibrium refinements at all but by the myopic short-run gains to type H from revealing its type. In summary, classical signaling equilibrium analysis does reasonably well when signals are not costly, but less well when signals are costly.

12.6.5 *Leader contributions to charitable giving: signaling or reciprocity?*

Evidence suggests that initial public contributions by a relatively big donor (the leader), positively influence subsequent contributions from smaller donors (the followers) (List and Lucking-Reiley, 2002). In a related context, church contributions are found to be higher when contributions are collected by the open basket method (public information) rather than by the closed basket method (Soetevent, 2005). Why is this so? One possibility is that the value of the donations might be private information to big donors, so their contributions may signal to others that such donations are valuable, hence, inducing greater contributions from others.⁹⁰ An alternative explanation is that people may have “other-regarding preferences,” thus, followers reciprocate the contributions of the leader. In this view, a high initial contribution from a leader induces positive reciprocity from the followers.

To fix ideas, consider, as in Potters et al. (2007), a two-player game with a leader (L) and a follower (F). Both players are risk neutral. Contributions to charitable giving, x_i , by each player $i = L, F$, can take two possible values: $x_i = 1$ or $x_i = 0$. The payoffs of the players are given by

$$\pi_i = 1 - x_i + r(x_L + x_F); \quad i = L, F, \quad r > 0, \quad (12.26)$$

where $r \geq 0$, is the per unit return to a player from the total charitable contributions and $r \in \{0, 0.75, 1.5\}$. Under full information, the follower knows the return r . However, under asymmetric information, the follower only knows that the three possible values of r are equiprobable.

The social maximum can be found by choosing x_L, x_F to maximize

$$\pi_L + \pi_F = 2 + (x_L + x_F)(2r - 1). \quad (12.27)$$

Proposition 12.19 *If $r = 0.75$ or 1.5 , then it is socially optimal for each player to contribute the maximum amount, $x_i = 1$. Otherwise, if $r = 0$, it is socially optimal to choose zero contributions, $x_i = 0$.*

⁹⁰ For theoretical models, see Hermlin (1998) and Vesterlund (2003).

Proof: Since the individual choices are binary, $x_i = 0$ or $x_i = 1$, $i = L, F$, we compare the optimum at the two possible values. If $2r - 1 > 0$, or equivalently $r > 0.5$, then, using (12.27), it is optimal to have $x_i = 1$, $i = L, F$. Otherwise, if $2r - 1 \leq 0$, or equivalently $r \leq 0.5$, then the optimal choice is $x_i = 0$, $i = L, F$. ■

Having established the social benchmark, now consider the full information case under privately optimal choices. Let both players simultaneously, or sequentially, decide to make their contribution decision; results are identical for these two cases under full information. In a Nash equilibrium, each player maximizes the objective function, given in (12.26), conditional on the contribution of the other player. Rewriting (12.26) we get,

$$\pi_i(x_i) = 1 + (r - 1)x_i + rx_j; j \neq i, j = L, F. \quad (12.28)$$

We assume that, holding fixed the contributions of player j , player i contributes if $\pi_i(x_i = 1) \geq \pi_i(x_i = 0)$. From (12.28), irrespective of the contribution of the other player, it pays a player to contribute only if $r \geq 1$. Hence, the following result is immediate.

Proposition 12.20 *Under full information, whether players contribute simultaneously or sequentially, it is optimal to contribute if $r = 1.5$ and to not contribute if $r = 0$ or $r = 0.75$.*

Under asymmetric information, nature first draws a value of r , which is observed by the leader but not by the follower; the follower only knows that each value of r is equiprobable. The leader chooses the level of contribution, x_L , which is observed by the follower who then chooses the level of contribution x_F .

Proposition 12.21 *In the unique sequential equilibrium, the leader contributes only when it is socially optimal to do so (i.e., when $r = 0.75, 1.5$). The follower makes the same contribution as the leader. The beliefs of the follower are that if a strictly positive contribution by the leader is observed, then $r = 1.5$, otherwise, $r = 0$.⁹¹*

Proof: We need to check that the stated strategies and beliefs are sequentially rational. First let us check that the strategy of the leader is optimal, given the subsequent strategies along the equilibrium path. Thus, let the leader believe that the follower will match his contribution, i.e., $x_L = x_F$. Then, using (12.28), the payoff of the leader can be written as $\pi_L = 1 + x_L(2r - 1)$. Since $x_L \in \{0, 1\}$, the leader makes a contribution when $2r - 1 \geq 0 \Leftrightarrow r \geq 0.5$, which occurs in the two cases, $r = 0.75, 1.5$. Let $r = 0.75$ or 1.5 . Given the equilibrium strategy of the leader ($x_L = 1$), and the follower's beliefs that require $r = 1.5$, the payoff of the follower, using (12.28), is given by

$$\pi_F = 1 + 1.5 + x_F(1.5 - 1),$$

which is maximized at $x_F = 1$. We leave it for the reader to check that if $x_L = 0$, then $x_F = 0$ is optimal. ■

So far, we have assumed that players care only about their own payoffs. What if players instead have social preferences? In order to incorporate social preferences, Potters et al. (2007) use the following form of Fehr and Schmidt (1999) preferences.

⁹¹ Bayes' rule in these games allows us to put any restrictions on out of equilibrium beliefs; see the game theory Appendix for this claim and for the basics of a sequential equilibrium.

$$U_i(x_i) = \pi_i - \beta_i \max\{\pi_i - \pi_j, 0\}, i = L, F, \quad (12.29)$$

where π_i, π_j are defined in (12.28). In this formulation, players care about their own payoffs as well as about *advantageous inequity*; $\beta_i > 0$ is known as the parameter of advantageous inequity.⁹² Only the full information case is considered under social preferences, so both players know r before they choose their contributions. Under social preferences, holding fixed the contribution of player j , player i contributes if $U_i(x_i = 1) \geq U_i(x_i = 0)$. We shall be particularly interested in the case $r = 0.75$.

Proposition 12.22

- (1) Consider first a simultaneous move game between the leader and follower when they have social preferences. The following are some interesting Nash equilibria.
 - (1a) Both players contribute if $1 - r \leq \beta_L$ and $1 - r \leq \beta_F$. In the special case of $r = 0.75$, this reduces to the condition $0.25 \leq \beta_L$ and $0.25 \leq \beta_F$.
 - (1b) If $r < 1$, then there is also a Nash equilibrium in which both players do not contribute. In particular, this also holds true for the case $r = 0.75$.
- (2) Now consider a sequential game in which the leader chooses x_L first. The follower observes the leader's contribution and then decides on his/her contribution, x_F . In this sequential game, the equilibrium is identical to that in part (1).

Proof: (1a) Suppose that we are interested in finding a Nash equilibrium in donations in which the leader and the follower simultaneously choose their respective contributions. Using (12.26), we get

$$\pi_i - \pi_j = x_j - x_i. \quad (12.30)$$

Suppose that the leader contributes, i.e., $x_L = 1$. Then the utility of the follower from contributing $x_F = 1$ is $U_F = 2r$ (using $x_L = x_F = 1$, (12.28), (12.29), (12.30)). On the other hand, if the follower does not contribute ($x_F = 0$), then from (12.29), (12.30), the follower's utility is $U_F = 1 + r - \beta_F$. Therefore, the follower prefers to match the contribution of the leader if $1 + r - \beta_F \leq 2r$, or, equivalently, if $1 - r \leq \beta_F$. Since $r = 0.75$, the sufficient condition is $0.25 \leq \beta_F$. Given the symmetry of the game, one can check analogously that if the follower contributes, then the leader also contributes if $1 - r \leq \beta_L$.

(1b) Suppose that the leader does not contribute, i.e., $x_L = 0$. The utility of the follower from the matching contribution, $x_F = 0$, is $U_F = 1$. From the non-matching contribution, $x_F = 1$, the utility of the follower is $U_F = 1 + (r - 1)$. Thus, it is optimal not to contribute if $r < 1$. Analogous reasoning applies to the leader if the follower does not contribute. Hence, for $r < 1$, $x_L = 0$ and $x_F = 0$ is also a Nash equilibrium.

(2) The proof follows by backward induction. Conditional on any level of the contributions by the leader, x_L , the payoff of the follower, given in (12.26), is

$$\pi_F(x_F = 1) = r(1 + x_L), \pi_F(x_F = 0) = 1 + rx_L. \quad (12.31)$$

First, let us find an equilibrium in which, having observed $x_L = 1$, the follower also finds it optimal to contribute $x_F = 1$. Using (12.26), (12.29), (12.31), we get $U_F(x_F = 1) = 2r$ and

⁹² Disadvantageous inequity, which is an essential part of Fehr-Schmidt preferences, is ignored here; see Part 2 for more details.

$U_F(x_F = 0) = 1 + r - \beta_F$. Hence, it is optimal for the follower to contribute if $1 - r \leq \beta_F$. Similarly, if the leader does not contribute, $x_L = 0$, then the utility of the follower from contribution and no-contribution is, respectively, $U_F(x_F = 1) = r$ and $U_F(x_F = 0) = 1$. Thus, it is optimal not to contribute if $r \leq 1$. ■

The experiments were run with 192 subjects over four sessions and there were 12 students in a session. Within each session, six pairs of students made decisions over 18 rounds.⁹³ Two kinds of treatments were run, full information and asymmetric information. Within each treatment, simultaneous and sequential move cases are separately run. The results are as follows.

1. The results for leader and follower contribution rates under asymmetric information are shown in Figure 12.43. The total contributions are 50% larger in the sequential move case, relative to the simultaneous move case. It is only in the sequential move case that the leader gets a chance to signal the value of r through his level of contributions. Hence, these results are consistent with the signaling explanation. In the sequential move game, a binomial test reveals that followers are more likely to contribute when leaders contribute. The followers do not exactly match the leader contributions, but there is a high enough match in the contributions to support the predicted equilibrium. When the leader contributes, the follower contributes 80.6% of the time and when the leader does not contribute, the follower also does not contribute 92.2% of the time.
2. Under asymmetric information, the behavior of individual subjects is rational in the following sense. For both simultaneous and sequential treatments, the leaders almost never contribute when $r = 0$ and almost always contribute when $r = 1.5$ (see Proposition 12.21). The most significant difference arises for the case $r = 0.75$. The leadership contributions




Figure 12.43 Contribution rates under asymmetric information.

Source: Potters et al. (2007).

⁹³ Subjects were not shown the functional forms for π and U . Rather, for two abstractly named decisions A and B, corresponding to $x = 0$ and $x = 1$, they were given probabilistic payoffs corresponding to different values of r .




Figure 12.44 Contribution rates under full information.

Source: Potters et al. (2007).

are 15% in the case of simultaneous moves, jumping up to 75% under sequential moves. The contributions of followers also increase in this case.

3. Consider now the case under full information shown in Figure 12.44. In this case, there are no signaling concerns. The classical predictions are given in Proposition 12.20, while the predictions under reciprocity are given in Proposition 12.22. In particular, in a sequential move game, if the follower's reciprocity is high enough (high β_F), then he is more likely to contribute, which also increases the leader's incentive to contribute. This is the main insight that is tested in the empirical results. The contribution rates are quite low when $r = 0$. However, when $r = 0.75$, a quarter of the population contributes, whereas in the absence of reciprocity, nobody should contribute. There is evidence of reciprocal behavior, which is suggestive of social preferences. When the leader contributes (in the sequential case), the follower also contributes in 33.3% of the cases; in the absence of reciprocity, there should be no such effect. That the followers' choices are influenced by those of the leaders is statistically verified using a binomial test. In the sequential case, the leader contributions (27%) are statistically larger than those in the simultaneous move case (16%). However, follower contributions are lower in the sequential case relative to the simultaneous move case. When we take account of the leader and follower contributions, the total contributions in the sequential and the simultaneous cases are not statistically different.

The main difference between the asymmetric and full information cases is this. When reciprocity seems to be present, it does not alter the total amount of contributions under full information when we compare the two cases of simultaneous and sequential contributions. However, under asymmetric information, there are significantly higher contributions in the sequential case. This lends credence to the signaling hypothesis, namely that leader contributions give rise to an increase in overall contributions when the leader has asymmetric information about the returns from contributing.

12.6.6 Multiple signals

Duffy and Feltovich (2002) consider the effect of multiple signals on behavior. They use the term “multiple signals” in the dual sense of “signals” and “messages/actions.” Suppose that a sender and a receiver interact repeatedly, playing a stage game over several rounds. The following three well-known games are considered.

Prisoner's Dilemma (PD)			Stag-Hunt (SH)			Chicken (CH)		
	C	D		C	D		C	D
C	70, 70	10, 80	C	70, 70	10, 55	C	70, 70	50, 80
D	80, 10	40, 40	D	55, 10	55, 55	D	80, 50	40, 40

Payoffs are expressed in points. Each point in the payoff matrix represents a 1% chance of winning \$1. In each of the three games, C stands for cooperation and D for defection. The cooperative outcome is (C, C), while the two coordinated outcomes, or Nash equilibria, are (C, C) and (D, D). In the prisoner's dilemma game (PD) there is a unique Nash equilibrium, (D, D). In the stag-hunt game (SH), there are two pure strategy Nash equilibria, (C, C), (D, D), and a mixed strategy Nash equilibrium in which each player plays C and D with respective probabilities 0.75 and 0.25. In the game of Chicken (CH), there are two pure strategy Nash equilibria, (C, D), (D, C), and a mixed strategy Nash equilibrium in which each player plays C and D with probability 0.5 each.

Three possible treatments are considered. In the *cheap talk only treatment*, the sender can send costless messages to the receiver. In the *observations only treatment*, players can observe the actions of others in the previous round. In the *control treatment*, neither observation of past actions nor cheap talk is allowed. Duffy and Feltovich (2006) introduce two further treatments. In the *words and deeds (WD) treatment*, senders send a cheap talk message to the receivers as an intention of the action that they might play. Receivers also observe the action of the sender in the previous round. The *words, deeds, and lies (WDL) treatment*, is similar except that, in addition, receivers also observe the message/action of the sender in the previous round. Thus, receivers can, by comparing the previous round action and message, check if the sender was lying in the past.

What outcomes do we expect? On the one hand, greater information might enable receivers to make better decisions. However, on the other hand, if multiple signals are *crossed* (e.g., the action in the current round is different from the message in the previous round), then the receiver might have been better off without the additional information. In the converse case, if the message and actions are identical, then signals are said to be *aligned*.

The authors use the framework of Farrell and Rabin (1996) to argue that the three games differ in the perceived truthfulness of messages sent by the sender. In the Farrell–Rabin framework, there are two conditions, which determine this perception. The first condition, *self-commitment*, is that the pair of actions (sender's message, receivers action) should be a Nash equilibrium. The second condition, *self-signaling*, has the following two requirements. (a) When a sender intends to be truthful, he prefers the receiver to play his best response to the message. For instance, in the SH game, the row player may wish the Pareto optimal outcome (C, C) to occur. To ensure this, the row player may send the message C as his intention to play the game. He would prefer the column player to play C, which is his best response. (b) When a sender intends to be deceptive, he does not prefer the receiver to play his best response to the message. For instance, in the PD game, if the sender sends the message C, then he prefers the receiver not to play his best response (because $10 < 70$).

Using this framework, messages in the SH game are most likely to be truthful and believed (they satisfy self-commitment and self-signaling). In comparison, messages sent in the CH game

Table 12.31 Relative frequencies of cooperation, coordination, and efficiency.

Game	Treatment	Cooperation	Coordination	Payoff efficiency	Efficiency gain
PD	WD	0.354 ^{ab}	–	0.230 ^{ab}	0.230 ^{ab}
	WDL	0.391 ^b	–	0.266 ^b	0.266 ^b
	<i>Cheap talk</i>	0.400 ^b	–	0.260 ^b	0.266 ^b
	<i>Observation</i>	0.404 ^b	–	0.266 ^b	0.266 ^b
	<i>Control</i>	0.222 ^a	–	0.113 ^a	0.113 ^a
SH	WD	0.782 ^{ab}	0.798 ^{bc}	0.752 ^{abc}	0.379 ^{abc}
	WDL	0.823 ^b	0.873 ^c	0.828 ^c	0.570 ^c
	<i>Cheap talk</i>	0.835 ^{ab}	0.840 ^c	0.803 ^{bc}	0.508 ^{bc}
	<i>Observation</i>	0.757 ^{ab}	0.667 ^{ab}	0.636 ^{ab}	0.090 ^{ab}
	<i>Control</i>	0.607 ^a	0.513 ^a	0.453 ^a	–0.368 ^a
CH	WD	0.577 ^{ab}	0.507 ^{ab}	0.746 ^{ab}	0.237 ^{ab}
	WDL	0.604 ^{ab}	0.498 ^{ab}	0.770 ^{ab}	0.311 ^{ab}
	<i>Cheap talk</i>	0.564 ^{ab}	0.532 ^b	0.741 ^{ab}	0.223 ^{ab}
	<i>Observation</i>	0.364 ^b	0.438 ^a	0.780 ^b	0.340 ^b
	<i>Control</i>	0.537 ^a	0.475 ^a	0.696 ^a	0.088 ^a

Italicized treatments from Duffy and Feltovich (2002). Within each game and statistic, entries with no superscripts in common are significantly different at the 10% level (two-sided robust rank-order test, session-level data); superscripts earlier in the alphabet correspond to significantly lower values. PD, Prisoners' Dilemma; SH, Stag Hunt; CH, Chicken; WD, Words and Deeds; WDL, Words, Deeds, and Lies.

Source: Duffy and Feltovich (2006).

satisfy self-commitment but not self-signaling, so they are less likely to be truthful or believed. Messages purporting to play C in the PD game satisfy neither self-commitment nor self-signaling, while those purporting to play D satisfy self-commitment but not self-signaling, hence, they are less likely to be believed relative to, say, the SH game.

We define an index of *payoff efficiency* as follows. The outcome with the maximum possible joint payoffs in the game is given an index of 1 and the one with the lowest possible joint payoff is given an index of 0 (for the prisoner's dilemma game, these are respectively (C, C) and (D, D)). The outcome (C, D) in the prisoner's dilemma game has payoff efficiency equal to $\frac{90}{140} = 0.64$. The *efficiency gain* of an outcome is defined as the increase in payoff efficiency relative to the payoff efficiency of the outcome with the lowest joint payoff. Thus, in the prisoner's dilemma, the efficiency gain of the outcome (C, C) is 1.0 and that of the outcome (C, D) is 0.64.

The results from Duffy and Feltovich (2002, 2006) are shown in Table 12.31. The italicized treatments are those from Duffy and Feltovich (2002). The additional, non-italicized, treatments, WD and WDL are from Duffy and Feltovich (2006). The subscripts on the statistics are explained under the table.⁹⁴

⁹⁴ For instance, for any of the games, the statistic corresponding to the control treatment is superscripted with an 'a'. If, say, the superscript on the statistic for the cheap talk treatment is a 'b' (e.g., in column 3, it is 0.400^b for the PD game), then the two statistics are significantly different at the 10% level. When there are double superscripts, it implies that one of the statistics is significantly lower than the other. For instance, the frequency of cooperation in the cheap talk treatment for the SH game is 0.835^{ab}. This implies that the frequency of cooperation in the control treatment (0.607^a) is significantly lower.


Duffy and Feltovich (2002) found that the differences in the levels of cooperation are more pronounced across games, rather than across the information treatments (cheap talk, observation, and control). They also found that when cheap talk was predicted to be credible (using the Farrell–Rabin criterion described above), the level of cooperation in the “cheap talk only” treatment was greater than that in the “observation only” and in the control treatment that allowed neither cheap talk nor observation. In this case, *words spoke louder than actions*. When cheap talk is predicted not to be credible, then it was found that *actions spoke louder than words*. For instance, in Table 12.31, restricting to the Duffy and Feltovich (2002) results, cooperation is much higher with cheap talk in the PD and SH games but lower in the CH game, relative to the other two treatments.

Turning now to the Duffy and Feltovich (2006) findings, relative to the three treatments in Duffy and Feltovich (2002), the WD treatment does not lead to greater cooperation. There is slightly greater cooperation in the WDL treatment. Hence, adding extra information helps somewhat but not a whole lot more. However, WDL still does not improve on the level of cooperation relative to any of the three treatments in Duffy and Feltovich (2002). One possibility is that multiple signals create greater cognitive difficulties of interpreting them.

Could the results be explained by distinguishing cases where the signals are crossed or aligned? Figure 12.45 answers this question. On the vertical axis are relative frequencies of cooperation, coordination, and efficiency. On the horizontal axis are the three games, PD, SH, and CH. The treatments (see key on the right hand side of the figure) are broken up by the two cases of aligned and crossed signals. The figure also shows the data from Duffy and Feltovich (2002).

In the PD and SH games, all three indicators, the levels of cooperation, coordination, and the efficiency gain, are enhanced when signals are aligned rather than crossed. In the aligned case, the level of cooperation is comparable to the three treatments in Duffy and Feltovich (2002). In the crossed case, the level of cooperation is much lower. Thus, it is not the level of information but the content of information that is vital. Results are less clear cut in the CH game.

In addition to the results described above, the individual-level data reveals the following findings. (1) Even when signals are predicted to be uninformative and/or false, they are truthful, and



* : Sig. diff. between aligned and crossed-signal combinations (one-tailed Wilcoxon summed-ranks test, session-level data, $p < 0.10$).
 ‡ : Sig. diff. between aligned and crossed-signal combinations (one-tailed Wilcoxon summed-ranks test, session-level data, $p < 0.05$).

Figure 12.45 Relative frequencies of cooperation, coordination, and efficiency.

Source: Duffy, J. and Feltovich, N., “Words, deeds, and lies: strategic behaviour in games with multiple signals,” *Review of Economic Studies* (2006) 73: 669–88, by permission of Oxford University Press.

are believed. (2) Overall, the truthfulness of messages varies with the incentive to lie. For instance, lies are more easily caught in the WDL treatment, hence, the senders are more truthful in this treatment. (3) Receivers seem to respond in a manner that is suggestive of payoff maximization and reciprocity.

12.6.7 *Two-sided asymmetric information*

A well-known prediction of classical game theory is that one-sided or two-sided private information may impede trade; these results are summarized in various *no-trade theorems*.⁹⁵ Carillo and Palfrey (2009, 2011) consider novel experiments in two-sided asymmetric information games. Their interest is in considering situations where it is not optimal to engage in some transaction, yet the presence of two-sided asymmetric information may induce trade, contrary to the predictions of classical game theory.

Consider two countries that might go to war. The two countries may keep the status quo (no-war) and enjoy a compromise payoff. Or they may go to war in which the winner gets a higher payoff than the compromise payoff and the loser gets a lower payoff. Yet, on account of two-sided private information, the relative strength of the countries is not mutually known. In the games considered in Carillo and Palfrey (2009), the prediction of a Perfect Bayesian Equilibrium is that both countries should go to war with probability 1 (the analog of the no-trade result). However, the experimental results show that the countries fight with probabilities in the interval [0.5, 0.7]. Thus, there is too much compromise relative to the predictions of theory. Further, and contrary to the predictions of theory, the probability of going to war in the experimental data depends on the magnitude of the compromise payoffs. The results also show limited evidence of any learning by the subjects.

12.7 Public signals and correlated equilibria

Cason and Sharma (2007) consider the following two-player symmetric, strategic form game with multiple Nash equilibria.

	Left (L)	Right (R)
Up (U)	3, 3	48, 9
Down (D)	9, 48	39, 39

Suppose that the monetary payoffs shown in the payoff matrix also represent the respective utilities of the players. This game has two pure strategy Nash equilibria, (D, L), (U, R), and one mixed strategy Nash equilibrium in which the row player plays Up with probability 0.6 and the column player plays Left with probability 0.6.

If the experimenter were to recommend the players to play their part of a Nash equilibrium, say, through a public signal, then, such a recommendation may be expected to be self-enforcing. If it turns out not to be the case, then there are two possibilities. (1) The monetary payoffs in the payoff matrix do not represent utilities of the players. For instance, the players might have social preferences, and, thus, may derive utility from their own monetary payoffs as well as the monetary payoffs of others. (2) The players may doubt the rationality of other players, and, so, believe

⁹⁵ See, for instance, Akerlof (1970) for the case of one-sided asymmetric information and Myerson and Satterthwaite (1983) for the case of two-sided private information.

that other players will not follow their part of the recommended strategy profile. Under these conditions, it may be optimal for a player to ignore the recommendations of the experimenter.

Cason and Sharma (2007) attempt to find out which of these two (or both) possibilities hold. They assign the players to play a *correlated equilibrium* in which the joint payoffs of the players exceed the highest payoffs from any Nash equilibrium. We will only develop the notion of a correlated equilibrium within this simple setup.⁹⁶

Suppose that the players observe some credible public signal of the following form: The probabilities of the strategy profiles (U, L), (U, R), (D, L) and (D, R) are respectively p_{UL} , p_{UR} , p_{DL} , p_{DR} , and $p_{UL} + p_{UR} + p_{DL} + p_{DR} = 1$. In the experiments, the experimenter may credibly and publicly announce these probabilities. The experimenter then randomizes between the four strategy profiles, according to the stated probabilities. Suppose that a particular realization of this randomization is (U, L). The experimenter does not reveal the realization to the players but privately asks each player to play his part of the strategy profile. Players are assumed to use Bayes' rule to form beliefs about the strategies of others.

Suppose, for instance, that the row player is given an instruction to play U. In order to determine whether to comply with the recommendation (play U) or not (play D), the row player needs to form beliefs, using Bayes' rule, about the probability with which the column player will play L. Conditional on the row player given a recommendation to play U, Bayes' rule dictates that the probability of the play of L by the column player is $P(L | U) = \frac{p_{UL}}{p_{UL} + p_{UR}}$. Analogously, one can compute that

$$P(R | U) = \frac{p_{UR}}{p_{UL} + p_{UR}}; P(L | D) = \frac{p_{DL}}{p_{DL} + p_{DR}}; P(R | D) = \frac{p_{DR}}{p_{DL} + p_{DR}}.$$

Then $(p_{UL}, p_{UR}, p_{DL}, p_{DR})$ is a correlated equilibrium if each player plays the assigned recommendation, conditional on the belief that all other players will also follow their recommendations. This requires the satisfaction of the following inequalities.

$$P(L | U) 3 + P(R | U) 48 \geq P(L | U) 9 + P(R | U) 39. \quad (12.32)$$

$$P(L | D) 9 + P(R | D) 39 \geq P(L | D) 3 + P(R | D) 48. \quad (12.33)$$

$$P(U | L) 3 + P(D | L) 48 \geq P(U | L) 9 + P(D | L) 39. \quad (12.34)$$

$$P(U | R) 9 + P(D | R) 39 \geq P(U | R) 3 + P(D | R) 48. \quad (12.35)$$

These inequalities are of the nature of incentive compatibility conditions. For instance, consider the inequality in (12.32). Assuming that (a) the column player plays the recommended strategy, and (b) the row player is given the recommendation to play U, the column player will play L with probability $P(L | U)$ and R with probability $P(R | U)$, irrespective of the strategy of the row player. Thus, (12.32) requires that the row player should get a higher payoff from complying with the recommendation to play U (the LHS), rather than deviating and playing D (the RHS). The inequality in (12.33) requires that the row player prefers to comply with the recommendation to play D rather than to deviate and play U. Inequalities (12.34), (12.35) give analogous conditions for the column player.

The set of all payoffs that can be implemented using a correlated equilibrium are enclosed by the dotted line in Figure 12.46. This is the set of all convex sets of payoffs that can be formed

⁹⁶ For further details on correlated equilibria, we refer the reader to any standard text in game theory, e.g., Fudenberg and Tirole (1991).




Figure 12.46 The set of possible payoffs in a correlated equilibrium.

Source: *Economic Theory* 33 (2007), 11–27, “Recommended play and correlated equilibria,” Cason, T. N. and Sharma, T. © Springer Verlag with kind permission from Springer Science and Business Media.

by using the four inequalities (12.32)–(12.35), the non-negativity conditions on probabilities, and the condition that probabilities sum to one. The Nash equilibrium payoff hull which is shown enclosed by the bold lines is a subset of the set of payoffs that can be obtained in a correlated equilibrium.⁹⁷ The cooperative outcome, (D, D), with the highest joint payoffs (39, 39) (labeled as ‘perfect cooperation’) lies outside the set of correlated equilibrium payoffs. The highest symmetric payoffs that can be obtained in the correlated equilibrium are (31.25, 31.25). This is shown as a point on the outer (northeast) boundary of the set of correlated equilibrium. However, players are indifferent between their pure strategies at this equilibrium, hence, Cason and Sharma (2007) decide to implement a correlated equilibrium just southwest of this point (shown in Figure 12.46). In this equilibrium, each player receives an expected payoff of 30.6, which exceeds the joint payoff from any Nash equilibrium.

The previous literature had found mixed evidence on the implementation of recommended play in a non-correlated equilibrium setting.⁹⁸ However, in a correlated equilibrium setting, Moreno and Wooders (1998) find evidence that is consistent with the play of recommended strategies. Crucially, they allow for non-binding preplay communication, while Cason and Sharma (2007) do not. This factor seems crucial in accounting for the difference in results because Cason and Sharma find that players do not always follow recommendations. Treatments A through F in Cason and Sharma (2007) are described in Table 12.32. Opponents were randomly matched in each of the 60 periods of play.

Treatments A and B consider the effect of recommendations on the strategies chosen by the players. Humans play against robot opponents in treatments C and D. These treatments control for social preferences, as well as beliefs about the rationality of others, that can come into play when humans play against human opponents. Treatments C and D differ with respect to

⁹⁷ The convex hull of the Nash equilibria is formed from the payoffs in the two pure strategy Nash equilibria (9,48), (48,9) as well as the payoff from the mixed strategy Nash equilibrium (21,21).

⁹⁸ See, for instance, Van Huyck et al. (1992), Brandts and MacLeod (1995), Seely et al. (2005).

Table 12.32 A description of the various treatments in Cason and Sharma (2007).

Treatment	Explanation of the treatment	Experience	Subjects
A	Human opponents, no recommendations	None	52
B	Human opponents with recommendations	None	80
C	Robot opponents with recommendations (no one receives robot earnings)	None	31
D	Robot opponents with recommendations (Humans receive robot earnings)	None	26
E	Human opponents with recommendations	Experienced in B	24
F	Human opponents with recommendations	Experienced in C	24

the receipt of robot earnings. If other humans receive robot earnings (as in treatment D), then social preferences might be activated once again. Subjects were told that robots will follow the recommendations without error. Treatments E and F consider the role of experience. These are similar to treatment B, with the exception that the subjects playing the treatment have had prior experience, respectively, in treatments B and C. Furthermore, experimenters elicit the beliefs of subjects about “their expectations of the play of others” and subjects were paid depending on the proximity of their beliefs with the actual play. The beliefs of the players were elicited after the recommendations were made to them. Players then submitted their actions and their beliefs.

In a comparison of treatments A and B, recommendations seem to enhance the play of the cooperative outcome associated with the (D, R) choice, but the significance level is marginal at 10%. However, across all treatments, the recommendations increased the payoffs of the players significantly. This arose not so much by enhancing the play of the cooperative outcome (D, R) as by reducing the play of the worst outcome (U, L).

The percentage of actions consistent with the recommended play are shown in Figure 12.47 for three treatments B, C, and D and for two recommendations: U/L (i.e., U to row player and L to column player) and D/R (D to row player and R to column player). The U/L recommendation is followed to a greater extent than the D/R recommendation. Subjects follow the recommendation more frequently (in a statistical sense) when they play robot as compared to human opponents. The likelihood of following the recommendations is very similar in the two robot treatments, so the results across these two treatments can be pooled. Looking at individual data, 75% of the subjects in the robot treatments followed the recommendations for 90% of the observations. Less than 50% of the subjects followed the recommendations for 90% of the observations when the opponent is a human.

In the two treatments with experienced subjects (E and F), the level of the payoffs and the frequency of following the recommendations are both statistically similar. Hence, these two treatments can also be pooled. Furthermore, the average ratio of actual to recommended payoffs in the




Figure 12.47 Percentage of actions consistent with the recommended play in the various treatments.

Source: *Economic Theory* 33 (2007), 11–27, “Recommended play and correlated equilibria,” Cason, T. N. and Sharma, T. © Springer Verlag with kind permission from Springer Science and Business Media.

experienced versus non-experienced subject treatments are statistically identical. So experience does not seem to alter payoffs.

In the correlated equilibrium that the authors choose to implement (see above), $p_{UL} = 0$. We can use this fact and the elicited beliefs of subjects to check if they believe that their opponents will play the recommended action. Suppose that the recommendation to the row player is to play U. Then, using the relevant formulae above, we can calculate that he should assign probability zero that the opponent will play L ($P(L | U) = 0$) and a probability 1 that the opponent will play R ($P(R | U) = 1$). Analogously if the column player is recommended to play L, then he should assign probability zero to the row player playing U and a probability 1 to the row player playing D.

Results in Figure 12.48, given for periods 31–60, show that the elicited beliefs of players did not always reveal such a pattern. The histogram in the left panel shows that for human opponents, inexperienced subjects had the correct beliefs about their opponents about 77% of the time (when the prediction is 100%). Beliefs get better with experience but only match the predictions when the opponent is a robot. The histogram in the right hand side panel in Figure 12.48 has a predicted height of 0% but actually varies between a height of 25% to 40%, depending on the conditions used. With human opponents, play does not conform to predictions because subjects believe that (human) opponents will not follow the recommendations. Cason and Sharma (2007, p. 26) conclude that: “Our results indicate that recommendations are not sufficient to induce players to switch to any chosen equilibrium. Subjects do not switch because they believe that their opponents will not switch.”

Cason and Sharma (2007) implement a correlated equilibrium in which the payoffs dominate those in a Nash equilibrium. This may potentially confound the coordinating role of recommendations of the mediator with a desire of subjects to improve their payoffs. Motivated by this find-




Figure 12.48 Beliefs for different recommendations for periods 31–60.

Source: *Economic Theory* 33 (2007), 11–27, “Recommended play and correlated equilibria,” Cason, T. N. and Sharma, T. © Springer Verlag with kind permission from Springer Science and Business Media.

	D	C
D	0, 0	9, 3
C	3, 9	7, 7

Figure 12.49 Symmetric game of chicken in Duffy and Feltovich (2010).

Source: “Corrected equilibria, good and bad: an experiments study,” Duffy, J. and Feltovich, N., *International Economic Review* 51: 701–21. Copyright © (2010) by the Economics Department of the University of Pennsylvania and the Osaka University Institute of Social and Economic Research Association.

ing, Duffy and Feltovich (2010) consider experimental evidence on correlated equilibrium in a symmetric game of chicken shown in Figure 12.49.

The following proposition summarizes the various equilibria of the game. The proof is similar to the one we outlined above for the Cason–Sharma paper, so it is relegated to the problem set.

Proposition 12.23 Consider the game of chicken in Figure 12.49.

- There are two pure strategy equilibria (D, C) and (C, D). There is also a mixed strategy equilibrium in which both players play strategy D with probability $2/5$. In the mixed strategy equilibrium, each player’s expected payoff is 5.4.
- Suppose that players could condition their play on a publicly available signal, say, the toss of a coin. Following heads, the row player plays D and the column player plays C. Following tails, the row player plays C and the column player plays D. Each player’s expected payoff in this case is 6, higher than the expected payoff from any Nash equilibrium.
- (Aumann, 1974) Suppose that players condition their choice of strategies on the outcome of three publicly observed random signals, X, Y, Z, each chosen with equal probability. Player 1 (row player) learns only if X was chosen or not, while player 2 (column

player) learns only if Z was chosen or not. Player 1 plays D if the public signal is known to be X and C otherwise. Player 2 plays D if the public signal is known to be Z and C otherwise. These correlated strategies are a mutual best response (a correlated equilibrium) and yield an expected payoff of $6\frac{1}{3}$ to each player.

- (d) Denote by p_{DD} , p_{DC} , p_{CD} , p_{CC} , the probabilities with which a mediator chooses, respectively, the four strategy profiles (D, D) , (D, C) , (C, D) , (C, C) , through some public randomization device. The mediator draws a strategy profile based on the stated probability distribution. However, the mediator does not tell the players which strategy profile was drawn. Players are only, privately, given recommendations to play their part of the strategy profile that is drawn.⁹⁹ Suppose that the following conditions hold.¹⁰⁰

$$2p_{DC} \geq 3p_{DD}; 3p_{CD} \geq 2p_{CC}; 2p_{CD} \geq 3p_{DD}; 3p_{DC} \geq 3p_{CC} \quad (12.36)$$

and

$$p_{DD} + p_{DC} + p_{CD} + p_{CC} = 1. \quad (12.37)$$

Then p_{DD} , p_{DC} , p_{CD} , p_{CC} constitute a correlated equilibrium. ■

Figure 12.50 shows the convex hull of the Nash equilibrium payoffs. This is the triangle labelled A with vertices that correspond to the payoffs from the three Nash equilibria; the two pure strategy Nash equilibria (C, D) and (D, C) with payoffs 3,9 and 9,3 and the mixed strategy equilibrium with payoffs 5.4,5.4. The set of correlated equilibrium payoffs that satisfy (12.36), (12.37) is given by the quadrilateral (union of the regions B_l , A , B_h) with vertices (3, 9), (4.5, 4.5), (9, 3), and $(6\frac{3}{7}, 6\frac{3}{7})$. The highest symmetric correlated equilibrium payoffs are $(6\frac{3}{7}, 6\frac{3}{7})$ and the lowest are (4.5, 4.5).¹⁰¹

Duffy and Feltovich (2010), attempt to implement the following outcomes; the first three are correlated equilibria, while the fourth is not.

1. *Nash-recommendations treatment*: The two pure strategy Nash equilibria (D, C) and (C, D) are recommended with a probability of 0.5 each. The expected payoff of each player is 6 in this case.
2. *Good-recommendations treatment*: Each of the three strategy profiles (D, C) , (C, C) , (C, D) is drawn with a probability of one third ((D, D) is drawn with zero probability) and then a recommendation is made by the mediator. The resulting equilibrium lies in the convex hull of the correlated equilibrium and the expected payoff is $6\frac{1}{3}$ for each player; this payoff exceeds any of the payoffs in the convex hull of the Nash equilibrium payoffs.
3. *Bad-recommendations treatment*: The strategy profiles (D, C) , (C, D) are selected by the mediator with probability 0.4 each and the profile (D, D) is selected with probability 0.2 ((C, C) is selected with probability zero). The expected payoff for each player is 4.8, which is lower than any payoff in the convex hull of the Nash equilibrium payoffs.

⁹⁹ For instance, if the strategy profile (C, D) is drawn, the row player is given a private recommendation to play C and the column player is given a private recommendation to play D .

¹⁰⁰ We assume that either the players are risk neutral or that the numbers in the payoff matrix in the game of chicken are utility values.

¹⁰¹ In general, the set of correlated equilibria contain the set of Nash equilibria. Hence, the convex hull of the set of correlated equilibria contain the convex hull of the set of Nash equilibria. Further, the set of payoffs obtained in the convex hull of Nash equilibria are contained within the set of payoffs obtained in the convex hull of correlated equilibria.




Figure 12.50 Equilibria in the game of chicken in Duffy-Feltovich.

Source: "Correlated equilibria, good and bad: an experimental study," Duffy, J. and Feltovich, N., *International Economic Review* 51: 701–21. Copyright © (2010) by the Economic Department of the University of Pennsylvania and the Osaka University Institute of Social and Economic Research Association.

4. *Very good recommendations treatment*: In the final treatment, an equilibrium that is outside the convex hull of the correlated equilibria is selected. In this case, the mediator selects the outcome (C, C) with probability 0.8 and (D, C), (C, D) each with probability 0.1. If players follow the recommendation of the mediator, then the expected payoff of each player is 6.8, which exceeds the highest possible symmetric correlated equilibrium payoffs of $(6\frac{3}{7}, 6\frac{3}{7})$.

A total of 16 experimental sessions were run, 4 for each treatment. Each session had 12 subjects in experiments. Each subject played 20 rounds with the recommendations and 20 rounds without the recommendations, resulting in a total of 7,580 observations. The results are shown in Table 12.33 for the case of no recommendations (under the heading, None) and for the four treatments under recommendations (respectively, under the headings, Nash, Good, Bad, and Very good).

For the case of no recommendations, subjects played without recommendations for 40 rounds. However, in the other cases, they played 20 rounds under the recommendation and 20 rounds without any recommendations. The combined results are also shown. In order to interpret the results in Table 12.33, consider, for instance, the treatment "None" and the combined play over the 40 rounds. In the mixed strategy Nash equilibrium (Proposition 12.23(a)), each player plays D with a probability of $\frac{2}{5}$, so that the joint probability of the play of the strategy profile (D, D) is $(\frac{2}{5})^2 = 0.16$, which appears quite close to the experimentally observed frequency of 0.159. Let us now consider the results in more detail.

1. Consider first the case with no recommendations. In several sessions and time periods, the observed play at the aggregate level is close to the mixed strategy equilibrium. Overall, the authors interpret their findings as supporting a mixed strategy Nash equilibrium, for

Table 12.33 Frequencies of observed outcomes at the aggregate level.

Recommendations	Ordering	Outcome				Average payoff
		(D, D)	(D, C)	(C, D)	(C, C)	
None	+	0.145	0.200	0.274	0.381	5.513
	—	0.174	0.226	0.295	0.305	5.261
	Combined	0.159	0.213	0.284	0.343	5.387
Good	+	0.154	0.258	0.363	0.225	5.300
	—	0.125	0.258	0.279	0.338	5.588
	Combined	0.140	0.258	0.321	0.281	5.444
Bad	+	0.213	0.267	0.246	0.275	5.000
	—	0.171	0.242	0.208	0.379	5.354
	Combined	0.192	0.254	0.227	0.327	5.177
Very good	+	0.138	0.175	0.254	0.433	5.608
	—	0.200	0.279	0.229	0.292	5.092
	Combined	0.169	0.227	0.242	0.363	5.350

Note: "+": recommendations received in rounds 21–40; "—": recommendations received in rounds 1–20.

Source: Duffy and Feltovich (2010).

- aggregate data, when no recommendations are given. The null hypothesis that behavior is generated by an iid mixed strategy profile is statistically rejected by using a chi-square test.¹⁰²
2. Recommendations matter, although they do so to a different extent in each treatment with recommendations. Overall, a chi-square test rejects the hypothesis that the distribution of outcomes in the no-recommendations and the recommendations treatment (but omitting the Very good recommendations treatment) are identical.
 3. When a correlated equilibrium is recommended to the players (i.e., ignoring the Very good recommendations treatment that lies outside the relevant convex hull of payoffs), play moves in the direction of the recommended equilibrium, relative to the mixed strategy equilibrium. In each case, the frequency of outcomes is very different from the mixed strategy equilibrium.
 4. In the Good-recommendations and the Nash-recommendations treatments, the recommendations were followed 70–75% of the time.
 5. When the recommendation does not arise from a correlated equilibrium, as in the Very good recommendations treatment, subjects learn to ignore the recommendations. Thus, based on this evidence, correlated equilibrium would seem to be a *necessary condition* for recommendations to be followed. This treatment also rules out the possibility that subjects were simply following the recommendations blindly, or to please the experimenter.
 6. Subjects do not follow the recommendations in the Bad-recommendations treatment, which is a correlated equilibrium, but the payoffs are dominated by the convex hull of the Nash equilibrium payoffs. Thus, correlated equilibrium recommendations is not a *sufficient condition* for the recommendations to be followed.

¹⁰² If, however, the outcomes (C, D) and (D, C) are pooled, then play is close to the predictions of the mixed strategy Nash. However, pooling the outcomes (C, D) and (D, C) is inconsistent with a mixed strategy equilibrium in the first place. Such pooling can neither confirm nor reject a mixed strategy equilibrium.

12.8 Strategic complements and strategic substitutes

Suppose that two individuals, or firms, indexed by i, j , are involved in some strategic interaction. They choose respective actions x_i, x_j from their respective action sets $X_i, X_j \subset \mathbb{R}$. The respective payoff functions of the two players are given by $\pi_i(x_i, x_j), \pi_j(x_i, x_j)$, where π_i, π_j are strictly concave and twice continuously differentiable. The Nash equilibrium (x_i^*, x_j^*) satisfies the first order conditions

$$\frac{\partial \pi_i(x_i^*, x_j^*)}{\partial x_i} = 0; \quad \frac{\partial \pi_j(x_i^*, x_j^*)}{\partial x_j} = 0. \quad (12.38)$$

Our assumptions guarantee that the optimal reaction functions of the two players are unique, and these are given, respectively, by $x_i = R_i(x_j)$ and $x_j = R_j(x_i)$.¹⁰³ The slopes of the reaction functions, given by

$$\frac{\partial R_i}{\partial x_j} = \frac{\frac{\partial \pi_i}{\partial x_i \partial x_j}}{-\frac{\partial^2 \pi_i}{\partial^2 x_i}}, \quad \frac{\partial R_j}{\partial x_i} = \frac{\frac{\partial \pi_j}{\partial x_i \partial x_j}}{-\frac{\partial^2 \pi_j}{\partial^2 x_j}}, \quad (12.39)$$

are completely determined by the respective signs of the numerators (the denominators are positive).

Definition 12.3 (*Strategic complements and strategic substitutes; Bulow et al., 1985*)

- (a) The actions of the two players are said to be *strategic complements* (SC) if the reaction functions are upward sloping, i.e. $\frac{\partial \pi_i}{\partial x_i \partial x_j} > 0$ and $\frac{\partial \pi_j}{\partial x_i \partial x_j} > 0$.
- (b) The actions of the players are said to be *strategic substitutes* (SS) if the reaction functions are downward sloping, i.e. $\frac{\partial \pi_i}{\partial x_i \partial x_j} < 0$ and $\frac{\partial \pi_j}{\partial x_i \partial x_j} < 0$.

Thus, actions are SC if an increase in the action by a player increases the marginal payoff to the other player from increasing his/her action. The converse happens under SS. An important implication of SC is that a change in one player's action gives the other player an incentive to move in the same direction, so we should expect to see positively correlated actions. Analogously, under SS, we should expect to observe negative correlation in actions. The concepts of strategic substitutes and strategic complements have great importance in many areas of economics; for instance, multiple equilibria in macroeconomics (Cooper and John, 1988), learning dynamics in games (Milgrom and Roberts, 1991), and supermodularity (Topkis, 1998).

In SC and SS games, because of the mutual dependence of payoffs on actions, *externalities* or *spillovers* arise, which can be positive or negative. In the presence of these externalities, the Nash equilibrium will typically not be Pareto optimal.

Definition 12.4 (*Cooper and John, 1988*): A game has (i) *positive spillovers/externalities* if $\frac{\partial \pi_i}{\partial x_j} > 0, i \neq j$, and (ii) *negative spillovers/externalities* if $\frac{\partial \pi_i}{\partial x_j} < 0, i \neq j$.

¹⁰³ The reaction functions satisfy $\frac{\partial \pi_i(R_i(x_j^*), x_j^*)}{\partial x_i} = 0$ and $\frac{\partial \pi_j(x_i^*, R_j(x_i^*))}{\partial x_j} = 0$.

Proposition 12.24 When $\frac{\partial \pi_i}{\partial x_j} > 0$, the Nash equilibrium will be Pareto dominated by higher actions. On the other hand, when $\frac{\partial \pi_i}{\partial x_j} < 0$, the Nash equilibrium will be dominated by lower actions.

12.8.1 Implications of SS and SC for the degree of cooperation

In Potters and Suetens (2009), pairs of subjects play a finite number of repetitions of a game that exhibits SC, SS, and positive and negative spillovers. They consider the following four treatments. (1) SS with negative externalities (SUBSTneg). (2) SS with positive externalities (SUBSTpos). (3) SC with negative externalities (COMPLneg). (4) SC with positive externalities (COMPLpos).

In the discussion below, we describe only the positive externality games, SUBSTpos and COMPLpos. These games can be converted into a negative externality game by a simple transformation of the variables in the following manner. If $x_i \in [0, m]$ is the decision variable in the positive externality game, then replacing x_i by $m - y_i$ changes this into a negative externality game, where y_i is the decision variable in the negative externality game.

The payoff of player i in the COMPLpos treatment, $\pi_i^{COMPLpos}$, and in the SUBSTpos treatment, $\pi_i^{SUBSTpos}$, is given, respectively, by

$$\pi_i^{COMPLpos}(x_i, x_j) = a + bx_i + cx_j - dx_i^2 + ex_j^2 + fx_ix_j; \quad b, c, d, f > 0, e \geq 0; x_i, x_j \geq 0 \quad (12.40)$$

$$\pi_i^{SUBSTpos} = \alpha + \beta x_i + \gamma x_j - \delta x_i^2 + \epsilon x_j^2 - \xi x_ix_j; \quad \beta, \gamma, \delta, \xi > 0, \epsilon \geq 0; x_i, x_j \geq 0. \quad (12.41)$$

The restrictions in (12.40) ensure that the actions of the two players are SC and that there are positive externalities. Analogously, the restrictions in (12.41) ensure that the actions of the two players are SS and that there are positive externalities. Given the assumptions made, we can use (12.38) and (12.40), (12.41) to find the Nash equilibrium. Similarly, by adding the payoffs of both players, in each case, one can find the joint payoff maximizing choice. The results are summarized in the following proposition; the proof is straightforward and it is omitted.

Proposition 12.25

- (1) Consider first the COMPLpos treatment. The unique, symmetric, Nash equilibrium (N) and the unique, symmetric, joint payoff maximizing choices (JPM) are respectively given by

$$x^N = \frac{b}{2d - f}; \quad x^{JPM} = \frac{b + c}{2(d - e - f)}.$$

- (2) In the SUBSTpos case, the respective unique, symmetric Nash and joint payoff maximizing solutions are given by

$$x^N = \frac{\beta}{2\delta + \xi}; \quad x^{JPM} = \frac{\beta + \gamma}{2(\delta - \epsilon + \xi)}.$$

The following six conditions are imposed across the treatments. (1) Identical Nash equilibrium strategies. (2) Identical Nash equilibrium payoffs. (3) Identical joint payoff maximizing (JPM) choice. (4) Identical JPM payoffs. (5) Identical payoffs from defection. (6) Identical slopes of the best response functions (in absolute value). All these conditions can be satisfied by imposing




Figure 12.51 The best response functions and equilibria in various cases.

Source: Potters, J. and Suetens, S., "Cooperation in experimental games of strategic complements and substitutes," *Review of Economic Studies* (2009) 3: 1125–47, by permission of Oxford University Press.

suitable restrictions on the parameters in (12.40), (12.41).¹⁰⁴ The Nash and JPM choices are shown in Figure 12.51 (readers should replace indices i, j by 1, 2 in the figure).

A total of 28 subjects (or 14 pairs of subjects) participated in each treatment that was run over 31 rounds. The pairings were maintained through the 31 rounds. At the end of each round, subjects were informed of the actions of their opponent, their earnings, and the earnings of the opponent.

Definition 12.5 The degree of cooperation for pair k in round t of the experiment, ρ_{kt} , is defined as

$$\rho_{kt} = \frac{\bar{x}^{kt} - x^N}{x^{JPM} - x^N},$$

where $\bar{x}^{kt}$ is the average choice of pair k at round t . If $\bar{x}^{kt} = x^N$, $\rho_{kt} = 0$ and if $\bar{x}^{kt} = x^{JPM}$, then $\rho_{kt} = 1$. If $x^{JPM} < \bar{x}^{kt} < x^N$, then $0 < \rho_{kt} < 1$.

¹⁰⁴ The interested reader can check that these restrictions are: $\alpha = a$, $\beta = \frac{b(2d-f)}{2d+f}$, $\gamma = c + \frac{2bf}{2d+f}$, $\delta = \frac{d(2d-f)^2}{(2d+f)^2}$, $\epsilon = e + \frac{2f^3}{(2d+f)^2}$, $\xi = \frac{f(2d-f)^2}{(2d+f)^2}$. Recall also our procedure for transforming a positive externality game into a negative externality game. We require $m = \frac{2b}{2d-f}$. In conjunction, these guarantee the six conditions.




Figure 12.52 Average degree of cooperation in each round for complements and substitutes.

Source: Potters, J. and Suetens, S., "Cooperation in experimental games of strategic complements and substitutes," *Review of Economic Studies* (2009) 3: 1125–47, by permission of Oxford University Press.

The results are as follows.

1. The sign of the externality does not affect the degree of cooperation. Thus, ρ_{kt} is statistically identical in the treatments SUBSTneg and SUBSTpos and in the treatments COMPLneg and COMPLpos. Therefore, the authors pool SUBSTneg and SUBSTpos into SUBST and similarly COMPLneg and COMPLpos into COMPL.
2. The main finding is that cooperation is higher in the complements treatment, COMPL, as compared to the substitutes treatment, SUBST, over all of the 30 rounds; see Figure 12.52. There are also *end-game effects* that are observed in other finitely repeated games (Ledyard, 1995). The endgame effect is larger in the substitutes treatment.
3. Figure 12.52 also shows that play does not converge to the Nash equilibrium of the stage game or to any other outcome (see Definition 12.5 for the conditions required on ρ_{kt}).
4. Recall from the discussion above that in the complements case, the actions of the two players are likely to be positively correlated, while in the substitutes treatment, they are likely to be negatively correlated. This is not observed in the individual data. In both the COMPL and SUBST treatments, the players change their actions in the same direction as the opponent. Social preferences might offer one potential explanation for these results. Suppose that individual players exhibit reciprocity. If one player takes an action directed towards the JPM choice, then the other player does too, motivated by *positive reciprocity*. Analogously, choices directed away from the JPM choice elicit *negative reciprocity*, inducing the opponent to also move away from the JPM choice. The statistical analysis suggests that such results are more pronounced in the case of COMPL as compared to the SUBST treatment.
5. Across both treatments, when both paired players in a round make JPM choices, they tend to stick to these choices in subsequent rounds of the experiment. Call these the "JPM pairs" and the rest as the "non-JPM pairs." Across both treatments, about 40% of the players are JPM pairs. The behavior of these pairs is shown in Figure 12.53. Consider JPM pairs, shown on the right panel of Figure 12.53. The degree of cooperation of these pairs is higher in the COMPL treatment in the first half of the experiment. Also these pairs reach the JPM choice earlier in the COMPL treatment relative to the SUBST treatment. Thus, the complements




Figure 12.53 Average degree of cooperation for non-JPM and JPM pairs.

Source: Potters, J. and Suetens, S., "Cooperation in experimental games of strategic complements and substitutes," *Review of Economic Studies* (2009) 3: 1125–47, by permission of Oxford University Press.

treatment is relatively more conducive to cooperation. Eventually, these pairs reach full cooperation ($\rho_{kt} = 1$) in both treatments. For non-JPM pairs, cooperation is still relatively higher in the COMPL treatment. However, unlike the JPM pairs, they do not achieve the JPM choice in either of the two treatments. Furthermore, the degree of cooperation diverges across the two treatments.

6. Choices of players in price setting experiments are typically SC, while in quantity setting games they are typically SS. The degree of collusion is also typically observed to be higher in the former. The results in this paper, which suggest that the degree of cooperation is likely to be higher in the case of SC is a possible explanation of these experimental findings.

12.8.2 The macroeconomics of nominal inertia

Issues of nominal inertia have been central to modern New Keynesian macroeconomics. A range of theories have been proposed as potential explanations,¹⁰⁵ but they have ignored the impact of behavioral factors. Fehr and Tyran (2008) introduce two behavioral factors with good empirical support: *Money illusion* and *anchoring*, in a model that allows considerations of SC and SS. They find that the extent of behavioral biases is endogenous and depends on the type of strategic interaction, i.e., SC or SS. Behavioral biases are less likely in those strategic environments where they are less costly in terms of forgone payoffs. The Fehr–Tyran findings are inconsistent with the current practice in theoretical macroeconomics; we elaborate below.

The basic structure of the experiment is described in Figure 12.54. Groups of n players are formed and play a pricing game for T time periods. The composition of the group members

¹⁰⁵ See, for instance, Blanchard and Fischer (1989, chapters 8, 9) and Romer (2006, chapter 6).




Figure 12.54 A partial representation of the Fehr–Tyran (2008) experimental setting.

Source: Fehr and Tyron (2008) with permission of The Econometric Society.

was not altered over the T periods. In the first half of the experiment (i.e., $T/2$ time periods), called the pre-shock phase, the money supply was fixed at a level M_0 . In the post-shock phase, which covers the next $T/2$ time periods, a negative monetary shock is implemented that cuts the money supply by a half, so that the new money supply is $M_1 = M_0/2$. The money supply shock is fully anticipated. This is achieved by announcing the shock at the beginning of the post-shock phase and then implementing it. In the experiments, $M_0 = 42$ and $T = 30$. The restart phase is described later.

The *real payoff* of the i th player, v_i , in a group of n players, is given by $v_i = v_i(P_i, \bar{P}_{-i}, M)$, where P_i is the nominal price chosen by the player, $\bar{P}_{-i}$ is the average nominal price chosen by the other $n - 1$ players, and M is the money supply. We assume that n is sufficiently large. It is assumed that v_i is homogeneous of degree zero in $P_i, \bar{P}_{-i}, M$, thus, $v_i(\lambda P_i, \lambda \bar{P}_{-i}, \lambda M) = \lambda^0 v_i$, where $\lambda > 0$. Setting $\lambda = 1/M$ gives

$$v_i = \phi \left(\frac{P_i}{M}, \frac{\bar{P}_{-i}}{M} \right), \text{ where } \phi \left(\frac{P_i}{M}, \frac{\bar{P}_{-i}}{M} \right) = v_i \left(\frac{P_i}{M}, \frac{\bar{P}_{-i}}{M}, 1 \right). \quad (12.42)$$

Half the players in the model are of type x and the remaining half are of type y . The types differ only in their payoffs from the same pair of numbers $P_i, \bar{P}_{-i}$. For player i , denote by $\bar{P}_{-i}^j$ the average price chosen by the other players of type $j = x, y$. Given the composition of the two types of players, and large n , $\bar{P}_{-i} = 0.5\bar{P}_{-i}^x + 0.5\bar{P}_{-i}^y$.

Since there is within-type homogeneity of players, we focus on the symmetric case where all players of type x choose a price $\bar{P}^x$, and all players of type y choose the price $\bar{P}^y$. Thus, the aggregate price level is given by

$$P = 0.5\bar{P}^x + 0.5\bar{P}^y. \quad (12.43)$$

If there are no costs of changing prices (e.g., no macroeconomic frictions or menu costs), then, as in traditional flexible price macroeconomic models, we assume that a proportional change in M leads to the same proportional change in $\bar{P}^x, \bar{P}^y$. Hence, money is neutral. Thus, we get the well-known result that monopolistic competition alone, in the absence of frictions, does not lead to nominal rigidity (Blanchard and Fischer, 1989, Section 8.1).

Holding M fixed, the payoffs of the individuals in (12.42) can be represented in $(P_i, \bar{P}_{-i})$ space. Subjects choose a price $P_i \in \{1, 2, \dots, 30\}$, and are shown a payoff table that gives the payoffs for each pair $(P_i, \bar{P}_{-i})$. The actions $P_i, \bar{P}_{-i}$ can either be strategic complements (the CT treatment) or strategic substitutes (the ST treatment). In any treatment, the equilibrium is identical. For both treatments, CT and ST, in the pre-shock period, $\bar{P}^x = 22$ and $\bar{P}^y = 28$, thus, using (12.43), $P = 25$. For the post shock period, $\bar{P}^x = 11$ and $\bar{P}^y = 14$, so $P = 12.5$.

While choosing their price in any period, players have to form an expectation about the average price of others, $\bar{P}_{-i}^e$. The actual value $\bar{P}_{-i}$ is revealed to players at the end of each round of the experiment. In any round, players had access to all past values of $\bar{P}_{-i}$ as well as their own choices P_i . Since one of the objectives of the paper is to test if money illusion plays an important role, the payoffs are shown to the subjects in nominal terms, $V_i = P v_i$, where v_i is the real payoff of player i . In only one treatment, a control treatment, were payoffs observed by the players in real terms. The payoff matrices of each type of player, are common knowledge and there are no issues of asymmetric information.

The payoff matrices for each of the several cases are very large. However, to get a sense of the game being played and the ease of calculating best replies in the model, we report one payoff matrix, in Figure 12.55, for type y players in the post-shock phase for the substitutes treatment. The selling price is P_i (numbers 1–30 in the first column) and the average price of other firms is $\bar{P}_{-i}$ (numbers 1–30 in the first row). The best reply function for player i for all possible levels of $\bar{P}_{-i}^e$ is shown by the shaded path in the payoff matrix (subjects did not see the shaded best reply path). Along this path, the payoff of player i is 39 throughout.

Corresponding to the payoff matrix for substitutes in the post-shock case shown in Figure 12.55 for type y players, there is an analogous payoff matrix for type x players (not shown) that differs in the payoffs. In equilibrium, type x and type y players choose respective prices 11 and 14, and their respective expectations of the average prices of the other type are 13 and 12. Thus, average price, using (12.43), is 12.5. From Figure 12.55, the best response of a type y player is 14, as claimed. Similarly the best response of a type x player is 11 (using the payoff matrix for type x players that is not shown).

For the case of complements, the equilibrium is identical to the substitutes case, for the post-shock period. In the pre-shock case, by contrast, for complements and substitutes, players of type x and y choose respective prices 22 and 28 (their respective price expectations of the other type are 26 and 24).

For each column in the payoff matrix, there is a unique optimal price for a player and the corresponding payoff is 39. All remaining payoffs in the column are very small numbers, enabling players to quickly spot the optimal value; this rules out the confounds of cognitive complexity. Furthermore, at the end of each period, the average price of other firms is revealed to each player. In light of this, it would be highly unexpected, therefore, if firms did not quickly converge to the equilibrium price in the pre-shock phase when money illusion and anchoring have not yet influenced actions. This intuition turns out to be correct.

At the beginning of period 16, the first period in the post-shock phase, the exogenous drop in money supply is announced publicly. Subjects are given enough time to study the new payoff tables (corresponding to $M_0/2 = 21$) before deciding on their nominal price. There are no informational frictions in the model, which are critical for nominal frictions in standard flexible price neoclassical models exemplified by Lucas (1972). There are also no nominal frictions (such as staggered wages; see Fischer, 1977) that are critical in New Keynesian models.

The *anchoring heuristic* is considered in more detail in Chapter 19. In the context of these experiments, anchoring implies that the equilibrium achieved in the pre-shock phase acts as a *reference*

Selling price	Average price of other firms																													
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	2	2	2	2	3	5	9	20	39	39	39	39
2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	2	2	2	2	3	5	9	20	39	20	20	20	20
3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	2	2	2	2	2	3	5	9	20	39	20	9	9	9	9
4	1	1	1	1	1	1	1	1	1	1	1	1	1	1	2	2	2	2	2	3	5	9	20	39	20	9	5	5	5	5
5	1	1	1	1	1	1	1	1	1	1	1	1	2	2	2	2	2	2	3	5	9	20	39	20	9	5	3	3	3	3
6	1	1	1	1	1	1	1	1	1	2	2	2	2	2	2	2	2	5	9	20	39	20	9	5	3	2	2	2	2	2
7	1	1	1	1	1	1	1	1	1	2	2	2	2	2	2	2	2	9	20	39	20	9	5	3	2	2	2	2	2	2
8	1	1	1	1	1	1	1	2	2	2	2	2	2	2	3	5	9	20	39	20	9	5	3	2	2	2	2	2	2	2
9	1	1	1	1	1	1	1	2	2	2	2	2	3	3	5	9	20	39	20	9	5	3	2	2	2	2	2	2	2	2
10	1	1	1	1	1	1	2	2	2	3	3	3	3	5	9	20	39	20	9	5	3	2	2	2	2	1	1	1	1	1
11	1	2	1	2	2	2	2	2	2	3	5	5	5	9	20	39	20	9	5	3	2	2	2	2	1	1	1	1	1	1
12	1	2	1	2	2	2	2	3	3	5	9	9	9	20	39	20	9	5	3	2	2	2	2	1	1	1	1	1	1	1
13	1	2	1	2	2	2	2	3	5	9	20	20	20	39	20	9	5	3	2	2	2	2	1	1	1	1	1	1	1	1
14	1	2	2	2	2	3	3	5	9	20	39	39	39	20	9	5	3	2	2	2	2	1	1	1	1	1	1	1	1	1
15	1	2	2	2	2	3	5	9	20	39	20	20	20	9	5	3	2	2	2	2	2	1	1	1	1	1	1	1	1	1
16	2	2	2	3	3	5	9	20	39	20	9	9	9	5	3	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1
17	2	2	2	3	5	9	20	39	20	9	5	5	5	3	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1
18	2	3	3	5	9	20	39	20	9	5	3	3	3	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1
19	2	3	5	9	20	39	20	9	5	3	2	3	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1
20	3	5	9	20	39	20	9	5	3	3	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
21	5	9	20	39	20	9	5	3	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
22	9	20	39	20	9	5	3	3	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
23	20	39	20	9	5	3	2	2	2	2	1	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
24	39	20	9	5	3	3	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
25	20	9	5	3	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
26	9	5	3	3	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
27	5	3	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
28	3	3	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
29	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
30	2	2	2	2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1

Figure 12.55 Payoff matrix for type y players in the substitutes case for the post-shock period.

Source: Fehr and Tyran (2002) with permission of The Econometric Society.

point or *anchor* in the post-shock phase. This could prevent subjects from adequately adjusting their price in the post-shock phase. This constitutes one possible behavioral explanation for nominal rigidity that is typically not considered in mainstream macroeconomics.

A second possible explanation for nominal rigidity is *money illusion*. An increasing amount of evidence shows that money illusion is important in a range of contexts.¹⁰⁶ In the context of the current experiments, money illusion implies that subjects could be more sensitive to changes in nominal payoffs $V_i = P v_i$ rather than real payoffs v_i . Suppose that P and M increase in the same proportion. Then, using (12.42), v_i is unchanged. However, a player who has money illusion may notice that $V_i = P v_i$ has increased, so he would prefer the change. Conversely, after a fall in M , such an individual might adjust P_i insufficiently downwards, fearing a fall in nominal income and, thus, leading to nominal rigidity.

The restart phase (see Figure 12.54) is used to identify if money illusion and anchoring are important. In this phase, the post-shock phase is simply restarted once again. There is evidence that restarts of this sort are not neutral.¹⁰⁷ In the complements case, in the CT-restart-real treatment, subjects were given payoffs in real terms for all three phases (pre-shock, post-shock, and restart). The analogous treatment for substitutes is denoted by ST-restart-real. In the

¹⁰⁶ In a macroeconomic context, Fehr and Tyran (2001) have already shown the presence of money illusion in a price setting game. In a range of other contexts, a significant fraction of individuals demonstrate money illusion. See, for instance, Shafir et al. (1997), Kooreman et al. (2004), Fehr and Tyran (2005, 2007), and Cohen et al. (2005).

¹⁰⁷ In the context of public goods games, a simple restart of the experiment can elicit high levels of cooperation despite the fact that prior to the restart, the cooperation levels had dropped to almost zero, on account of free riding. See, for instance, Andreoni (1988) and Croson (2000).




Figure 12.56 Average nominal price across time in the complements and substitutes treatment.

Source: Fehr and Tyran (2008) with permission of The Econometric Society.

CT-restart-nominal treatment, subjects are in the complements case and they receive real payoffs in the first two phases (pre-shock, post-shock), but nominal payoffs in the third phase (restart). The corresponding treatment with substitutes is called as ST-restart-nominal.

The results are as follows.

1. Across both types of subjects, x, y , the predicted average equilibrium price for complements and substitutes in the pre-shock phase (first 15 periods) is 25, while in the post-shock phase (periods 16 through 30), it is 12.5. Figure 12.56 shows the pre- and post-shock prices observed in the experiments across the complements treatment (CT) and the substitutes treatment (ST). In the pre-shock phase, in both treatments, and as expected (see the discussion above), price quickly converges to the equilibrium. In the post-shock phase, however, actual prices in both treatments are off (and typically above) the equilibrium price. In the complements case, the actual price is significantly above equilibrium and takes relatively longer to approach it, while in the substitutes treatment prices slightly overshoot the equilibrium price.¹⁰⁸ In each case, there is evidence of nominal rigidity of prices.¹⁰⁹ The strategic environment (complements or substitutes) seems to have an important bearing on the results on nominal rigidity.

¹⁰⁸ In particular, in the CT treatment, for the first eight periods of the post-shock phase, one can reject, at 1%, the hypothesis that the equilibrium prices and the observed average prices are equal. By contrast, in the case of the ST treatment, one cannot statistically reject this hypothesis.

¹⁰⁹ Field evidence on price stickiness is harder to find. Sellers might change a non-price dimension (quantity discounts) or offer underhand discounts, or sales, which might not be reflected in the data from catalogue prices that are often used to check for evidence on price stickiness. For a discussion of these sorts of issues see Blinder et al. (1998). Some of the empirical evidence is also discussed in Romer (2006).




Figure 12.57 Actual average prices and average best replies in the complements and substitutes treatments (Periods 16–18).

Source: Fehr and Tyran (2008) with permission of The Econometric Society.

- The individual choices for type x subjects in period 16, the period in which the anticipated monetary shock takes place, are as follows. In the substitutes case, 72% of the subjects jump immediately to the new equilibrium. But in the complements case, only 25% manage to do so, while 50% choose a price significantly higher than the equilibrium price. The choices of type y subjects are qualitatively similar.
- Subjects were asked about their expectations of the average price chosen by the other players. Thus, one can directly check from the payoff tables if subjects played their best response to the expectations that they held. Subjects played a best response in their nominal price, P_i , conditional on their expectations of the average price, $\bar{P}_{-i}^e$, in the complements treatment and the substitutes treatment (see Figure 12.57).
Thus, the presence of nominal rigidities in the post-shock phase, particularly in the complements treatment, cannot be ascribed to deviations from the best response by players. Rather the differences would then seem to arise from incorrect expectations, $\bar{P}_{-i}^e$. The experimental data indicates that in the pre-shock phase, $\bar{P}_{-i}^e$ is close to the actual price level. However, in the post-shock phase, price expectations are very sticky in the complements treatment, while they align very closely to the equilibrium expectations in the substitutes treatment.
- So what accounts for nominal rigidities? Consider the two treatments CT-restart-real and ST-restart-real. In these treatments, the subjects see payoff tables in real terms in the pre-shock and the post-shock phase. So money illusion is absent. Since there is no need to adjust prices, anchoring is also absent in these treatments. Most subjects conform to the equilibrium predictions towards the end of the post-shock phase in both treatments (in the last ten periods). If the restart-phase results in disequilibrium in this treatment, then factors other than money illusion and anchoring are likely to be contributing to nominal rigidities. However, in the restart-phase, actual and expected prices are instantly in equilibrium. This suggests that, in these experiments, there are no factors other than money illusion and anchoring which could account for nominal rigidities.
- Consider the two treatments CT-restart-nominal and ST-restart-nominal. In these treatments, the subjects were given payoffs in real terms in the pre-shock and the post-shock

phases. In the restart phase they were, however, given payoffs in nominal terms. In the restart phase, the price expectations of subjects are above equilibrium in the complements and the substitutes treatments despite the fact that the same equilibrium should prevail in the restart phase as in the post-shock phase. This is even more surprising because in the last ten periods of the post-shock phase, subjects played out the expected equilibrium. Thus, simply restarting the experiment in terms of nominal payoffs inflates price expectations and actual price levels. This is downward nominal rigidity driven by money illusion. Money illusion was seen to have a much larger effect in the complements case relative to the substitutes case.

6. Consider the treatment CT-real. In this treatment, the payoffs are in real terms in the pre-shock and the post-shock phases. Thus, there is no role for money illusion. There is a statistically significant deviation from equilibrium play in the first four rounds of the post-shock phase. There is no such deviation in the case of the ST-real treatment. Thus, anchoring (to the pre-shock equilibrium) plays a much stronger role in the complements as compared to the substitutes case.
7. Why is equilibrium reached more quickly in the substitutes treatment relative to the complements treatment? Haltiwanger and Waldman (1989), who discussed similar issues, proposed a solution that involved a joint consideration of adaptive and rational players. In order to test this, Fehr and Tyran (2008) ran some simulations. Adaptive players, in these simulations, simply expect the price in the previous period to prevail (perhaps due to money illusion or anchoring). Rational players know the fraction of adaptive players. However, the Haltiwanger and Waldman (1989) explanation is not able to fully explain the experimental findings (although it explains some aspects of it¹¹⁰). One problem is that the same mix of adaptive and rational players that does reasonably well in one treatment does not do well in the other.
8. The last observation in the previous point raises an intriguing possibility. What if the degree of bounded rationality is endogenous and the mix of adaptive and rational players in the two treatments (complements and substitutes) are endogenously distinct? The data suggests that in the complements treatment, the percentage of players who display behavior “close” to adaptive expectations is 45% for type x players and 60% for type y players. In the substitutes treatment, on the other hand, these percentages are, respectively, 6% and 0%. These results are consistent with the view that bounded rationality is relatively more pervasive under complements. Whilst the equilibria are identical in the complements and substitutes treatments, the main difference between the two is the slope of the best reply function in $(P_i, \bar{P}_{-i})$ space. This causes the loss in payoffs from adaptive expectations in the substitutes treatment to be higher than the complements treatment. By implication, subjects in the substitutes treatment stand to gain more by moving from adaptive to rational expectations. Thus, the extent of bounded rationality seems to respond to the cost of making cognitive errors.
9. The implications for existing macroeconomic models of nominal rigidity are profound. Models that invoke informational or contractual frictions or models of purely statistical learning,¹¹¹ do not invoke the cognitive costs of making decisions. Thus, they seem to be

¹¹⁰ For instance, it can explain the greater nominal rigidity in the complements treatment relative to the substitutes treatment.

¹¹¹ For models in the tradition of information asymmetries see Lucas (1972). The new Keynesian literature on contractual asymmetries began with the work of Fischer (1997). For a discussion of statistical learning models, see Evans and Honkapojha (2001).

missing a fundamental explanation of such rigidities based on money illusion and anchoring. Indeed, traditional macroeconomic models are unable to explain these experimental findings. Existing macroeconomic models in which there is a small degree of bounded rationality, such as Akerlof and Yellen (1985), and models of sticky information, such as Mankiw and Reis (2002), are also not fully able to explain the facts.

These comments should give the existing macroeconomic theories a sobering cause for concern. However, the model of Reis (2006), which incorporates a cost of searching for and absorbing information, and of Evans and Ramey (1992, 1998), Brock and Hommes (1997), which relate the formation of rational expectations to cognitive limitations, would seem to be promising in explaining the experimental data.

12.9 A digression on competitive equilibrium experiments

A discussion of the experiments on competitive equilibrium (CE) may appear to be out of place in this chapter. However, there are three justifications for this section. First, the CE is often a limiting case, as the number of participants in strategic situations become very large (e.g., the solution under Cournot competition when the number of firms is large). Second, the CE is often used as a benchmark for comparing the outcomes in many strategic market situations. Third, and intriguingly, the double auction experiments that are used to test the predictions of a CE often claim that the results approximate a CE, even with a very small number of participants (as low as even three); conditions that one may associate with strategic situations.

Economists have developed an elegant theory of *competitive equilibrium*. A vector of the prices of all inputs and all outputs, p^* , is a CE if, given these prices, all economic agents maximize utility/profits and all markets clear. With appropriate dating of goods, this static version can be easily extended to problems that have a time dimension. The existence and uniqueness of such a price vector, p^* , requires strong assumptions (Varian, 1992; Mas-Collel et al., 1995).

An indispensable part of the theory of CE is a fictitious actor, the *Walrasian auctioneer*, who contemplates fictitious out-of-equilibrium trades and then arrives at equilibrium prices. Actual trading then takes place *only* at the equilibrium prices.¹¹² One may imagine that the Walrasian auctioneer starts out with an initial price vector, then invites market participants to truthfully report their intended trades in all markets at that price. If all the excess demands are zero at that initial price, the initial price is the CE price vector, p^* . If the excess demands in some markets are positive (negative), then the auctioneer increases (decreases) the price in those markets. This experimental process goes on until the auctioneer finds a price vector, p^* , at which all excess demands are zero. This is the CE price vector.

Global uniqueness of p^* requires, as a sufficient condition, that all goods be gross substitutes at all prices.¹¹³ Furthermore, we require a large number of atomistic buyers and sellers, none of whom are large enough to influence the actual market outcome. Thus, all individuals are *price takers*. In the classical model, buyers, sellers, and the Walrasian auctioneer, have full information about all relevant aspects. However, because each participant has no effect on the actual outcome, none of the participants can use their information in any strategic manner. Furthermore,

¹¹² In particular, no trades are allowed to take place at the intermediate trial prices that the auctioneer may have experimented and/or iterated with.

¹¹³ Two goods are gross substitutes if an increase in the price of one good leads to an increase in the excess demand for the other.

we require that there be no market frictions or market imperfections, no externalities, and no public goods.

The assumptions behind a CE are unlikely to be met in actual practice. Hence, it is important to test the predictions of the theory in the lab. Edward Chamberlin (1948), one of the pioneers of monopolistic competition, performed a series of seminal experiments to test for a CE. Experimental data gathered from Harvard economics students to test the theory was not even remotely close to the predictions. Vernon Smith (1962) recast these experiments in terms of a double auction format and found support for a CE. Hundreds of experiments using double auctions followed, and found support for a CE.

Gode and Sunder (1993) used the format of the double auction experiments with zero intelligence (ZI) robots, and confirmed the predictions of the CE. In more complex settings, as we shall see, ZI robots do not do as well as humans. The positive results for a CE using ZI robots have raised some fundamental questions. What degree of rationality does the market require on the part of the market participants? Does individual irrationality wash out in the aggregate? Do behavioral biases have any bite in a competitive market equilibrium?

Davis and Holt (1993, p. 167) echo the mainstream view on double auction CE experiments as follows: “In markets organized under double auction trading rules, the predictions of the competitive model appear to be robust . . . Competitive price theory also does a good job of organizing data in some enriched double auction market structures.” In a similar vein, Kagel and Roth (1995, p. 370) write: “The striking competitive tendency of the double auction institution, which has been confirmed by at least a thousand market sessions in a variety of designs, indicates that neither complete information nor large numbers of traders is a necessary condition for convergence to CE outcomes.”

To some readers, the experiments of Chamberlin (1948) that allow buyers and sellers to freely meet and negotiate prices and quantities in an unstructured manner may seem closer to their own visualization of competitive markets. We discuss these issues below. Given excellent survey articles on the subject,¹¹⁴ the aim below is to consider some selected, but seminal, papers only.

12.9.1 *The original market experiments by Chamberlin (1948)*

Chamberlin (1948) performed a series of seminal market experiments with economics students at Harvard. He divided students into an equal number of buyers (B) and sellers (S). Sellers can sell a single unit of a good and buyers have unit demands. Buyers, $j = 1, 2, \dots$ are given cards with a number B_j that equals the buyers' valuation for a single unit. Sellers, $i = 1, 2, \dots$ are given cards marked with a number S_i that equals the cost to the seller of supplying the good. Each buyer and seller has private information about their valuation or cost. The expectation is that buyer j will be willing to pay a price no higher than B_j and seller i will not be willing to accept any price lower than S_i .

In contrast to later experimental practice, Chamberlin allowed the subjects in experiments, in the spirit of actual markets, to mingle with each other and to freely strike bargains with each other. If buyer j encounters a seller i and $S_i < B_j$, then the pair has an opportunity to strike a mutually profitable deal. For instance, if they agree on a price $p : S_i \leq p \leq B_j$, then each gets a non-negative payoff (the buyer's payoff is $B_j - p$ and the seller's payoff is $p - S_i$). The market closes after a certain amount of time. The same experimental market is repeated several times.

¹¹⁴ For the details, the interested reader can look up Davis and Holt (1993) and Volume 1 of the *Handbook of Experimental Economics Results* by Plott and Smith (2008).

Table 12.34 Buyers' valuations, sellers' costs, and market data.

Valuations and costs assigned to experimental subjects			Actual data on market trading in an experimental session			
Index	Buyer's valuation	Seller's cost	Number	Buyer's valuation	Seller's cost	Price
1	104	18	1	56	18	55
2	102	20	2	54	26	40
3	94	26	3	72	30	50
4	90	28	4	84	34	45
5	86	30	5	44	44	44
6	84	32	6	102	42	42
7	82	34	7	80	20	40
8	80	36	8	60	28	55
9	76	40	9	48	40	45
10	74	42	10	76	36	45
11	72	44	11	94	52	55
12	68	46	12	68	58	62
13	66	50	13	66	46	55
14	60	52	14	82	32	58
15	58	54	15	90	72	72
16	56	58	16	104	54	54
17	54	62	17	52	50	50
18	52	64	18	86	64	64
19	50	66	19	74	62	69
20	48	68	Participants unable to trade			
21	44	70	20	38	68	
22	38	72	21	50	66	
23	34	74	22	28	82	
24	32	78	23	32	88	
25	30	80	24	18	90	
26	28	82	25	26	84	
27	26	84	26	22	104	
28	24	88	27	24	78	
29	22	90	28	30	80	
30	20	98	29	20	98	
31	18	104	30	34	74	
			31	58	70	

Source: Chamberlin (1948).

The experimental data are shown in Table 12.34. The left hand side panel in Table 12.34 shows the valuations of buyers, arranged in decreasing order, and the cost of sellers, arranged in increasing order. These data are plotted in Figure 12.58. The demand curve, DD, reflects the valuations of the buyers and the supply curve, SS, reflects the costs of sellers. The zig-zag shape arises because sellers have unit supplies and buyers have unit demands. The equilibrium quantity is shown by




Figure 12.58 A plot of the demand and supply curves in Table 12.34.

q^* . The equilibrium price is indeterminate over a range of values. One possible equilibrium price, p^* , is shown in the figure.

From the left hand side panel in Table 12.34, for the first 15 units, shown in a darker background, the cost is lower than the valuation, so there is a possibility of mutually beneficial trades. The prediction of the CE model is $q^* = 15$, and the equilibrium price p^* could be any number between 56 and 58. For the 15th unit, bargaining between the buyer and the seller would ensure a price $p \in \{54, 55, 56, 57, 58\}$. However, given the unit demands and unit supplies, it is not efficient for the buyer with valuation $B_{16} = 56$ to buy. Hence, in equilibrium, $p^* \in \{56, 57, 58\}$.

In Chamberlin's experimental design, subjects mix freely in a room and strike bargains, much as we would expect in many real-world markets. This raises some interesting issues. Consider sellers with costs S_k, S_m and buyers with valuations B_j, B_n , as shown in Figure 12.58. Suppose that the seller with cost S_k (who is not expected to trade in standard equilibrium theory¹¹⁵) meets the buyer with valuation B_j (who is expected to trade in standard equilibrium theory). The buyer's valuation is greater than the seller's cost, so there is scope for mutually advantageous trade. The buyer does not know the valuations of other buyers or the costs of other sellers. If he turns down this opportunity to trade, he might not get another opportunity, thus, he might be willing to strike a bargain with the seller. In an alternative case, if the seller with cost S_m happens to meet a buyer with valuation B_n (who is not supposed to trade in equilibrium theory), then he might strike a bargain for similar reasons. This intuition turns out to be correct as we see below.

¹¹⁵ Such sellers are sometimes referred to as *extramarginal sellers*. Similarly, buyers whose valuations are lower than the competitive equilibrium price are sometimes known as *extramarginal buyers*. The converse of such traders (with costs lower than and valuations higher than the competitive equilibrium price) are sometimes known as *inframarginal traders*.




Figure 12.59 The time path of prices over successive rounds of the experiment.

Source: *Journal of Political Economy*. Chamberlin, Edward H. (1948). "An experimental imperfect market." 56: 95–108. © 1948 University of Chicago Press.

Consider the experimental findings reported in the upper right hand side panel of Table 12.34 from one of the timed sessions. This data describes the buyers' valuations and the sellers' costs from actual trades that took place in the session. Sellers who are in the position of a seller with cost S_k in Figure 12.58 are sellers with costs 58, 62, 64, 72. Buyers who are in the position of a buyer with valuation B_n in Figure 12.58 are buyers with valuations 56, 54, 48, 42. We should not observe any of these market participants trading if the predictions of the CE model were to hold. Yet, we do observe that they trade.

Over 46 repetitions of the experiment, the traded quantity is never below the CE quantity. Further, the actual quantity traded is above the equilibrium quantity in 91.3% of the cases. By contrast, the average price is higher than the equilibrium price 7 times, and lower 39 times.

The typical experimental results are shown in Figure 12.59; the equilibrium price is shown at 57. The data in each round of the experiment is plotted as the curve labeled "Actual Prices." It is clear that there is no tendency for prices to move towards the competitive outcome. Even in the later periods, no apparent convergence is obvious.¹¹⁶ Readers may wish to ponder the following quote, in full, from Chamberlin (1948, p. 102).

My own skepticism as to why actual prices should in any literal sense tend toward equilibrium during the course of a market has been increased not so much by the actual data of the experiment before us which are certainly open to limitations as by failure, upon reflection stimulated by the problem, to find any reason why it should be so. It would appear that, in asserting such a tendency, economists may have been led unconsciously to share their unique knowledge of the equilibrium point with their theoretical creatures, the buyers and sellers, who, of course, in real life have no knowledge of it whatsoever.

Chamberlin (1948), like others after him, allows "out of equilibrium trading," although this is not allowed in theory. However, out of equilibrium trading is needed to substitute for the fictitious

¹¹⁶ As Chamberlin (1948, p.101) notes: "Among the other forty-five examples, the most diverse patterns appear, with no apparently predominant tendencies to be noted."

actor, in the model, the Walrasian auctioneer. The presence of this actor makes theory very difficult to test. Chamberlin's implementation of the market mechanism seems eminently plausible and natural, although it was abandoned.

It is still nevertheless useful to see if one could design institutions that mimic the CE outcome. We consider this next, although we should also test if the proposed institutions are widely observed or not.

12.9.2 Double auction tests of a competitive equilibrium

In a series of papers, beginning with Smith (1962), Vernon Smith used an alternative implementation of a CE, one that used *double auctions* instead. Unlike Chamberlin's setup in which subjects in experiments freely met in order to strike mutually advantageous deals as in most real markets, Smith's subjects only make public offers/asks (in their role as sellers) or public bids (in their role as buyers).

Consider m sellers and n buyers with unit supplies and unit demands. The costs of sellers and the valuations of buyers are not typically public knowledge in double auction experiments. Suppose that one of the sellers moves first and publicly calls out an asking price (or "ask"), a_1 . This can be immediately accepted by a buyer, in which case both parties who are involved in the exchange exit the market and have no further role to play. The remaining parties bid/ask anew. Alternatively, one of the buyers might choose a bid, b_1 . The bid can be immediately accepted by a seller, in which case both parties exit the game and the remaining parties bid/ask anew. If not, then one of the sellers makes another ask, a_2 . This can be followed by fresh rounds of bids and asks, and so on until the time runs out in the experiment.

Proceeding in this manner, we obtain a sequence of asks $a_1, a_2, a_3, \dots$ and a sequence of bids $b_1, b_2, b_3, \dots$ until an object is exchanged. If an object is exchanged at an agreed price p between a buyer who has a valuation v and a seller who has a cost c , then their respective payoffs are $v - p$ and $p - c$. Multilateral public bids distinguish the double auction method from the bilateral private bargaining between individuals on both sides of the market in Chamberlin (1948).

An important component of the double auction experiments is *the market improvement rule*.

Definition 12.6 (*Market improvement rule: Davis and Holt, 1993, p. 128; Kagel and Roth, 1995, p. 369*) *In most auction experiments, non-improving bids/asks are not allowed. So for the sale of any one unit of an object, it is required that $a_1 > a_2 > a_3 > \dots$ and $b_1 < b_2 < b_3 < \dots$*

From Definition 12.6 the asking price of sellers needs to drop over successive asks and the bidding price of the buyers needs to increase over successive bids until an exchange takes place. Table 12.35 gives an illustration of the operation of a double auction during the sale of one unit of a good, taken from Davis and Holt (1993, p. 129). There are four buyers, B1 through B4, and four sellers, S1 through S4. The entire experiment lasts 300 seconds and time is given in seconds in the first column. Seller S2 first announces an ask of \$5, followed by a bid from buyer B4 of \$4.10. The market improvement rule prohibits any subsequent asks above \$5 and any bids below \$4.10. As shown in Table 12.35, subsequent sellers offer decreasing asks and subsequent buyers make increasing bids until, finally, buyer B4 accepts the offer of seller S1 to trade one unit at a price of \$4.30.

The market improvement rule is intended to substitute for a Walrasian auctioneer and it is widely used in double auction experiments with humans and non-human subjects, as in agent-based models. A skeptic may argue that the theory makes no allowance for such a rule and it

Table 12.35 Some data from a double auction experiment reported in Davis and Holt (1993, p. 129).

Time	Bid	Offer
296		S2 offers \$5.00
294	B4 bids \$4.10	
293		S3 offers \$4.50
291	B1 bids \$4.20	
285	B2 bids \$4.21	
284		S1 offers \$4.40
279		S3 offers \$4.39
276		S4 offers \$4.35
271	B2 bids \$4.25	
267	B2 bids \$4.26	
265		S3 offers \$4.34
261		S1 offers \$4.30
254	B4 accepts	

forces trades in the direction of a CE, which is the very phenomenon that it seeks to explain. However, many in the profession consider the results from the double auction experiments to be of fundamental importance.¹¹⁷

In his pioneering work, Smith (1962) used the double auction method, described above, to test for a CE. Each *market period* ends when bids and asks no longer lead to further successful transactions between buyers and sellers. At the end of each market period, say period n , the market is immediately closed and reopened for the next day, which is referred to as period $n + 1$. Each market participant has private information about his/her valuation or cost for the good. A variety of demand and supply curves are considered. The results from one of the 10 “tests,” which corresponds to one such set of demand and supply curves, is shown in Figure 12.60.¹¹⁸ In any test, the demand and supply conditions are held fixed, so that subjects are able to “learn” their way to a CE; however, the supply and demand curves are altered from one test to another.

The demand and supply schedules are shown in the left panel of Figure 12.60. The right panel shows the dynamics of the actual path of prices over the various periods. Buyers have unit demands, given to them on a piece of card. The buyer with the highest valuation in Figure 12.60 has a valuation of \$3.25, and the next highest valuation is \$3.00. So if the price were \$3.00, then, at most, two units would be demanded. These valuations are used to generate the demand curve

¹¹⁷ The following quote from Eckel (2007) is not unrepresentative: “When Vernon tells about this experiment, he maintains he was astonished at the rapid convergence to CE that he observed, and says he had to replicate it a few times, with different shaped demand and supply arrangements, to convince himself it was correct. Experimentalists now know that the institution he implemented, the Oral Double Auction, is about the most powerful institution there could be for inducing convergence to CE. It works like a charm, with groups as small as 3 buyers and sellers, even with low or hypothetical stakes, it reliably converges to equilibrium. Every experimentalist I know conducts this experiment in classrooms in almost every class they teach. It is an unforgettable experience for students. They come away believing that markets (can) work, and that the theory of supply and demand has teeth (or jaws, as Plott would say).”

¹¹⁸ The word “test” in Smith (1962) seems to correspond to an experimental session and each test is comprised of several periods (see Figure 12.60).




Figure 12.60 Price dynamics.

Source: *Journal of Political Economy*. Smith, V. L. (1962). "An experimental study of competitive market behavior." 70: 111–37. © 1962 University of Chicago Press.

DD. Similarly, sellers are given cards that give the minimum price at which they should be willing to sell a unit, possibly reflecting their costs. This generates the supply curve SS. The minimum price that sellers are willing to accept and the valuation of the buyers is private information. The number of trading periods varies from one experiment to the other. The double auction procedure, described above, is then implemented. The theoretical prediction is an equilibrium price of \$2 and the equilibrium quantity is 6 units.

In Figure 12.60, in the first of five periods, five successful transactions took place, whereas in period 4, seven successful transactions took place. The *parameter of convergence*, α , is defined as $\alpha = \frac{\sigma}{p^*} \times 100$, where σ is the standard deviation of the actual experimental price and p^* is the predicted equilibrium price. The variation in prices, as measured by α , tends to fall in successive periods. While α falls in each period, except for the last few transactions in period 5, a simple eyeball test does not conclusively establish that prices have converged to the equilibrium level p^* . Recall that the various tests differ in the demand/supply curves. In Test 4, for instance, the supply curve is perfectly elastic. In this case, the actual price converges to about \$0.20 above p^* .

The results in Smith (1962) from a range of demand and supply curves are consistent with the competitive market equilibrium. However, one of the predictions of the CE model, the *Walrasian hypothesis*, does not hold.¹¹⁹ Hundreds of subsequent studies that have been done with the basic double auction template provided by Smith (1962) have found similar results.¹²⁰ Features

¹¹⁹ The Walrasian hypothesis implies that the unit price is increasing in excess demands; see Smith (1962, p.116).

¹²⁰ The interested reader may consult the survey by Davis and Holt (1993) for the relevant evidence. Smith (1964) extends his earlier work by looking at prices in various rounds of the experiments when the market institutions are altered. For instance, only sellers may be permitted to announce asks with the buyers passively accepting or rejecting; or the sellers may be permitted to announce asks and the buyers could simultaneously announce bids; or only the buyers may announce bids and the sellers passively accept or reject bids.

of the transition dynamics that lead to the eventual convergence to a CE are still an active area of research (Plott, 2008).

One method of evaluating the efficiency of the observed outcome is to compute a measure of allocative efficiency, E . Define total social surplus as the sum of consumer surplus and producer surplus. Let S_P denote the predicted social surplus under the competitive outcome and S_A the actual social surplus in experimental conditions. Then the allocative efficiency of the observed outcome is $E = (S_A/S_P) \times 100$. The allocative efficiency in double auction experiments with human subjects is very high, typically ranging between 97.3% and 99.4%. The interested reader can see a summary of several studies in Davis and Holt (1993, Table 3.4, p. 136).

List (2004) constructs an experimental design that has shared features with Chamberlin (1948) and Smith (1962). In an artefactual field experiment, he takes experienced traders from the sports cards and collector pin markets. The results are similar, so we consider only sports card dealers. Such traders are either buyers (non-dealers) or sellers (dealers) in their respective markets; these endogenously determined real-life roles are maintained in the artefactual experiments. As in Smith (1962), buyers are given cards with their assigned valuations and sellers are given cards with their assigned marginal costs. Buyers do not find it profitable to pay more than their valuations and sellers do not find it profitable to sell below their marginal costs. The valuations of buyers and the costs of sellers are private information. Typically, there are 12 buyers and 12 sellers in each experimental session that is comprised of five market periods, each period lasting for 10 minutes.

As in Chamberlin (1948), buyers and sellers are allowed to engage in privately optimal exchanges by meeting, bargaining and haggling. The homogeneous commodity that is used in the experiments is multiple copies of a sports card that had no value outside the experiments. A critical feature of the experimental design is in the spirit of Smith (1962) but distinguishes these experiments from the experiments in Chamberlin (1948). Namely, as soon as a buyer and a seller exchanged an object, the price and the contract was publicly announced to all other individuals in the session. It is likely that the most recent exchange price acts as an anchor around which bidding in subsequent rounds takes place.¹²¹ Also, unlike Chamberlin (1948), discussions among groups of sellers and groups of buyers that could lead to a collusive outcome were prohibited.

Three kinds of market demand and supply curves were considered; in each, the equilibrium price was \$13.50 and the equilibrium quantity was 7. In the S12 treatment, the demand and supply curves are symmetric; in the PED12 treatment, the demand curve is perfectly elastic at a price of \$13.50; and in the PES12 treatment, the supply curve is perfectly elastic at a price of \$13.50. The experimental outcomes quickly converge to the predictions of the competitive market outcomes. In the fifth round, in the treatments S12, PED12, PES12, and averaged across all sessions, the respective prices are \$13.12, \$13.15, and \$13.68; the corresponding quantities are 7.3, 6.3, and 6.3.

These convergence results are in stark contrast to the results in Chamberlin (1948) that show no tendency for prices to converge to the competitive equilibrium prices. List (2004) ascribes these results to the role of experience among traders. The more experienced the traders, the greater the rents that they earn in the experimental exchanges, conditional on trade taking place. By contrast, when the experiments are run with children, convergence to a competitive equilibrium does not take place. These conclusions are open to alternative interpretations. However, the value of this study is that it demonstrates that a double auction is just one method of achieving consistency with a competitive outcome. A second method is the provision of information to each participant following each successful exchange in these experiments.

¹²¹ For details on the *anchoring heuristic*, see Chapter 19.

12.9.3 Agent-based models of competitive double auction markets

There has been growing interest in *agent-based computational economics* (ACE). Duffy (2006, p. 951) provides a useful set of characteristics of this approach: “agent-based computational economic (ACE) researchers consider decentralized, dynamic environments with populations of evolving, heterogeneous, boundedly rational agents who interact with one another, typically locally. These models do not give rise to closed-form solutions, and so results are obtained using simulations. ACE researchers are interested in the aggregate outcomes or norms of behavior that emerge and are sustained over time, as the artificial agents make decisions and react to the consequences of those decisions.”

An important application of ACE to double auctions is provided by Gode and Sunder (1993).¹²² They were interested in finding if the convergence properties of double auction experiments on CE with human subjects could also be obtained with non-human subjects, such as zero intelligence (ZI) robots. The conjecture was that the rules of the double auction experiments, such as the market improvement rule, and not the rationality of human subjects, is responsible for the supposed convergence results with human subjects.

ZI robots, unlike human subjects, bid randomly from an appropriate distribution of asks or bids. A robot buyer, j , with valuation B_j is allowed to bid randomly, any number in the interval $[0, B_j]$. A robot seller, k , with cost S_k , asks randomly any price in the interval $[S_k, m]$, where m is the maximum over all buyer valuations. Crucially, the ZI robots follow the market improvement rule given in Definition 12.6, which plays an important role in the observed success of the double auction market experiments. ZI robots, defined in this way, are said to respect the budget constraint, because they do not trade at a loss. We call these constrained robots ZI-C traders. In a control treatment, all buyer and seller robots could bid randomly from the interval $[0, m]$. Since, these traders do not respect the budget constraint and will sometimes make losses, they are called the ZI-U traders.

The price dynamics are shown in Figure 12.61. The left panel shows the demand and supply schedules with a CE price prediction of 80 and a quantity prediction of 24. The right panel shows the results over six periods of the experiment with ZI-U and ZI-C robots, and human subjects; each period has several market rounds. Human subjects stay fairly close to the predictions of the theory. The ZI-C subjects show greater fluctuation around the equilibrium prediction than human subjects, however, the variation in prices declines over successive market periods.

One has to be careful in interpreting this decreased variation in prices for ZI-C subjects, because, by definition, robot subjects do not engage in learning. The reason for this reduction in variance is that the higher valuation/lower cost units tend to be traded first, which is a feature that is in-built into the market improvement rule. The greatest variation is found in the case of ZI-U robot traders who show no tendency to converge to the equilibrium predictions. The allocative efficiency, E (defined above), for ZI-C robots is 98.7%, slightly higher than the corresponding figure for human subjects of 97.6%.¹²³

¹²² The precursor to the Gode and Sunder (1993) paper is the paper by Hurwicz et al. (1975). The interested reader can explore this literature further by consulting Tesfatsion and Judd (2006).

¹²³ Readers who wish to follow up the extensive literature may consult Duffy (2006).




Figure 12.61 Price dynamics with ZI-C and ZI-U robot traders and humans.

Source: *Journal of Political Economy*. Gode, Dhananjay K. and Shyam Sunder (1993). "Allocative efficiency of markets with zero intelligence traders: market as a partial substitute for individual rationality." 101: 119–37. © 1993 University of Chicago Press.

12.9.4 *Evaluating the experimental evidence on competitive markets*

It is not possible to observe the behavior of the fictitious player in economic theory, the Walrasian auctioneer, who is critical in achieving a CE. This makes it difficult to falsify the theory. Experimenters are left with the open question of what proxy to use for the Walrasian auctioneer in their experiments. Chamberlin (1948) allowed buyers and sellers to freely bargain with each other and agree on mutually beneficial trades. Smith (1962) chose a double auction setting in which the bilateral dealings in Chamberlin are replaced with public bids and asks. Chamberlin's results reject the Walrasian competitive outcome, while Smith's results support it. The subsequent literature has followed Smith's double auction experimental method.¹²⁴ Furthermore, a CE requires

¹²⁴ For instance, Miller et al. (1977) extend Smith's (1962) work to an intertemporal, two season market. On p. 299, they open with the following quote: "However, it is now well established by many experimental replications that knowledge of market supply and demand is not a necessary and may not be a sufficient condition for classical

that there should be no “out of equilibrium trading.” All trades should only be carried out at the equilibrium price which is arrived at by the Walrasian auctioneer. However, in all experiments, out of equilibrium trading is crucial and essential.

One interpretation of the experimental literature in this area is that its main interest lies in discovering if there exist institutions that mimic the outcome of CE. This is of interest in its own right, insofar as a CE possesses desirable efficiency properties. The double auction setting, augmented with the market improvement rule, is one such institution. Let us call this the *augmented double auction institution*.

The augmented double auction institution achieves nearly 100% allocative efficiency. However, despite its similarity to trading in stock markets, it begs the question of why do we not observe it more often in practice. For instance, regulatory and antitrust agencies desiring a competitive, efficient outcome should simply announce and implement the augmented double auction for all cases where they suspect a non-competitive outcome. But they typically don't. One possible explanation is that there could be important transaction costs, or other important frictions associated with this institution that are not modeled in experiments. Furthermore, the auction or stock market trading situation need not be typical. For instance, there may be practical difficulties in bringing all participants in a transaction to the table every time an exchange takes place.

The market improvement rule seems critical to the success of the double auction setting, both in the human and the robot setting. By constraining movements in bids and asks (see description above), one may wonder if it nudges the results in the direction of the predictions of the competitive model. Indeed, zero intelligence robots also arrive at the competitive outcome when a market improvement rule is imposed on them. On the other hand, a strength of the augmented double auction results is that even with a minimum degree of rationality (as with zero intelligence robots) one typically obtains the desirable competitive outcome, which is a particular strength of this mechanism.

The static double auction setting is quite special. The equivalence of results generated with ZI robots and humans can break down in a more complex double auction setting. Furthermore, in both human and robot treatments, prices might not converge to the CE prices. Van Boening and Wilcox (1996) relaxed the assumption of fixed costs for sellers and introduced *avoidable costs*. These are costs that are actually incurred only if an exchange takes place. The seller's entry decision is more complicated here. They have to guess the identity of the other sellers who might also enter (because costs are avoidable if they defer entry). In this more complicated setting, the efficiency of the observed outcome for human subjects is significantly lower than the competitive outcome. Furthermore, the performance of ZI robots is even worse than human subjects.

Miller (2008) considers the performance of ZI robots in the two-season setup introduced by Miller et al. (1977). In these experiments, the subjects have to make an informed decision on how many units to carry forward from one season to the next, which requires a degree of cognitive sophistication. Predictably, ZI robots fail to converge to the CE. Brewer et al. (2002) consider the possibility that the demand and supply curves are *continuously refreshed*. By this it is meant that they are altered from one trading period to another. In this setting, ZI-C robots chose very variable prices that do not converge to the CE. These results caution against drawing the common conclusion that a CE can easily be obtained in lab experiments.

competitive equilibria to be achieved when a stationary market is organized under auction rules.” A similar report of the evidence is given in Davis and Holt (1993) and Kagel and Roth (1995).

CHAPTER 13

Models of Behavioral Game Theory

13.1 Introduction

Consider a perfect Bayesian Nash equilibrium (PBNE) in a dynamic, extensive form game.¹ Roughly speaking, a profile of strategies and beliefs is a PBNE if for each player at each decision node, two main conditions hold. First, conditional on what has happened so far in the game, players play a best response, conditional on their beliefs about other players and the subsequent equilibrium strategies of the players. Second, the beliefs of players about the strategies of other players are formed using Bayes' rule, and equilibrium actions and beliefs are consistent with each other. Simpler games, such as static games of full information, and static games of incomplete information, require appropriate simplifications of a PBNE, respectively, a Nash equilibrium and a Bayesian Nash equilibrium.

Chapter 12 showed that the Nash equilibrium concept and its refinements in more complex games fail to fully account for the experimental evidence, particularly from the early rounds of games. A natural response is to attempt to relax one or both conditions of a PBNE. In particular, this has led to the development of two main behavioral alternatives.

The first of these alternatives, *quantal response equilibrium* (QRE), is considered in Section 13.2. In a QRE, the first condition of a PBNE is relaxed. As in *stochastic utility models*, players play noisy best replies, conditional on their beliefs about the strategies or types of others. The noise in best replies may arise from privately observed payoff shocks or taste shocks. However, players correctly anticipate the noisy best replies of others. Noisy best replies do not imply random play. Rather, the actions of players that give a higher expected payoff are played with a higher probability. A QRE is a fixed point in the strategies and the beliefs of the players.

We consider QRE in normal form games (McKelvey and Palfrey, 1995), and, briefly, in extensive form games (McKelvey and Palfrey, 1998). In many cases, the predictions of QRE are in better accord with the findings of experimental games, as compared to a PBNE. In particular, a QRE can account for the effect of changes in the magnitude of payoffs on the chosen strategies when a Nash equilibrium predicts no such effect. We consider two successful empirical applications. The first is due to McKelvey et al. (2000) and a second is an application to Basu's traveller's dilemma game, due to Capra et al. (1999).

¹ See the appendix on game theory if you are rusty with this concept.

Haile et al. (2008) showed that a QRE does not impose falsifiable restrictions if the noise with which players play their strategies is not independently and identically distributed. We outline their objections, and the response of Goeree et al. (2005) who proposed the *regular quantal response equilibrium* (RQRE).

It would appear that the cognitive requirements for a QRE that require computing a fixed point in “actions” and “beliefs about the noisy play of others” are quite high. Relaxing the second feature of a PBNE is likely to be less cognitively taxing. In this case, players’ beliefs are not required to be consistent with the actions of other players, yet all players play a best response to their, possibly mistaken, beliefs. Within this class of models, *level- k models*, considered in Section 13.3, provide the leading behavioral alternative to a Nash equilibrium and its refinements. Level- k models were initially introduced by Stahl and Wilson (1994, 1995) and Nagel (1995) to explain certain experimental findings that were difficult to reconcile with existing theories.

In level- k models, players have a *hierarchy of beliefs*. Level-0 players (or type $L0$) are non-strategic, and in many applications they choose uniformly randomly over the set of available actions. However, when the underlying actions are framed in a non-neutral manner or have different degrees of salience, then this need not be the case. Level- k , $k \geq 1$ players (or types Lk , $k = 1, 2, \dots$) assume that all other players are of type $Lk - 1$. Type $L0$ players need not physically exist. At a minimum, they serve as a model of other players in the minds of type $L1$ players. Since each type best responds to the optimal actions of the immediately lower type in a recursive manner, type $L0$ is the *anchoring type* and the specification of its actions is often critical to the results. The beliefs of players are not assumed to be in equilibrium, so level- k models would seem to apply particularly to new and unfamiliar situations, to games that one plays infrequently, and particularly, to data from the early rounds of experiments. Arguably, the domain of such decisions is rich, varied, and important (e.g., marriage, divorce, purchase of consumer durables, house purchases, decisions on pension plans, and mortgage/insurance decisions).

The inability of level- k models to specify the optimal actions of type $L0$ players may be considered to be a drawback by some. However, this problem is no different from the specification of a reference point in prospect theory, or the specification of a prior in Bayesian models, or the specification of a utility function in a model. The assumptions that one makes about type $L0$ players can be tested for their validity, particularly in out of sample data. The evidence on the physical existence of type $L0$ players is mixed; Sections 13.4.5 and 13.4.6 specifically discuss this evidence, although further evidence is scattered in many places in the chapter.

Section 13.3.2 considers a well-known application of level- k models, the *p -beauty contest* (Nagel, 1995). In the p -beauty contest, participants simultaneously make a guess in an interval $[0, 100]$; the winning guess is closest to p times the average guess. When $p < 1$, the Nash equilibrium is the lower limit of that interval, 0, and when $p > 1$ the Nash equilibrium is the upper limit, 100. The level- k model makes very precise predictions for this game. Let $p = \frac{2}{3}$, and assume that type $L0$ players choose uniformly randomly over $[0, 100]$. Then, the level- k prediction is that we should observe spikes in guesses at the numbers $50(\frac{2}{3})^k$ for $k = 1, 2, 3, \dots$. The thrust of the experimental findings for data from early rounds of experiments is that one rarely observes the Nash outcome, and one does observe spikes in the data at the points predicted by the level- k model. These findings hold even for subjects whom one would expect to have high cognitive abilities.²

Section 13.3.3 considers the *cognitive hierarchy model* (CH), which is closely related to the level- k model (Camerer et al., 2004a). In CH models, each type Lk player, $k \geq 1$ assumes that there

² See for instance, Stahl and Wilson (1994, 1995), Nagel (1995), Ho et al. (1998), Bosh-Domènech et al. (2002), Camerer (2003a,b).

is a non-degenerate distribution of all lower types of players, $Lk - 1, Lk - 2, \dots, L0$. Players form the correct inferences about the distribution of lower types, normalized by the truncation of the distribution at level $Lk - 1$. Typically, the underlying distribution of players is a Poisson distribution, which is a one parameter family of distributions, allowing for a tractable estimation of such models. In a CH model, the rationality of the players increases as k increases. In most games, CH and level- k models make similar predictions. More work is needed to successfully discriminate between these two models (Crawford et al., 2013).

Section 13.4 considers further applications of level- k and CH models. When there are multiple equilibria in coordination games, classical game theory proposes several refinements, such as salience of labels, payoff dominance, and risk dominance, to explain the choice among multiple equilibria. Section 13.4.1 considers the salience of labels in 2×2 coordination games with multiple equilibria (Crawford et al., 2008). Label salience, in the sense of Schelling (1960), helps one to reconcile experimental data when one chooses symmetric games. However, in asymmetric games, payoff asymmetries may conflict with label salience. Indeed in such cases, label salience alone is unable to predict experimental choices. By contrast, a level- k model makes predictions that are in better accord with the evidence relative to classical game theory.

Section 13.4.2 considers the role of communication in enhancing coordination in a Battle of Sexes (BOS) game (Crawford, 2007). Section 13.4.3 considers a stag-hunt game in which payoff dominance and risk dominance are traditionally used to choose among multiple equilibria (Crawford, 2007). Several messages emerge, although they are sensitive to the specification of type $L0$ players. The level of coordination in level- k models is often higher than that under a Nash equilibrium. Communication can enhance the level of coordination in level- k models, even when it is not predicted to do so in classical game theory. Risk dominance and payoff dominance play no role in level- k models unlike in classical game theory. Coordination, when it occurs, is entirely accidental.

Section 13.4.4 considers the *market entry game* in which there is a market of fixed integer size, say $d > 0$, and there are $N > d$ potential entrants who can each supply one unit (Selten and Güth, 1982). If $n \leq d$ entrants enter, then they all have a payoff of 1. Otherwise, if $n > d$ entrants enter, then all entrants get zero. The Nash equilibrium is that $n = d$ entrants should enter. Most experiments show that, with some learning, the predictions of the Nash equilibrium hold, a finding that Kahneman (1988) found to be “magical”. This was often held up as an example of the ability of classical game theory to make successful predictions (Sundali et al., 1995). However, it turns out that CH models can explain this magical finding equally well (Camerer et al., 2004a; Crawford et al., 2013). The emerging field evidence supports an explanation based on level- k /CH models (Goldfarb and Yang, 2009; Goldfarb and Xiao, 2011).

Empirical evidence for level- k models is considered in Section 13.4.5 (Costa-Gomes et al., 2001; Costa-Gomes and Crawford, 2006). These papers examine the *choice data* of the subjects as well as the *search and lookup data* collected using the mouselab software.³ The authors allow for the presence of a large number of possible types of players, including non-level- k players. However, once the choice and the lookup data are combined, the most common types are the level- k types, $L1, L2$, and $L3$. Type $L0$ players are rare.⁴ Some evidence is also found for subjects who play the Nash equilibrium and those who use simple heuristics. One such simple heuristic is to carry out one round of iterated dominance, and then play a best response on the assumption that opponents

³ See Section 12.5.4 in Chapter 12 for a description of the relevant method.

⁴ This finding is replicated in several of the other papers considered in this chapter; for instance, see Crawford and Iriberri (2007b) in an auction context.

randomize uniformly over their strategies. There is a fair bit of stability of types across similar games (but see Hargreaves et al. (2014) and the discussion below).

Georganas et al. (2012) find that for one class of games there is stability of types, but for the other class in their experiments the level- k types are not stable. They conjecture that (i) the level- k heuristic may be activated in some games and not in others, but it is not clear what causes this activation, and (ii) the likely source of the difference in their results from earlier ones is that in Costa-Gomes et al. (2001) and Costa-Gomes and Crawford (2006) the instructions are too detailed, and subjects may have been primed to play best responses due to the number and nature of the practice rounds.

Clearly, the issue of priming of subjects prior to an experiment by playing best response games with them is an important one for experimental economists. Perhaps it may be considered to be a proxy for achieving greater external validity in the experiments. Nevertheless, this issue has also been picked up by several other emerging papers who avoid such priming of subjects. We consider some of these papers in Section 13.4.6 under the title “Do non-strategic players really exist?”⁵ The message arising from these papers is that type $L0$ players, and those who just make non-strategic choices (such as choosing dominant actions in p-beauty contest games), need not just be an artifact in the minds of level $L1$ players. Rather, such non-strategic players seem to comprise a significant fraction of the players—around 40% in some studies. Furthermore, there is significant heterogeneity in the choices of these players. Clearly, this area is very new and requires further research before a fuller picture emerges.

In classical game theory, and in many behavioral models, players derive utility from the strategies or actions of all players. However, in many contexts, players may also derive utility *directly* from the beliefs about the choices and actions of others. For instance, people may have beliefs about the beliefs of others about their actions, and they may dislike falling short of these beliefs (*guilt*). This, in turn, may influence their choice of action. The standard framework cannot be adapted by simply adding the psychological payoffs to the material payoffs, because beliefs themselves are endogenous. These considerations require the development of new techniques that is embodied in *psychological game theory*. *Psychological games* and the relevant equilibrium concepts are considered in Section 13.5.

The seminal contribution in psychological games is Geanakoplos et al. (1989). In order to build intuition for more complicated classes of games, we outline this contribution in Section 13.5.1 for one-stage games. Recall that in Part 2 of the book, we considered the empirical evidence for the role of reciprocity and intentions in influencing outcomes. A formal treatment of these issues requires psychological game theory. Drawing on the work of Geanakoplos et al. (1989), the seminal application to reciprocity, based on a normal form two-player game, is Rabin (1993). While Rabin’s model is the first to formalize issues of intentions and reciprocity, it is not designed for sequential games that are of considerable importance in modeling issues of intentions and reciprocity.⁶

One drawback of the framework proposed by Geanakoplos et al. (1989) is that the beliefs of players are *initial beliefs* even when the underlying game is a sequential game. In sequential games, the actions of players at any node may reveal their intentions, hence, the beliefs of other players about choices and beliefs may change as the game unfolds through subsequent nodes. Further-

⁵ Some of the literature considered in this section includes: Costa-Gomes and Weizsäcker (2008), Agranov et al. (2012), de Sousa et al. (2013), and Burchardi and Penczynski (2014).

⁶ For instance, Rabin’s model can explain rejections in ultimatum games; however, the model is consistent with many other equilibria as well, including overfair offers (e.g., proposer offers more than 50%), which are generally not observed.

more, in this framework, players derive utility solely from their own beliefs and payoffs depend only on the set of strategies that lead to a terminal node, and not the plans that players may have had.

In Section 13.5.2 we consider the more general framework of Battigalli and Dufwenberg (2009) that takes account of these objections. In Section 13.5.3, we give an application of this general framework to a consideration of reciprocity in sequential games, and introduce the notion of a *sequential reciprocity equilibrium* (Dufwenberg and Kirchsteiger, 2004). This allows for a better fit of the predictions with the relevant evidence. For instance, while Rabin's model predicts unconditional cooperation in the sequential prisoner's dilemma game, a sequential reciprocity equilibrium makes the plausible prediction of conditional cooperation in the same situation. We also consider some of the evidence for this framework from public goods games (Dufwenberg et al., 2011).

Dufwenberg and Kirchsteiger (2004) were mainly concerned with the modeling of reciprocity. In Section 13.5.4, we combine the role of reciprocity with considerations of social preferences in sequential games (Falk and Fischbacher, 2006). This more general framework has even better explanatory power. For instance, it can explain the experimental results arising from the ultimatum game, the dictator game, the gift exchange game, the public goods game, and the prisoner's dilemma game.

Section 13.5.5 considers the role of *guilt* in games (Battigalli and Dufwenberg, 2007). We consider a range of examples and evidence that illustrate the important role of guilt: from public goods games (Dufwenberg et al., 2011), to *promises* and *partnerships* (Charness and Dufwenberg, 2006). A growing number of empirical papers support a positive correlation between the degree of cooperative actions of players and their *second order beliefs*.⁷ This is consistent with players exhibiting *guilt aversion*. We also consider the recent work of Khalmetski et al. (2015) who reconcile the negative findings of Ellingsen et al. (2010) for guilt aversion with earlier work that reports positive findings.

Section 13.6 considers some aspects of the important work of Gintis (2009).⁸ He argues for retaining the rationality assumption in economics, i.e., humans have a complete and transitive preference relation between the relevant choices. His starting point is that there are no plausible *epistemic foundations* for the assumption of *common knowledge* of rationality in classical game theory. Hence, the very plausibility of a Nash equilibrium based on the epistemic conditions (Aumann and Brandenburger, 1995) is in doubt.⁹ It is possible that *social norms* and *conventions* may serve to bring about common priors and ensure common knowledge about an event. However, Gintis argues that it is a mistake to think of Nash equilibria as social norms (a common argument given by many game theorists) for two reasons. First, the epistemic conditions required for a Nash equilibrium are too stringent to be plausible. Second, many social norms are, in fact, not Nash equilibria.

Imagine a *choreographer* who coordinates human behavior. First, Gintis argues that social norms and conventions play the role of a choreographer. The foundation for this argument goes

⁷ First order beliefs are beliefs about the likely actions of others. Second order beliefs are beliefs about first order beliefs of others.

⁸ Our focus is mainly on Gintis's 2009 book but we encourage the reader to follow up the references in the book that form the basis of his arguments, for a more well-rounded understanding of the issues, as well as a better appreciation of the contribution by Herbert Gintis.

⁹ Aumann and Brandenburger (1995) showed that in two-player games, mutual knowledge of payoffs and of rationality is sufficient for a Nash equilibrium. However, for three or more players, this mutual knowledge assumption must be supplemented by common knowledge of conjectures.

back to Aumann (1987) who showed that under common priors Bayesian rationality guarantees that in any epistemic game economic agents play a correlated equilibrium. A second insight is to assign the role of creating the common priors for the choreographer. A necessary condition is that humans should be predisposed to following the recommendations of the choreographer to play a correlated equilibrium if they can do no better than by following the recommendations. Evolutionary arguments are then used to argue that such predispositions must exist.

Section 13.7 considers a range of other behavioral game theory models. These models often rely on plausible arguments that are often, but not always, underpinned by some form of psychological evidence. However, the psychological underpinnings of many of these models need to be better established. We consider existing evidence for these models but further empirical testing needs to be carried out to establish their relevance.

Section 13.7.1 considers *analogy-based expectations equilibrium* (ABEE) (Jehiel, 2005). In an analogy-based model, players have an information partition of the set of nodes of a game at which other players move. Each of these information partitions is known as an *analogy class*. On account of cognitive limitations, players assign the average action of the opponent to all the nodes in an analogy class and best respond to these average beliefs. Players need not know the payoffs of other players, nor the analogy classes of others. This lightens the cognitive burden on players from playing the game. However, at the moment we do not have a good account of how players construct analogy classes.

In Section 13.7.2 we consider the *subjective heterogeneous quantal response equilibrium* (SQRE) (Rogers et al., 2009). The QRE and the CH models relax different features of a Nash equilibrium. The QRE relaxes the best response feature and level- k /CH models relax the feature of mutually consistent beliefs. An SQRE, on the other hand, has the ambitious agenda of combining insights from both models and to show how these models are special cases of a more general overarching model.

Section 13.7.3 introduces the concept of a χ -*cursed equilibrium* (Eyster and Rabin, 2005), which applies to finite Bayesian games. In this model, χ -*cursed* players believe that with probability $\chi \in [0, 1]$, all the types of an opponent play the average optimal action, and with probability $1 - \chi$, the opponent's types play their type-contingent optimal actions. Thus, with probability χ , players do not show full awareness of the correlation between the types and the optimal actions of players. A player is fully cursed if $\chi = 1$, and the benchmark case of a Bayesian Nash equilibrium obtains when $\chi = 0$.

For finite Bayesian games, a χ -cursed equilibrium always exists. The main application of a χ -cursed equilibrium that we consider is to auctions in Section 13.8.3. For games of full information, the χ -cursed equilibrium reduces to a standard Nash equilibrium, which is often rejected by the data for new games and for early rounds of play in experimental games. The psychological underpinnings for the cursed equilibrium model are less clear than some of the other behavioral models, but the similarity in some settings with the ABEE, a model with arguably better psychological foundations, is encouraging.

Section 13.7.4 considers *evidential reasoning* and *evidential games* (al-Nowaihi and Dhami, 2015). This model applies to static games of full and incomplete information, although, in principle, it can be extended to dynamic games. Suppose that players choose their strategies simultaneously in full information or incomplete information settings. Players are typically uncertain about the actions that other players will choose in static games, perhaps because it is a new game/situation, or it is played infrequently, or the cognitive limitations of players may prevent the deductive reasoning needed for a Nash equilibrium. A great deal of evidence from psychology indicates that in such cases, players assign *diagnostic significance* to their own actions in making

inferences about the actions of other like-minded players. For instance, players who intend to cooperate in a prisoner's dilemma game, assign very high probability that others will cooperate. Voters who intend to vote for Democrats (resp. Republicans) in US elections assign very high probability that other voters will vote for Democrats (resp. Republicans) despite sharing identical public information on polls. A corrupt official may, likewise, overestimate the probability that others are corrupt, again taking his own actions as diagnostic of the actions of others.

Each player then best responds to his *evidential beliefs* about others, by solving essentially a decision theoretic representation of the underlying strategic situation. This affords enormous cognitive simplification of a complex underlying situation. Evidential games provide a general framework within which to think of many other belief formation processes such as level- k models, when restricted to static games. This line of work may be thought of as providing a *heuristics* and *biases* approach to strategic interaction, much like a similar approach in decision theory due to Daniel Kahneman and Amos Tversky that helped explain the evidence much better (see Part 7). There should be no presumption that evidential reasoning is rational, in the same way that judgment heuristics are not judged on their satisfaction of the relevant axioms of rationality. However, it provides a fast and frugal method of dealing with an uncertain situation, which is a typical feature of good heuristics. The relevant issue is not the a-priori rationality of such heuristics, but the conformity of empirical evidence with such modes of reasoning.

Section 13.7.5 considers the model of *noisy introspection* (Goeree and Holt, 2004). In a QRE, recall that players play a noisy best response but their beliefs are mutually consistent and each player takes account of the noise with which the other players choose their actions. Models of noisy introspection relax the mutual consistency of beliefs by relying on iterated conjectures in a noisy version of the standard notion of rationalizability. Finally, Section 13.7.6 gives a brief note on other models that include *team reasoning* (Bacharach, 1999, 2006) and *strong Stackelberg reasoning* (Colman et al., 2014).

Section 13.8 considers auction settings in which bidders bid for a single unit of an object. The evidence indicates that the standard predictions of auction theory are often violated by experiments. For instance, the actual bids of bidders are often different from their predicted bids in first and second price auctions, and the revenue equivalence theorem often does not hold. Particular interest has focused on the explanation of a *winner's curse* in common value auctions. Consider an object that has a common value for each bidder but each bidder receives a private signal of the value. The bidder with the most favorable signal wins the object but finds, ex-post, that the bid was too high relative to the common value; this is known as the winner's curse. We look at explanations of the winner's curse using two kinds of theoretical models—level- k models and the χ -cursed equilibrium. We also consider the empirical evidence for these alternative explanations.

13.2 Quantal response equilibrium (QRE)

In a *quantal response equilibrium* (QRE), players do not play best replies, but play noisy best replies, given their beliefs about the strategies of the opponents. All players have accurate beliefs about the noisy play of their opponents. QRE has the desirable feature that more rewarding strategies are played with relatively greater probabilities; in particular the best reply is played with the greatest probability. There are several possible reasons why players might play in this manner. Their payoffs might be subject to privately observed random shocks. Hence, other players may observe the play of a probabilistic best response. Yet another reason has its roots in the refinement literature in which players play their strategies with some noise (Van Damme, 1987).

One implication of the probabilistic choice of strategies in a QRE is that kinks, if any, in the best reply functions of players are smoothed. A second implication is that since all strategies are played with strictly positive probability, the vexed problem in classical game theory of assigning beliefs to “out of equilibrium paths of play” is avoided. For instance, in signaling games, issues arising from the failure of Bayes’ rule to assign probabilities to nodes that are not on the equilibrium path of play, do not arise.

13.2.1 QRE in normal form games

We first consider QRE in normal form games, as in McKelvey and Palfrey (1995), followed by the application of QRE to extensive form games.

Suppose that there are n players. The pure strategy set of player i is given by $S_i = \{s_{i1}, s_{i2}, \dots, s_{in_i}\}$, where n_i is the cardinality of the set. To minimize notation, we denote a particular member of the set S_i by s_i . Thus, any pure strategy profile for the n players is denoted by $s = \{s_1, s_2, \dots, s_n\} \in S = \times_i S_i$. The pure strategy set of all players other than player i is denoted by S_{-i} , with the typical element denoted by s_{-i} .

Let Σ_i be the set of all probability distributions over the set S_i , with typical element σ_i . We will often denote the probability with which player i plays pure strategy j by σ_{ij} . A mixed strategy profile of the n players is denoted by $\sigma = \{\sigma_1, \sigma_2, \dots, \sigma_n\}$, where σ is a member of the set $\Sigma = \times_i \Sigma_i$. We will sometimes denote the mixed strategy profile of the game by $\sigma = (\sigma_i, \sigma_{-i})$, where $\sigma_i \in \Sigma_i$ and $\sigma_{-i} \in \times_{j \neq i} \Sigma_j$.

The von Neumann–Morgenstern utility of player i from the pure strategy profile $s = (s_i, s_{-i})$ is given by $u_i(s) = u_i(s_i, s_{-i}) : S \rightarrow \mathbb{R}$. The normal form of the game is denoted by $G = \{n, S_1, \dots, S_n, u_1, \dots, u_n\}$. Consider the mixed strategy profile σ , where $\sigma(s)$ gives the probability with which the pure strategy profile s is played. Then the expected payoff of player i is given by

$$u_i(\sigma) = \sum_{s \in S} \sigma(s) u_i(s). \quad (13.1)$$

Consider player i and begin with some mixed strategy profile of the game, $\sigma = (\sigma_i, \sigma_{-i})$. Denote the mean utility of player i when he plays the pure strategy s_{ij} and others play the mixed strategy vector σ_{-i} by $\bar{u}_{ij}(\sigma) = u_i(s_{ij}, \sigma_{-i})$; this is found by applying (13.1). Let $\bar{u}_i(\sigma) = \{\bar{u}_{i1}, \bar{u}_{i2}, \dots, \bar{u}_{in_i}\}$, so that $\bar{u}(\sigma) = \{\bar{u}_1, \bar{u}_2, \dots, \bar{u}_n\}$. QRE introduces additive payoff shocks (as in *stochastic utility models*), so that the *true utility* of player i from $\sigma = (s_{ij}, \sigma_{-i})$ is given by

$$v_{ij}(\sigma) = \bar{u}_{ij}(\sigma) + \epsilon_{ij}, \quad j = 1, 2, \dots, n_i. \quad (13.2)$$

$E\epsilon_{ij} = 0$, where the expectation is taken over all pure strategies of player i ; this holds for all i . The vector of shocks $\epsilon_i = (\epsilon_{i1}, \epsilon_{i2}, \dots, \epsilon_{in_i}) \in \mathbb{R}^{n_i}$ has a joint probability density $f(\epsilon_i)$ such that the marginal distribution $f(\epsilon_{ij})$ exists for all j . Such an f is said to be *admissible*. These shocks may also arise from mistakes in implementing choices. The vector of payoffs of player i arising from his various pure strategies, when others play σ_{-i} is given by

$$v_i(\sigma) = (v_{i1}(\sigma), v_{i2}(\sigma), \dots, v_{in_i}(\sigma)) \in \mathbb{R}^{n_i}.$$

The behavioral assumption made in QRE is that player i plays action s_{ij} only if it gives a better payoff relative to the other available actions, i.e., if

$$v_{ij}(s_{ij}, \sigma_{-i}) \geq v_{ik}(s_{ik}, \sigma_{-i}) \text{ for all } s_{ik} \in S_i. \quad (13.3)$$

Given σ , denote the set of vectors, $\varepsilon_i \in \mathbb{R}^{n_i}$, for which (13.3) holds, by Φ_{ij} .

$$\Phi_{ij} = \{\varepsilon_i \in \mathbb{R}^{n_i} : v_{ij}(s_{ij}, \sigma_{-i}) \geq v_{ik}(s_{ik}, \sigma_{-i}) \text{ for all } s_{ik} \in S_i\}. \quad (13.4)$$

Substituting (13.2) in (13.4) we get,

$$\Phi_{ij}(\bar{u}_i(\sigma)) = \{\varepsilon_i \in \mathbb{R}^{n_i} : \bar{u}_{ij}(\sigma) + \varepsilon_{ij} \geq \bar{u}_{ik}(\sigma) + \varepsilon_{ik} \text{ for all } s_{ik} \in S_i\}. \quad (13.5)$$

The probability that $\varepsilon_i \in \Phi_{ij}(\bar{u}_i(\sigma))$ is given by $\pi_{ij}(\bar{u}_i(\sigma))$ such that

$$\pi_{ij}(\bar{u}_i(\sigma)) = \int_{\Phi_{ij}(\bar{u}_i(\sigma))} f(\varepsilon_i) d\varepsilon_i. \quad (13.6)$$

Thus, for a given σ , the pure strategies of player i , $(s_{i1}, s_{i2}, \dots, s_{in_i})$, are played with respective probabilities $\pi_i(\sigma) = (\pi_{i1}(\sigma), \pi_{i2}(\sigma), \dots, \pi_{in_i}(\sigma))$. In turn, these probabilities are induced by the probability distribution of the error term ε_i . Analogous expressions can be found for the play of pure strategies of the other players. Given σ , denote the vector of such probabilities by $\pi(\sigma) = (\pi_1(\sigma), \pi_2(\sigma), \dots, \pi_n(\sigma))$; this vector gives the probability with which each player plays each of the pure strategies.

Definition 13.7 (*Quantal response equilibrium*): A QRE, σ^* , is a fixed point of the mapping $\pi(\sigma^*) = \sigma^*$. Alternatively, for all players, $i = 1, 2, \dots, n$ and for all strategies $j = 1, 2, \dots, n_i$, we must have $\sigma_{ij} = \pi_{ij}(\bar{u}_i(\sigma))$.

In words, this is what happens. Each player decides to play all his pure strategies with *choice probabilities* conditional on his *beliefs* about the probabilistic play of others. A QRE arises when all the choice probabilities and the beliefs coincide for all players. Hence, the idea is very similar to a rational expectations equilibrium. To reiterate, each individual plays his noisy best response in a QRE, taking full account of the errors made by the opponents that are reflected in their probabilistic play of pure strategies.

The function $\pi_i(\bar{u}_i(\sigma)) = (\pi_{i1}, \pi_{i2}, \dots, \pi_{in_i})$ is known as a *quantal response function*. The function $\pi = (\pi_1, \pi_2, \dots, \pi_n)$ is defined over a compact interval and each π_i is continuous. Thus, it follows from Brouwer's fixed point theorem that the fixed point in Definition 13.7 exists. From (13.5), the set $\Phi_{ij}(\bar{u}_i(\sigma))$ expands for larger values of $\bar{u}_{ij}$, hence, from (13.6), π_{ij} is a non-decreasing function of $\bar{u}_{ij}$. Thus, if the mean payoff from a pure strategy increases, the probability with which it is played cannot decrease.

If every ε_{ij} is an independent draw from an extreme value distribution,¹⁰ then the choice probabilities take the *logit form*,¹¹ i.e.,

$$\pi_{ij}(\bar{u}_i(\sigma)) = \frac{\exp(\lambda \bar{u}_{ij}(\sigma))}{\sum_{k=1}^{n_i} \exp(\lambda \bar{u}_{ik}(\sigma))}. \quad (13.7)$$

¹⁰ The extreme value distribution has a cumulative distribution function given by $F(\varepsilon) = \exp[-\exp(-\lambda\varepsilon)]$. The mean of the distribution is given by $E\varepsilon = \gamma/\lambda$ where $\gamma \approx 0.577$ is Euler's constant.

¹¹ For a derivation, see Appendix B in Part 5.

	L	R
T	-1, 1	x, -1
B	-1, -1	-1, 1

Figure 13.1 Two-player matching pennies game.

The logit density has a single parameter λ . From observed values of σ , the parameter λ can be estimated by maximum likelihood methods. λ is inversely proportional to the variance of the errors. As $\lambda \rightarrow \infty$ the variance of the shocks goes to zero and behavior approaches a Nash equilibrium. On the other hand, when $\lambda = 0$, people choose randomly among their pure strategies. Thus, the fitted values of λ can inform us about the degree of noise in the game.

QRE has often been successful in explaining behavior in experimental games where the Nash equilibrium predictions have been unsuccessful. We shall review some of this evidence below from different kinds of games such as auction games, bargaining games, and coordination games. A strong case for the QRE as an alternative to Nash equilibrium is made by Camerer et al. (2004b). “Quantal response equilibrium (QRE), a statistical generalization of Nash, almost always explains the deviations from Nash and should replace Nash as the static benchmark to which other models are routinely compared.”

Consider the two-player matching pennies game shown in Figure 13.1. There is no pure strategy Nash equilibrium in this game. In most versions of the matching pennies game, $x = 1$, but here we allow for the effects of different levels of stakes. Let us first compute the mixed strategy Nash equilibrium (MSE) of the game. Let the column player play R with probability π_R and the row player play T with probability π_T . The row player prefers to play T over B if and only if

$$x\pi_R - (1 - \pi_R) \geq -\pi_R + (1 - \pi_R) \Leftrightarrow \pi_R \geq \frac{2}{3+x}, \quad (13.8)$$

which, for $x = 1$, requires that $\pi_R \geq 1/2$. The column player prefers to play L over R if and only if

$$\pi_T - (1 - \pi_T) \geq -\pi_T + (1 - \pi_T) \Leftrightarrow \pi_T \geq \frac{1}{2}. \quad (13.9)$$

Each of the players is indifferent among the respective pure strategies, if

$$\pi_T^* = \frac{1}{2}; \pi_R^* = \frac{2}{3+x}, \quad (13.10)$$

which gives the MSE.

Let us now compute the logit QRE using (13.7). For expositional purposes, consider the row and column players, respectively, as players 1 and 2. Given that the column player plays R with probability π_R , the respective expected payoffs to the row player from playing the two pure strategies, T, B , are

$$\bar{u}_{1T}(\pi_R) = x\pi_R - (1 - \pi_R) \text{ and } \bar{u}_{1B}(\pi_R) = -\pi_R + (1 - \pi_R). \quad (13.11)$$

Similarly, conditional on the row player playing T with probability π_T , the respective expected payoff of the column player from the strategies L, R , are




Figure 13.2 QRE and Nash equilibrium for the cases $x = 1$ and $x = 3$.

$$\bar{u}_{2L}(\pi_T) = \pi_T - (1 - \pi_T) \text{ and } \bar{u}_{2R}(\pi_T) = -\pi_T + (1 - \pi_T). \quad (13.12)$$

Given (13.11), (13.12), and using (13.7), the respective probabilities with which the row and the column players play strategies, T, R , in a logit QRE are given by:

$$\pi_T = \frac{\exp(\lambda \bar{u}_{1T}(\pi_R))}{\exp(\lambda \bar{u}_{1T}(\pi_R)) + \exp(\lambda \bar{u}_{1B}(\pi_R))}, \quad (13.13)$$

$$\pi_R = \frac{\exp(\lambda \bar{u}_{2R}(\pi_T))}{\exp(\lambda \bar{u}_{2R}(\pi_T)) + \exp(\lambda \bar{u}_{2L}(\pi_T))}. \quad (13.14)$$

The solution to (13.13), (13.14) gives the QRE: π_T^{qre}, π_R^{qre} .

In Figure 13.2, the best reply functions are shown as the continuous step functions; their intersection gives the MSE in (13.10). Two cases are shown, $x = 1$ and $x = 3$. The MSE for the case $x = 1$ is shown at point A, where $\pi_T^* = \pi_R^* = \frac{1}{2}$. When x increases from $x = 1$ to $x = 3$, the best reply function for the column player does not shift, but that for the row player shifts to the left (shown as the dotted step function). In the new Nash equilibrium, shown at point B, $\pi_T^* = \frac{1}{2}, \pi_R^* = \frac{1}{3}$. An increase in x does not alter the probability, π_T^* , with which the row player plays T . This is a testable restriction of an MSE.

The logit probabilities under QRE, derived from (13.13), (13.14), are shown as the smooth curves. For any non-degenerate choice of the opponent's mixed strategy, each player plays both strategies with non-zero probabilities. The row (column) player's quantal best reply starts in the southwest (northwest) corner of the unit square and ends in the northeast (southeast) corner. The initial QRE, corresponding to $x = 1$ is shown at point A, as $\pi_T^{qre} = \pi_R^{qre} = 0.5$ (check that this holds for any λ).

When x increases to $x = 3$, the quantal reply function for the row player (shown as the dotted curve) shifts left, while that of the column player is unaffected; since the payoff from T increases,

	L	R		L	R		L	R
U	9, 0	0, 1	U	9, 0	0, 4	U	36, 0	0, 4
D	0, 1	1, 0	D	0, 4	1, 0	D	0, 4	4, 0
Game A			Game B			Game C		

Figure 13.3 Three normal form games.

Source: Reprinted from *Journal of Economic Behavior and Organization* 42: 523–48. McKelvey, R., Palfrey, T., and Weber, R., “The effects of payoff magnitude and heterogeneity on behavior in 2×2 games with mixed strategy equilibria.” Copyright © 2000, with permission from Elsevier.

it is played with a higher probability. In the new QRE, shown at point C, $\pi_R^{qre} < 0.5$, while $\pi_T^{qre} > 0.5$; this is a “better response” by both players. Hence, the prediction of the QRE is different from the Nash prediction. As x increases the QRE changes, but there is no change in a Nash equilibrium. This, *own-payoff effect* on observed probabilities of the play of strategies is observed (Ochs, 1995; McKelvey et al., 2000; Goeree et al., 2003); which is consistent with QRE but not Nash.

It is likely that the degree of errors people make, as captured in the parameter λ , varies across people. McKelvey et al. (2000) confirm this to be the case. They find that a heterogeneous QRE model, HQRE (see Section 13.7.2), accounts much better for the data and outperforms a Nash equilibrium. They consider the three games, A, B, and C shown in Figure 13.3.

The generic game is Game A. Game B is derived from A by multiplying the payoffs of the column player by 4. Game C is derived from A by multiplying the payoffs of all players by 4. Since the Nash equilibrium is invariant to positive affine transformations of the payoffs, the Nash equilibrium in all three games is identical.¹² Suppose that the row player plays U with probability p and the column player plays L with probability q . In the unique Nash equilibrium of the game, $(p^*, q^*) = (0.5, 0.10)$.

In contrast to the Nash equilibrium, the QRE is not invariant to affine transformations of payoffs. Figure 13.4 shows values of (p, q) in a logit QRE for games A, B, C when we vary λ (subscripts on p, q denote the game played). When $\lambda = 0$, the mixing probabilities in a QRE, p, q , correspond to random choice. One may conjecture that as players gain experience, the variance of errors with which they play their strategies decreases (or, λ increases). As λ increases, p first increases then decreases before converging towards the Nash equilibrium value of 0.5. Thus, unlike a Nash equilibrium, the prediction of QRE is that the row player will first overplay U, before converging to the Nash prediction. The behavior of q in Figure 13.4, in contrast to the behavior of p , exhibits no hump-shaped profiles. Rather, starting from random choice at $\lambda = 0$, q gradually moves downwards towards the Nash prediction of 0.10.

The three games differ with respect to the speed with which the behavior converges to the Nash equilibrium. Games with the highest payoffs are more conducive to quicker convergence. Thus, the speed of convergence is highest in Game C, followed, respectively, by Games B and A (McKelvey et al., 2000, p. 526).¹³ In particular, for any λ , it is true that $0.5 = p^* < p_B < p_A$ and $0.5 = p^* < p_B < p_C$.

¹² If (σ_i, σ_{-i}) is a Nash equilibrium then for all i we have $u(\sigma_i, \sigma_{-i}) \geq u(\sigma'_i, \sigma_{-i})$ for all possible mixed strategies σ'_i of player i . Define $v = a + bu$, where $b > 0$. Then we also have $v(\sigma_i, \sigma_{-i}) \geq v(\sigma'_i, \sigma_{-i})$. For instance, Game B is derived from Game A by letting $a = 0$ and $b = 4$ for the column player.

¹³ The difference between the speed of convergence between Games B and C is not entirely transparent from looking at Figure 13.4.




Figure 13.4 Predictions of the quantal response equilibrium for games A, B, and C for various values of λ .

Source: Reprinted from *Journal of Economic Behavior and Organization* 42: 523–48. McKelvey, R., Palfrey, T., and Weber, R., "The effects of payoff magnitude and heterogeneity on behavior in 2×2 games with mixed strategy equilibria." Copyright © 2000, with permission from Elsevier.

Since a Nash equilibrium corresponds to no errors ($\lambda \rightarrow \infty$), one may view the movement from left to right as a process of learning. Indeed, one purpose of these experiments is to check if the errors with which people play their strategies diminish with experience.

In the first instance, three candidate models are pitted against each other in explaining the data. (1) Random play (i.e., $p = q = 0.5$), (2) QRE, and (3) Noisy Nash model (NNM). In the NNM, players play a linear combination of random play and Nash play. Thus, the row and the column players play the respective strategies, $\hat{p} = (1 - \gamma)0.5 + \gamma p^*$ and $\hat{q} = (1 - \gamma)0.5 + \gamma q^*$, $\gamma \in [0, 1]$. In all cases, $\hat{p} = 0.5$ because $p^* = 0.5$, but as γ increases, $\hat{q}$ falls from a maximum value of 0.5 to $q^* = 0.10$. For a fixed value of γ , the NNM makes identical predictions for all three games, A, B, and C.

For the aggregate data in the experiments, we get $(p, q) = (0.643, 0.241)$ for Game A; $(p, q) = (0.630, 0.244)$ for Game B; and $(p, q) = (0.594, 0.257)$ for Game C. Thus, the Nash equilibrium is not supported by the data from any of the games. Random play is rejected for all games. In all games, the row player overplays U and the column player overplays L (relative to a Nash), which is consistent with the predictions of a QRE. Affine transformations in payoffs not only create departures from the Nash predictions but create inter-game differences in strategies. For instance, the play of D by the row player is statistically different between games C and A, and between games C and B. There are no trends over time in the estimated values of λ , ruling out any role for learning.¹⁴ Based on a likelihood ratio test, NNM beats random play and Nash, and for games A and C it also beats QRE, although for game B, QRE does best.

The choices of the row and the column players at the individual level exhibit significant heterogeneity, suggesting that the value of λ across the players is variable. Thus, the analysis carried out on the basis of a fixed λ across the players is based on a mis-specified model. In the second

¹⁴ Eight sessions were run. In each session, subjects played two games against anonymous opponents. Each game was played over 50 rounds.

part of the paper, the authors allow for heterogeneity across players in λ (heterogenous quantal response equilibrium, HQRE).¹⁵ The HQRE fits the data for all games much better than QRE and performs better than NNM for all games that are considered.

13.2.2 QRE in extensive form games

McKelvey and Palfrey (1998) extended the concept of QRE to extensive form games. We shall only give a heuristic treatment of this topic, and one example. Consider an extensive form game in which the set of information sets is given by H . Pick any arbitrary information set, $h \in H$ and let A_h be the set of actions available to a player at that information set. As in normal form games, we assume that each player plays with some noise, so each possible action at each information set is played probabilistically. For any action $a_{hj} \in A_h$, we wish to calculate the expected payoff of the player, $\bar{u}_{hj}$ (this calculation will be explicit in the example, below). Since players play their actions with noise, the “true” utility from action a_{hj} is given by

$$v_{hj} = \bar{u}_{hj} + \varepsilon_{hj},$$

where ε_{hj} is the noise associated with action a_{hj} . Our behavioral assumption is that the player who moves at the information set A_h , chooses action a_{hj} , if and only if

$$v_{hj} \geq v_{hk} \text{ for all } a_{hk} \in A_h. \quad (13.15)$$

In the logit form, the probability of playing strategy $a_{hj} \in A_h$ is given by

$$\pi_{hj} = \frac{\exp(\lambda \bar{u}_{hj})}{\sum_{a_{hk} \in A_h} \exp(\lambda \bar{u}_{hk})}. \quad (13.16)$$

The RHS of (13.16) depends on beliefs because it involves the calculation of expected payoffs from various actions. The LHS of (13.16), on the other hand, is the actual probability with which the player chooses action a_{hj} . In a QRE, beliefs and choice probabilities coincide.

Consider, as an example, the bilateral crises game in Signorino (1999). There are two countries, 1 and 2. Each country has two actions at each of its decision nodes, fight (F) or not fight (N). The sequence of moves is shown in the extensive form representation of the game in Figure 13.5. The final outcomes at each of the terminal nodes are represented by acronyms (payoffs are suppressed). For instance, if 1 and 2 both choose N, then the outcome is the status quo (SQ). If 1 does not fight at any node but 2 fights, then the outcome is that player 1 capitulates (C1). If player 2 initiates a fight and 1 sequentially decides to fight, the outcome is a player 2 initiated war (W2). C2 and W1 are defined in an analogous manner. Define the utility of player $i = 1, 2$ from outcome $j = \text{SQ, C1, C2, W1, W2}$, by u_{ij} .

Figure 13.5 also shows the probabilities with which each player plays his pure strategies. For instance, the probability that player 1 fights at the initial node is p_1 (thus, the probability that he does not fight is $1 - p_1$). If player 1 does not fight then player 2 fights with probability q_2 followed by player 1 who fights with probability p_4 . If player 1 fights at the initial node, then player 2 fights with probability q_3 .

¹⁵ This was achieved by assuming that the distribution of λ is lognormal across the players. Maximum likelihood was then used to estimate the mean and the variance of the normal distribution.




Figure 13.5 The bilateral crises game.

Source: Signorino, C. S. (1999) "Strategic interaction and the statistical analysis of international conflict." *American Political Science Review*, 93(2): 279–97, reproduced with permission from Cambridge University Press.

Using (13.16), the logit choice probabilities are recursively given by the following set of equations.

$$q_3 = \frac{\exp(\lambda u_{2W_1})}{\exp(\lambda u_{2W_1}) + \exp(\lambda u_{2C_2})} \quad (13.17)$$

$$p_4 = \frac{\exp(\lambda u_{1W_2})}{\exp(\lambda u_{1W_2}) + \exp(\lambda u_{1C_1})} \quad (13.18)$$

$$q_2 = \frac{\exp(\lambda [p_4 u_{2W_2} + (1 - p_4) u_{2C_1}])}{\exp(\lambda [p_4 u_{2W_2} + (1 - p_4) u_{2C_1}]) + \exp(\lambda u_{2SQ})} \quad (13.19)$$

$$p_1 = \frac{\exp(\lambda [q_3 u_{1W_1} + (1 - q_3) u_{1C_2}])}{\exp(\lambda [q_3 u_{1W_1} + (1 - q_3) u_{1C_2}]) + \exp(\lambda [q_2 [p_4 u_{1W_2} + (1 - p_4) u_{1C_1}] + (1 - q_2) u_{1SQ}])} \quad (13.20)$$

A QRE of the bilateral crises game is the set of values $(p_1^{qre}(\lambda), q_2^{qre}(\lambda), p_4^{qre}(\lambda), q_3^{qre}(\lambda))$ in which the choices and the beliefs of the players are mutually consistent. Once the equilibrium choice probabilities are computed, then we can directly compute the probability of each outcome by observing the path of play. For instance, the probability of the status quo equals $(1 - p_1^{qre}(\lambda)) (1 - q_2^{qre}(\lambda))$ and the probability of war initiated by player 1 equals $p_1^{qre}(\lambda) q_3^{qre}(\lambda)$. Thus, the total probability of war is

$$p_W(\lambda) = p_1^{qre}(\lambda) q_3^{qre}(\lambda) + (1 - p_1^{qre}(\lambda)) q_2^{qre}(\lambda) p_4^{qre}(\lambda). \quad (13.21)$$

The size of λ reflects errors made by the players. Players choose randomly if $\lambda = 0$, and as $\lambda \rightarrow \infty$ they are fully rational in the sense of making no errors and complying with the Nash predictions. Given the observed probabilities of play of the pure strategies, λ can be estimated using maximum likelihood techniques. For instance, suppose that we have N data points on whether wars occurred. Let d_i be a dummy variable for the i th observation that takes a value 1 if a war occurred, and 0 otherwise. The likelihood function is

$$L(\lambda) = \prod_{i=1}^N p_W^{d_i}(\lambda) [1 - p_W(\lambda)]^{1-d_i}.$$

The likelihood function can now be maximized with respect to λ to get the estimated value of λ . Since an increase in λ is associated with a decrease in errors in decision making, experiments repeated over time can be used to check if λ increases, which would indicate learning.

One application of these ideas is to the theory of *information cascades*. In such models, nature first chooses some state of the world that is not revealed to a sequence of T players. The players then choose some action in a sequential manner. Each player takes some action, based on a private signal about the state of the world that might have been initially chosen by nature. The signals of the players are uncorrelated with each other. At the time of choosing an action, each player observes the history of the game, i.e., the actions of the sequentially preceding players, but not their signals. An *information cascade* is said to take place when after some set of players have chosen their actions, each subsequent player rationally ignores their own private signal, so learning completely ceases (Banerjee, 1992; Bikhchandani et al., 1992).

When players cease to learn in this manner, there are potentially serious welfare consequences, e.g., bank runs, mass hysteria, and stock market bubbles. Information cascades are distinguished from *herd behavior* in that while the actions of the players might be observationally equivalent, under herd behavior, players do not necessarily ignore their private information.¹⁶ Using experimental data, Goeree et al. (2007) show that when cascades are predicted in 100% of the sequences, i.e., behavior is conformable to a Nash equilibrium, they only observe cascades in less than 7% of the sequences (out of 116 possible sequences). Kübler and Weizsäcker (2004) also report negative findings, similar to these. A QRE model is able to explain the experimental findings by placing restrictions on the inferences of subsequent players if a deviation from a cascade occurs. In a QRE, learning continues to take place even after an information cascade has taken place.

13.2.3 Does the QRE impose falsifiable restrictions?

Haile et al. (2008) show that unless one makes specific assumptions on the distribution of the errors, QRE imposes no falsifiable restrictions. If the errors ε_{ij} (associated with the pure strategy j of player i) are independently and identically distributed then there are no problems for QRE. However, in the absence of such information, the results from QRE need to be interpreted with caution. We demonstrate this result by using examples from Goeree et al. (2005) who responded to the working paper version of Haile et al. (2008). While the focus of Haile et al. (2008) is on strategic interaction, the pedagogical merit of the examples in Goeree et al. (2005) is that they focus on a single decision maker (player 1) playing a game against nature.

Suppose that a single decision maker has two strategies s_{11} and s_{12} ; the respective mean payoffs from these strategies are $\bar{u}_{11} = 1$ and $\bar{u}_{12} = 2$. Under the setup of the QRE, the actual payoffs are given by $v_{1j} = \bar{u}_{1j} + \varepsilon_{1j}$, $j = 1, 2$. Thus, the probability that strategy s_{11} is chosen is $\pi_{11} = \Pr[v_{11} \geq v_{12}] \Leftrightarrow \Pr[1 + \varepsilon_{11} \geq 2 + \varepsilon_{12}]$. Thus,

$$\pi_{11} = \Pr[\varepsilon_{11} - \varepsilon_{12} \geq 1] \text{ and } \pi_{12} = 1 - \pi_{11}. \quad (13.22)$$

The joint density of $\varepsilon_{11}, \varepsilon_{12}$ is given by $f(\varepsilon_{11}, \varepsilon_{12})$. Let us now drop the assumption that $\varepsilon_{11}, \varepsilon_{12}$ are independently and identically distributed. There are two cases to consider.

¹⁶ So essentially, information cascades imply herding but the converse is not true. See Smith and Sørensen (2000) for this distinction. Boğaçhan and Kariv (2004) implement this distinction in their experimental design.

1. ε_{11} and ε_{12} are independently but not identically distributed.
2. ε_{11} and ε_{12} are identically but not independently distributed.

It can be shown that in either of these two cases, for any $p \in (0, 1)$, we can find some $f(\varepsilon_{11}, \varepsilon_{12})$ such that $\pi_{11} = p$ and $\pi_{12} = 1 - p$. Since we started off with an arbitrary $p \in (0, 1)$, in these cases, QRE is not falsifiable and any observed frequency of play of the two strategies can be rationalized.

Consider the first case where ε_{11} and ε_{12} are independently but not identically distributed. Suppose that an experimenter observes that the frequency of play of strategy s_{11} is p . He then imposes the following distributional assumptions. (i) ε_{11} takes a value 2 with probability p and a value $\frac{-2p}{1-p}$ with probability $1-p$. (ii) ε_{12} takes a value 0 with probability 1. Under these assumptions it follows that

$$\pi_{11} = \Pr[\varepsilon_{11} - \varepsilon_{12} \geq 1] = p,$$

which gives a perfect fit of QRE with the evidence.

Now consider the second case of identical but not independent marginal distributions. Suppose that ε_{12} is distributed uniformly over $\left[0, \frac{2}{1-p}\right]$. The marginal distribution of ε_{11} is not independent. It takes the following form.

$$\varepsilon_{11} = \begin{cases} \varepsilon_{12} + 2 & \text{when } \varepsilon_{12} \in \left[0, \frac{2}{1-p} - 2\right], \\ \varepsilon_{12} + 2 - \frac{2}{1-p} & \text{otherwise.} \end{cases}$$

It follows that the condition $\varepsilon_{11} - \varepsilon_{12} \geq 1$ holds only over the interval $\left[0, \frac{2}{1-p} - 2\right]$, where the marginal distributions of ε_{11} and ε_{12} are both uniform. Thus,

$$\Pr[\varepsilon_{11} - \varepsilon_{12} \geq 1] = \Pr\left[\varepsilon_{12} \leq \frac{2}{1-p} - 2\right] = p.$$

Thus, if an experimenter observes the frequency of play of strategy s_{11} to be p and if he had the freedom to choose the distributions of $\varepsilon_{11}, \varepsilon_{12}$ in the manner described above, then QRE would give a perfect fit with the data. These observations hold for an arbitrary number of players and strategies (Haile et al., 2008).

In light of this critique, researchers who wish to use the QRE have at least two options. First, they could continue to assume that the ε_{ij} are independently and identically distributed. Second, they might wish to use a variant of the QRE, the *regular QRE*, which was proposed by Goeree et al. (2005). The regular QRE makes no assumptions about the distribution of errors but, rather, gives an axiomatic foundation for the QRE, based on a set of plausible axioms, as we now describe.

From (13.6), we know that π_{ij} is a function of $\bar{u}_i = (\bar{u}_{i1}(\sigma), \dots, \bar{u}_{ij}(\sigma), \dots, \bar{u}_{in_i}(\sigma))$.

Definition 13.8 (*Regular quantal response function*; Goeree et al. 2005) *The regular quantal response function satisfies the following axioms for all $\bar{u}_i(\sigma) \in \mathbb{R}^{n_i}$.*

1. *Interiority:* $\pi_{ij}(\bar{u}_i(\sigma)) > 0$ for all i, j .
2. *Continuity:* $\pi_{ij}(\bar{u}_i(\sigma))$ is continuous and differentiable in $\bar{u}_i(\sigma)$.

3. *Responsiveness*: $\frac{\partial \pi_{ik}(\bar{u}_i(\sigma))}{\partial \bar{u}_{ik}(\sigma)} > 0$, $k = 1, 2, \dots, n_i$.
4. *Monotonicity*: If $\bar{u}_{ij}(\sigma) > \bar{u}_{ik}(\sigma)$, then $\pi_{ij} > \pi_{ik}$ for all i, j, k .¹⁷

Interiority implies that all pure strategies are played with strictly positive probability. Continuity is a purely technical restriction, which ensures that slight changes in mean payoffs do not lead to jumps in the probability with which the associated strategies are played. Responsiveness requires that strategies whose expected payoff increases are played with a greater probability. Monotonicity requires that strategies whose mean payoffs are higher are also played with a higher probability. Responsiveness and monotonicity appear to be very similar but they are distinct; one of the exercises gives an example that demonstrates their difference. A regular QRE requires that the probabilities with which players play their pure strategies satisfy Definition 13.8 and the beliefs and choices of players coincide.

Definition 13.9 (Regular QRE): Suppose that π_{ij} is regular (Definition 13.8). A regular QRE arises for all players, $i = 1, 2, \dots, n$ and all strategies $j = 1, 2, \dots, n_i$ if $\sigma_{ij} = \pi_{ij}(\bar{u}_i(\sigma))$.

A regular quantal response function preserves continuity of $\pi_{ij}(\bar{u}_i(\sigma))$, hence, Brouwer's fixed point theorem can be used to ensure the existence of a regular QRE.

13.2.4 An application of QRE to the traveller's dilemma game

Consider Basu's (1994) *traveller's dilemma game*. In this game, an airline loses two identical pieces of luggage of two distinct travellers, who are returning from a holiday from an identical destination. It is public knowledge that each piece of luggage has an identical value,¹⁸ but the exact value is private knowledge to the travellers. The airline is concerned that the passengers might make an excessive claim. Hence, it makes the following announcement to the two travellers: "We know that the bags have identical contents, and we will entertain any claim between \$80 and \$200, but you will each be reimbursed an amount that equals the minimum of the two claims submitted. If the two claims differ, we will also pay a reward of \$R to the person making the smaller claim and we will deduct a penalty of \$R from the reimbursement to the person making the larger claim."

The unique Nash equilibrium of this game is for each individual to make an identical claim of \$80. To prove this, assume to the contrary that travellers 1 and 2 make respective claims x_1, x_2 such that $x_2 > 80$. Assume risk neutrality and $0 < \varepsilon < R$. If $x_1 = x_2 + \varepsilon > x_2$, then the payoff of traveller 1 is $x_2 - R$. If traveller 1 makes a claim $x_1 = x_2 - \varepsilon < x_2$, then his payoff is $x_2 - \varepsilon + R > x_2 - R$. Thus, it is a strictly dominant strategy to make a lower claim than the other player. This argument applies symmetrically to player 2. Iteratively deleting strictly dominated strategies in this manner, in the unique Nash equilibrium, $x_1 = x_2 = 80$, independent of the penalty R . This is the unique rationalizable equilibrium outcome.

However, one's intuition may suggest that as R becomes small, the travellers will attempt to make the maximum possible claims. Capra et al. (1999) conduct lab experiments to see if the outcome is affected by the penalty, R . In their experiments, the values of the items and the penalties are denominated in cents (so, e.g., \$80 is replaced by 80 cents).¹⁹

¹⁷ Goeree et al. (2005) show that *interchangeability* is a sufficient condition for monotonicity. Interchangeability, in turn, implies that the joint distribution of errors $f_i(\varepsilon_{i1}, \varepsilon_{i2}, \dots, \varepsilon_{in_i})$ for player i is invariant to permutations of the ε_{ij} terms. This implies that the marginal distributions are identical and so are covariances and conditional moments.

¹⁸ Perhaps each traveller has been able to produce a witness who was present when the luggage was packed and based on an estimate, the value of the two pieces of lost luggage is deemed identical.

¹⁹ Groups of 9–12 economics students were recruited to play a series of traveller's dilemma games. They were given a \$6 participation fee and could earn \$24–\$44 for a session.




Figure 13.6 Claims, over time, for various values of R .

Source: Capra et al. (1999) with permission from the American Economic Association.

The experimental findings are as follows.

1. The behavior of the subjects was heavily influenced by the penalty, R , as shown in Figure 13.6. For low values of R , behavior diverged significantly from the Nash equilibrium and no evidence of convergence was found. For instance, for low reward/penalty rates ($R = 5$ and $R = 10$), the average claim was 180, close to the maximum possible claim of 200. In the two intermediate cases ($R = 25$ and $R = 20$), the average claims were respectively, 120 and 145. For high penalty rates ($R = 50$ and $R = 80$), behavior did converge towards the Nash equilibrium. Thus, unless the reward/penalty is high enough, behavior departs from a Nash equilibrium.
2. The authors also find that *learning direction theory* (making adjustments in the current period that are a best response to the play of the opponent in the previous period) explains to an extent the current actions of players. However, it cannot account for the effect of R on the pattern of play.
3. A natural candidate explanation is QRE. We have seen several examples above, in which changes in the magnitude of the payoffs is predicted, under QRE, to alter the probabilities with which players play their strategies, even when the Nash equilibrium predicts no such effect. Indeed, the *responsiveness axiom* in Definition 13.8 includes this feature as an essential element of the regular quantal response function. When R is low, then errors that players make under a QRE are less costly in terms of payoffs, hence, players make higher claims with greater probability. The authors implement a learning model in which players update their beliefs about the opponent by observing their actions each period in a manner that is consistent with *fictitious play*.²⁰ A logit model is used to compute the probabilities of making various claims, as in (13.16); the maximum likelihood estimate of the noise parameter is $\lambda = 0.092$. In Figure 13.7, we show the average claims in each treatment (dots),

²⁰ See Part 5 of the book for models of fictitious play.




Figure 13.7 Average claims, logit predictions, and the Nash prediction in the traveller's dilemma game.

Source: Capra et al. (1999) with permission from the American Economic Association.

the Nash prediction (the horizontal dotted line at 80) and the logit equilibrium prediction (dark continuous curve). The logit prediction clearly explains the data much better than the Nash equilibrium.

13.3 Level- k and cognitive hierarchy models

Compared to a Nash equilibrium, in a QRE one relaxes the assumption that players play best replies with probability one. In a second class of models, however, players play their best replies with probability one but the beliefs and choices of players need not be mutually consistent. This is arguably less cognitively challenging, relative to computing a fixed point in beliefs and actions. During the early rounds of play in experiments when players play a new game against other players whose choices they are uncertain of, there is no reason why their beliefs should be consistent with the actions of other players. Hence, this class of models is likely to be more successful in explaining behavior in these situations. In this section we consider two closely related models in this class—the *level- k model* and the *cognitive hierarchy model*.

13.3.1 The level- k model

In order to explain the structure of level- k models, consider a set of players in some game. Players are heterogeneous and can be grouped into types (or levels) such as $L_0, L_1, L_2, \dots, L_K$ and a typical type/level is written as type/level $L_k, k = 0, 1, \dots, K$. Here the term “types” bears no relation to the ordinary usage of this term in neoclassical economics (e.g., types in adverse selection problems). Only players of type L_0 are *non-strategic*, in the sense that they do not play a best response to the actions of any other type. Typically, but not necessarily, type L_0 players are assumed to randomize over the set of available actions. In level- k models, for all $k = 1, 2, \dots, K$, type L_k players play a best response to the belief that all other players are of type L_{k-1} .

One generalization of the level- k model is the *cognitive hierarchy model* (CH) of Camerer et al. (2004). In the CH model, for types $k = 1, 2, \dots, K$, a type Lk player assumes that there is a distribution (typically Poisson) of all the lower type players $Lk - 1, Lk - 2, \dots, L0$. Each type then plays a best response to the distribution of lower types. The CH model is considered below in Section 13.3.3. For the moment, we restrict ourselves to the level- k model.

We may refer to type $L0$ as the *anchoring type*. The motivation for this terminology lies in the hierarchical structure of the model. The computation of the best response for type Lk players, $k \geq 1$ requires finding the best response of type $Lk - 1$ players, which in turn requires finding the best response of type $Lk - 2$ players, and so on until we arrive at type $L0$ players. Thus type $L1$ players directly anchor on type $L0$ players, while all other types of players do so indirectly. There might, of course, be no non-strategic $L0$ types in the population. Type $L0$ players could be entirely fictitious and exist only as models of other players in the minds of type $L1$ players. The exact specification of type $L0$ is important in applications.

The frequencies of the various types of players can be estimated from the data. A great deal of initial work found that the proportion of type $L0$ players is very small but some of the emerging empirical evidence estimates a higher fraction of such players. The broad message of the results is that the greatest proportion of players is of type $L1$, a somewhat smaller fraction is $L2$ players, and a relatively small fraction is of type $L3$. Types higher than $L3$ are rare. Since iterated beliefs are truncated at fairly low integer value in most experimental results ($k \leq 3$), these non-equilibrium models also seem to address the criticism of circularity of traditional game theoretic equilibrium concepts. For instance, Selten (1998) wrote: “Basic concepts in game theory are often circular in the sense that they are based on definitions by implicit properties...Boundedly rational strategic reasoning seems to avoid circular concepts. It directly results in a procedure by which a problem solution is found.”²¹

A well-known finding from many experimental games is that the initial response of players may be different from their later responses. Furthermore, in the presence of multiple equilibria, the choice of strategies in the initial rounds often determine which equilibrium outcome will be reached (see Section 12.4.9). One motivation for level- k models is that they can potentially inform us about the initial pattern of play, and the choice among multiple equilibria.

Critics argue that level- k models may allow for too much flexibility in the specification of the actions of types $L0$. One possible response to this criticism is to consider the fit of out-of-sample data under the relevant assumptions made on the behavior of type $L0$ players. Another response is to invoke various notions of *portability*. In the context of the level- k model, this involves the application of the level- k model to similar problems and checking to see if the predictive success in one problem carries over to a comparable problem. This is the strategy employed by Crawford and Iriberri (2007a) in the context of the hide-and-seek game that we outline in the problem set, below. They determine that O'Neill's (1987) card matching game and a related game in Rapoport and Boebel (1992) are similar to the hide-and-seek game. They find that when they apply the level- k model to these related games, the actions of type $L0$ are non-strategic. The similarity of the behavior of type $L0$ players across related games provides some, albeit imperfect, degree of confidence in the robustness of the results.²²

²¹ However, Selten may not consider level- k models to be models of bounded rationality since these models involve explicit optimization in order to compute best replies; see, Harstad and Selten (2013).

²² See also the discussion on portability across games in the context of the paper by Hargreaves Heap et al. (2014) in the introductory chapter.

Another possible criticism of the level- k models is that, they do not allow for the possibility of learning. In level- k models, since the beliefs of players are not in equilibrium, why should players continue to hold on to these beliefs in repeated play? There are several responses. First, the objective of level- k models is to explain the behavior of players in the early rounds of the experiments (Crawford et al., 2013).²³ Second, the real-world is not stationary, unlike the typical experimental environment in repeated play, hence, real world learning may be quite difficult in some cases. Third, not all learning experiments result in rational expectations in beliefs (see Part 5).

13.3.2 The p -beauty contest

The basic idea of the p -beauty contest game is outlined in John Maynard Keynes (1936, pp. 155–6):

professional traders may be likened to those newspaper competitions in which the competitors have to pick out the six prettiest faces from a hundred photographs, the prize being awarded to the competitor whose choice most nearly corresponds to the average preferences of the competitors as a whole . . . It is not a case of choosing those which, to the best of one's judgement, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practise the fourth, fifth and higher degrees.

The quote is highly suggestive of a level- k model. Keynes also seems to think that people predominantly classify as types $L0$ to $L3$ with higher types being rarer; this is in agreement with the experimental evidence from a range of games.

The outcome in p -beauty contest games depends on whether the equilibrium can be achieved in a finite or infinite number of steps. We follow Ho et al. (1998) in describing both cases. There are n participants. Participant $i = 1, 2, \dots, n$ is asked to choose a number, x_i , from some closed interval $[a, b]$. Denote the simple average of the guesses by $\bar{x} = \sum_{i=1}^n x_i$, and define the number $y = p\bar{x}$, where $p \geq 0$. The winners are those individuals whose guess is closest to y . Formally, the winners belong to the set

$$W = \{i : |y - x_i| \leq |y - x_j|; j = 1, 2, \dots, n\}. \quad (13.23)$$

The total prize money is equal to $n\pi$ where $\pi > 0$.²⁴ If the set W consists of $m \geq 1$ winners, then each winner receives a prize equal to $n\pi/m$. Notice the similarity of this setup with Keynes's example. To win, one must guess "the average opinion about the average guess."

The equilibrium outcome in the p -beauty contest depends on whether $p < 1$ or $p > 1$, as shown in the next proposition.

²³ In the real world, players often make decisions on a fairly irregular basis, e.g., household decisions on pension plans, marriage, number of children, purchase of consumer durables such as cars, and partly illiquid assets such as houses. Thus, learning opportunities are not unlimited. Similarly, many decisions of firms are also of a very irregular nature that afford little learning opportunities, such as merger decisions, bidding in an auction for bandwidth rights, and fundamental corporate restructuring. Arguably, irregular and infrequent decisions might account for a large part of economic activity.

²⁴ In most experiments, the reward does not depend on the number of participants. However, Ho et al. (1998) were interested in testing for the effect of group size.

Proposition 13.26 Suppose that there are n participants in the p -beauty contest and each chooses a number $x_i \in [a, b]$. The set of winners satisfies (13.23). (i) When $p < 1$, the Nash equilibrium outcome is $x_i = a$ for all i . (ii) When $p > 1$, the Nash equilibrium outcome is $x_i = b$ for all i .

Proof: Since all individuals are identical we can restrict attention to symmetric equilibria.

(i) Let $p < 1$. Suppose that for all i , $x_i = x$, so $\bar{x} = x$. There are two possibilities. First, $x = a$, so $y = pa$. In this case, since $p < 1$ no player can improve his payoff by choosing a smaller number, so we are done. However, if $x > a$, then unilateral deviations by players can improve payoffs. Suppose that player i chooses a guess $x' = x - \epsilon \geq a$, where $\epsilon > 0$ is a small number, and all other players choose a guess x . Then, $y = p \left[\frac{(n-1)x}{n} + \frac{x'}{n} \right]$. Reducing his guess by ϵ , player i induces a reduction in y of $\frac{\epsilon p}{n} < \epsilon$. Hence, the guess of player i is closest to the new value of y , making him the unique winner. Thus, the unique symmetric equilibrium is one where $x_i = a$, $i = 1, 2, \dots, n$. (ii) The proof is analogous to (i) so it is omitted. ■

In experiments, it is often the case that subjects do not take account of the effect of their actions on the average action. Proposition 13.26 continues to hold in this case as well.

Next, we consider the predictions of the level- k model. Suppose that there are n players in the p -beauty contest, where n is very large, and each player ignores the effect of his guess on the average guess. Each player chooses a number $x_i \in [a, b]$. The set of winners satisfies (13.23). All players are level- k thinkers, so they belong to types $L_0, L_1, L_2, \dots$

Suppose that $[a, b] = [0, 100]$, and $p = 2/3$. Suppose that type L_0 players choose uniformly randomly over $[0, 100]$. Type L_1 players believe that all other players are type L_0 , whose average choice is 50, so, they believe that $y = 50(\frac{2}{3}) = 33.33$, which is also their optimal guess. Type L_2 players believe that all others are type L_1 , so $y = 50(\frac{2}{3})^2 = 22.22$, which is also their optimal guess. In general, type L_k players guess the number $50(\frac{2}{3})^k$. Hence, if these types of players are present in the data we should expect to observe spikes at the numbers $50(\frac{2}{3})^k$, $k = 1, 2, \dots$. Type L_∞ players will choose a guess of 0 that coincides with a Nash equilibrium.

Now suppose that $[a, b] = [0, 100]$, $p > 1$, and continue to assume that type L_0 players choose uniformly randomly over $[0, 100]$. If $50p > 100$, then type L_1 players will win for sure if they guess 100, so this is their optimal guess; all higher types will also optimally choose to guess 100. If $50p < 100$, then type L_1 players will choose the guess $50p$. Type L_2 players believe that all others are type L_1 players, so they choose the guess $\min[50p^2, 100]$. If $\min[50p^2, 100] = 100$, then all higher types will also choose to guess 100. Similarly, type L_3 players will choose $\min[50p^3, 100]$ and if $\min[50p^3, 100] = 100$, then all higher types also choose to guess 100. Continuing with this recursion, type L_∞ will choose the guess 100, which is also the Nash equilibrium. Hence, the sequence of guesses will be $50, 50p, 50p^2, \dots, 100$.

In the earliest test of the p -beauty contest, Nagel (1995) found that for $[a, b] = [0, 100]$ and $p = 2/3$, the mean guess with a subject pool of 14–16 German students was 35. There were spikes in the data at the choices 50, 33, 22, which identified level 0, 1, and 2 thinkers, and the respective percentage of such types in her dataset was 13, 44, and 39. Figure 13.8 shows a histogram of choices from Nagel's data for the case $p = 2/3$ (she also reports results for other values of p such as $1/2$ and $4/3$).

Ho et al. (1998) extend Nagel's results. To illustrate their experimental findings, consider two games.

1. An infinite threshold game IT in which $a = 0$, $b = 100$, $p = 0.7$; this is denoted by $IT([0, 100], 0.7, n)$. The highest number that a type L_0 may choose is 100. Hence, given




Figure 13.8 Results for the p -beauty contest game when $p = 2/3$.

Source: Nagel (1995) with permission from the American Economic Association.

$p = 0.7$, type $L1$ will not choose a number between $(70, 100]$, type $L2$ will not choose a number between $(49, 100]$, type $L3$ will not choose a number between $(34, 100]$ and so on. In general type Lk will not choose a number between $(0.7^k 100, 100]$. The Nash equilibrium choice of 0 is achieved as $0.7^k \rightarrow 0$, i.e., by type $L\infty$, and requires an infinite number of steps of iterated dominance.

2. A finite threshold game FT in which $a = 100$, $b = 200$, $p = 1.3$; this is denoted by $FT([100, 200], 1.3, n)$. The lowest number that a type $L0$ may choose is 0. Hence, given $p = 1.3$, type $L1$ never chooses a guess from the interval $[100, 130)$, type $L2$ never chooses to guess from the interval $[100, 169)$, type $L3$ never chooses to guess from the interval $[100, 219)$. But since $219 > 200$ the optimal guess of type $L3$ and all higher types is 200. Hence, we need only three rounds of iterated deletion of dominated strategies to achieve the equilibrium.

Ho et al. (1998) give results for a range of values of p , such as 0.7, 0.9, 1.1, 1.3. Some of their findings can be summarized as follows.²⁵

First, based on first round data, the choices made by the players are clustered around the midpoint of the interval $[a, b]$. For the case $p < 1$, across all conditions (i.e., varying the size of the interval $[a, b]$ and the number of participants, n), the respective percentages of types $L0, L1, L2, L3$ were 16, 21, 13, 50. The corresponding numbers for the case $p > 1$ are 22, 31, 13, 34. Hence, all the players in the sample can be accounted for by types $L0$ to $L3$; in particular, there are not types $L4$ or above. By contrast, for any p , we obtain the Nash equilibrium prediction for type $L\infty$. Choices converge towards equilibrium only over time. However, this does not prove that players attain higher levels of iterated rationality. Rather it could simply be due to learning effects. A learning model that is estimated by the authors confirms this. In this model, there are two main types of

²⁵ See also Camerer (2003a) for a nice exposition of these findings.




Figure 13.9 Results for the p-beauty contest using diverse subject pools with $p = 2/3$, $a = 0$, $b = 100$.

Source: Reprinted from *Trends in Cognitive Sciences*, 7: 225–31. Camerer, C.F. “Behavioral studies on strategic thinking in games.” Copyright © 2003, with permission from Elsevier.

players: Level-0 players who simply adapt to past play and level-1 players who best respond to the play of level-0 players.

Second, choices are closer to equilibrium in finite threshold games, and converge faster to equilibrium relative to infinite threshold games. This is due possibly to the higher degree of rationality required in the latter. These differences are statistically significant. Larger groups ($n = 7$ rather than $n = 3$) make choices that are closer to equilibrium.

Third, we report in Figure 13.9, taken from Camerer (2003b), the choices made by Caltech graduates, portfolio managers, and economics Ph.D.s for the case $p = 2/3$, $a = 0$, $b = 100$. Most people would impute a high degree of cognitive ability to this group (and possibly rationality too). Yet these special groups show significant spikes around the numbers 50, 33, 22 that are also consistent with level- k subjects.

Fourth, experienced and inexperienced subjects demonstrate similar choices in the first round of the experiments, but the choices of the former converge more quickly to equilibrium. Experienced subjects must have played a similar game in the past. For instance players who play the IT game could have played the FT game before and vice versa.

Bosch-Domènech et al. (2002) compare the results of p-beauty contest games from two different kinds of experiments. (1) Experiments with lab subjects. (2) Three different sets of experiments with newspaper subjects carried out separately in 1997 by Thaler; Bosch-Domènech and Nagel; and Selten and Nagel. The newspaper experiments are treated by the authors as field experiments in which one has far less control over the subject pool. Two kinds of data are considered: choice data, i.e., actual guesses made by players in the experiments and data arising from unstructured comments of the participants to justify the choices that they make. The main observations are as follows. The results from choice and comments data from the independent newspaper experiments are similar and show spikes at 50, 33, and 22 that are consistent with level- k subjects.

The subjects in experiments in the lab are typically not allowed to communicate with each other. However, some of the newspaper subjects attempted to form coalitions but it did not affect the results (see, however, the findings of Burchardi and Penczynski (2014), below). The frequency of equilibrium choices is highest among theorists in the lab (15%), followed by newspaper subjects (6%) and undergraduate lab subjects (1%). The most successful newspaper subjects (in terms of winning the prize) were those who chose to run their own experiments and chose to submit numbers close to the winning numbers in their private experiments. The authors comment (p. 1698) that “To our shame, we were taught the very lesson that we, experimentalists, are trying to teach our fellow economists: when in doubt, run experiments.”

Grosskopf and Nagel (2008) introduce the two-person p -beauty contest game for $p = 2/3$. The solution to this game is even less cognitively challenging than a n -player game, $n \geq 3$. The reason is that if the winning number is the one which is closer to $p < 1$ times the average, then the smaller number always wins, irrespective of the size of the larger number chosen. For this reason, it is cognitively easier to spot that choosing the number zero is a weakly dominant strategy.

Two distinct subject pools were chosen. First, students who did not have a prior exposure to the p -beauty contest. Second, professionals (including game theorists) at conferences where this paper was presented; four of these five conferences were economics conferences. The results are as follows. Professionals do better than students in choosing the Nash outcome (36.2% versus 9.85%). However, another way of looking at this result is that about 90% of the students and about 63% of professionals play weakly dominated strategies (i.e., a choice that is not equal to zero). Professionals choose zero more often (in a statistically significant sense) in the two-player game than in the $n > 2$ player game (data for the $n > 2$ comes from Bosch-Doménech et al., 2002), hence, they benefit relatively more from a reduction in the cognitive requirements of a two-player game. However, there is no statistical difference in the cumulative density functions of choices of students and professionals in the cases $n = 2$ and $n > 2$.

Breitmoser (2012) questions the suitability of a level- k approach for the p -beauty contest. He argues for an explanation of the data that is based on quantal response equilibrium (see Section 13.2) and models of noisy introspection (see Section 13.7.5). His main point is that by ignoring the effect of one's own guess on the average guess, the predictions of level- k models are mis-specified. However, it appears to be the case, in most experiments, that subjects ignore the effects of their own actions on the average. Thus, the predictions in the standard level- k model are made under the joint hypothesis that subjects follow level- k models and use the heuristic of ignoring the effects of their actions on the average.²⁶ The empirical tests that we report in Section 13.4.5, for instance, by Costa-Gomes et al. (2001) and Costa-Gomes and Crawford (2006) assume that subjects follow this heuristic. Yet, they find strong support for level- k models.

13.3.3 The cognitive hierarchy model (CH)

Camerer et al. (2004a) generalize the level- k model by assuming that each type Lk player, $k = 1, 2, \dots$ believes that other players belong to a distribution of lower types $Lk - 1, Lk - 2, \dots, L0$. This is the *cognitive hierarchy model* (CH). The authors give two motivations for the CH model. First, as k increases, the expectations of each type in the level- k model may appear increasingly irrational because they ignore the presence of an increasingly larger number of types, while in the CH model, they appear increasingly rational. Second, the predictions of the CH model appear

²⁶ I am grateful to Vincent Crawford for clarifying these issues in a personal communication.

as precise as those of the level- k model in a large class of games, but in some, such as the entry game, they appear to be better. However, in a recent survey, Crawford et al. (2013) write that:²⁷

It would be of interest to discriminate between level- k and CH models, but so far no study has clearly separated their predicted decisions. But most level- k estimates that make adequate allowance for errors assign $L0$ a much lower frequency than CH estimates do, often zero. Thus a CH model's Poisson constraint will often be binding, and a level- k model, which imposes no such restriction, will fit better (see however Camerer, Ho, and Chong 2004, Section II.B; and Chong, Camerer, and Ho 2005).

We now explain the CH model in more detail. Suppose that players are exclusively of types Lk with $k = 0, 1, 2, \dots$ and there is a distribution $f(k)$ of these types. In the CH model, a type Lk player believes that all other players are of types $Lk - 1, Lk - 2, \dots, L0$. Given the evidence on cognitive limits, it is plausible to assume that the fraction of higher types is lower relative to the fraction of lower types. One way of capturing this is to assume that the ratio $f(k)/f(k - 1)$ declines at a rate that is proportional to k . This property is satisfied by the Poisson distribution,²⁸

$$f(k) = \frac{e^{-\tau} \tau^k}{k!}; k = 0, 1, 2, \dots, \quad (13.24)$$

which is a one-parameter family of distributions such that the mean and variance are given by τ .

The belief of type Lk players about the proportion of players of a lower type Lh ($h < k$) is denoted by $g_k(h)$ and is given by the following normalized distribution

$$g_k(h) = \frac{f(h)}{\sum_{l=0}^{k-1} f(l)}. \quad (13.25)$$

Thus, players use the correct distribution f in forming beliefs about the fractions of players who perform relatively lower steps of thinking. As k becomes larger, $\sum_{l=0}^{k-1} f(l)$ approaches 1, so the gap between $g_k(h)$ and $f(h)$ diminishes. In this sense, higher types have increasingly rational expectations about others.

Denote by s_k^j , strategy j of a type Lk player and by s_h^l the strategy of an opponent, which gives the strategy profile (s_k^j, s_h^l) . Suppose that a player of type Lk has J_k pure strategies. Denote by $p(s_h^l)$ the probability that a player of type Lh , $h < k$, chooses the strategy s_h^l . Players of type $L0$ are assumed to randomize uniformly over their strategies, hence, the probability that any player of type $L0$ chooses strategy s_0^j is denoted by

$$p(s_0^j) = 1/J_0. \quad (13.26)$$

A player of type Lk knows that the opponent who plays s_h^l is likely to be a type Lh player, $h < k$, with probability $g_k(h)$. In turn, a type Lh player plays s_h^l with probability $p(s_h^l)$. Thus, in the mind of a player of type Lk , the probability of the strategy profile (s_k^j, s_h^l) is given by

²⁷ In the quote we have omitted an explanatory footnote.

²⁸ Crawford et al. (2013, p. 15) point out that while the Poisson distribution is not very restrictive in some domains, it is quite restrictive in others.

$$P_k(s_h^l) = \sum_{h=0}^{k-1} g_k(h) p(s_h^l). \quad (13.27)$$

The expected payoff of a player of type Lk from the strategy profile (s_k^j, s_h^l) is given by

$$E\pi_k(s_k^j, s_h^l) = \sum_{s_h^l \in S_{-k}} \pi_k(s_k^j, s_h^l) P_k(s_h^l), \quad (13.28)$$

where $\pi_k(s_k^j, s_h^l)$ is the payoff of player k and S_{-k} is the set of pure strategies of the opponent.

Players are assumed to best respond, conditional on their (non-equilibrium) beliefs. Thus, if $s_k^* \in \arg \max E\pi_k(s_k^j, s_h^l)$, then $p(s_k^*) = 1$. The model is solved recursively. From (13.26), the strategy of type $L0$ players is given by $p(s_0^j)$. Using this data, we can solve for g_1 and P_1 , which enables us to solve for the optimal strategy of type $L1$ players. Given the optimal strategies of types $L0$ and $L1$ we compute g_2 and P_2 that allows us to solve for the optimal strategy of type $L2$ players and so on.

Consider the class of dominance solvable games. The prediction of the CH model with a Poisson $f(k)$ when the single parameter of the Poisson distribution $\tau \rightarrow \infty$ is identical to the Nash equilibrium. This prediction is reached by finite deletion of dominated strategies. To see this, use (13.24) to write:

$$\frac{f(k-1)}{f(k-j)} = \frac{(k-j)!}{(k-1)!} \tau^{j-1}, j \geq 1. \quad (13.29)$$

Consider a type Lk player. As $\tau \rightarrow \infty$, type Lk imagines that all other types are type $Lk-1$. In this case, the model becomes essentially the simpler level- k model. We show in Section 13.4.4 that the outcome of a dominance solvable game with level- k players is identical to the Nash equilibrium outcome.

In coordination games such as the stag-hunt game, the CH model gives important insights into the effects of group size. It turns out that when group size increases, the likelihood that the inefficient risk dominant equilibrium gets played increases. This is able to explain the empirical results of Huyck et al. (1990); see one of the exercises at the end of the chapter.

From (13.29), an increase in τ , for any player of type Lk , increases the relative proportion of higher types among the set of types $Lk-j$, $j \geq 1$. The parameter τ can be estimated from the data. In p-beauty contest games, collating results over several studies, Camerer et al. (2004a) report mean and median values of τ equal to, respectively, 1.30 and 1.61. These values are supportive of a CH model; recall that level- k models, by contrast, would require much higher values of τ . Also, subject pools that comprise individuals with greater analytical skills or greater education are also characterized by higher values of τ . An increase in stakes also increases the estimated value of τ .

Camerer et al. (2004a) consider early round data, with no feedback, for a very large number of games (33 matrix games and 22 games that have a mixed strategy equilibrium). They pit the predictions of a Poisson CH model and the Nash equilibrium against each other. The two main findings are as follows. First, when the Nash predictions are good approximations, the CH model is almost as accurate. Second, when the Nash predictions are poor approximations of the experimental results, then the predictions of the CH model are usually more accurate (Crawford et al. (2013)).

Table 13.1 Game 8 from Costa-Gomes et al. (2001).

	L	R	Data	Nash	CH
T	45, 66	82, 31	.92	1	.82
TM	22, 14	57, 55	0	0	.06
BM	30, 42	28, 37	0	0	.06
B	15, 60	61, 88	.08	0	.06
Data	.64	.36			
Nash	1	0			
CH	.55	.45			

Source: Camerer et al. (2004a).

Consider, for instance, the results given by Camerer et al. (2004a) for game 8 from Costa-Gomes et al. (2001) in Table 13.1. This is a two-player normal form game in which the players simultaneously choose their strategies. The unique Nash equilibrium of the game in pure strategies is (T,L). The experimental data, reported in the right hand side and the lower panels shows that most row players (92%) play T, but 36% of the column players do not play their part of the Nash profile, L, preferring instead to play the alternative strategy R. Suppose that $\tau = 1.7$ which is not unrepresentative of the estimates (Camerer et al. (2004a)). Then, from (13.24), we get that

$$f(0) = \frac{e^{-1.7}(1.7)^0}{0!} \approx 0.18; f(1) = \frac{e^{-1.7}(1.7)^1}{1!} \approx 0.31; f(2) = \frac{e^{-1.7}(1.7)^2}{2!} \approx 0.26. \quad (13.30)$$

Column players of type $L1$ believe that they are facing type $L0$ row players who randomize equally between their strategies T, TM, BM, B. Hence, the expected payoffs of a type $L1$ column player from playing L and R, respectively, are 45.5 and 52.75, so, they play R. Now consider type $L1$ row players who believe that they are facing type $L0$ column players who randomize equally among L and R. Since T is a dominant strategy, it is also the best reply for type $L1$ row players.

Row and column players of type $L2$ in the CH model believe that they face a heterogeneous population of type $L0$ and type $L1$ opponents. Using (13.25), (13.30), the normalized fractions of opponents in the mind of a type $L2$ player are

$$g_2(0) = \frac{0.18}{0.18 + 0.31} = 0.36; g_2(1) = 1 - g_2(0) = 0.64. \quad (13.31)$$

Thus, the payoffs of a column player of type $L2$ from the two strategies, L and R are respectively

$$\begin{aligned} L : 0.36 \left(\frac{66 + 14 + 42 + 60}{4} \right) + 0.64(66) &= 58.62. \\ R : 0.36 \left(\frac{31 + 55 + 37 + 88}{4} \right) + 0.64(31) &= 38.83. \end{aligned}$$

Thus, it is optimal for a column player of type $L2$ to play L. Playing L is also optimal for any column player of type Lk , $k \geq 3$. Row players of any type Lk , $k \geq 1$ find it optimal to play the strictly dominant strategy T. Since type $L0$ column players randomize equally between L and R,

so half of them choose R. Also, as shown above, all type $L1$ column players also play R. All other types of column players play L. Using (13.30), the total fraction of column players who play the strategy R is $\frac{1}{2}(0.18) + 0.31 = 0.40$. This is close to the findings of the experimental data that 36% of the column players play R.

Type $L0$ row players play T with probability 0.25. Since $f(0) \approx 0.18$ and all other types of row players play T under a Poisson CH, a fraction 95.5% of the players are predicted to play T. The experimental finding of 92% of the players playing T is in line with this prediction. Thus, in this illustrative example, where Nash predicts well (for row players), the CH model predicts equally well and where Nash does not predict well (for column players), the prediction of the CH model is superior. An exercise at the end of this chapter gives another example from Game 4 of Binmore et al. (2001) that illustrates a similar point.

13.4 Applications of level- k and CH models

In this section we consider further applications of level- k models.

13.4.1 *Is the power of focal points limited?*

Conventional game theory does not allow for the names of strategies or labels on decisions to have “salience.” Indeed, the equilibrium of the game is independent of the names that one might choose for various strategies. Schelling (1960) argued that some actions and labels could be more salient than others and may facilitate coordination in games. We have considered the evidence on these issues in the previous chapter. In Schelling’s experiments, two subjects were simultaneously asked to choose a place where they would meet the other subject in New York. No communication was allowed between the subjects. If the choices of the two subjects coincided, then they both received an identical positive payoff. Otherwise they both received zero. Despite the very large number of places that one could name in New York, a common choice was the Grand Central Station. This allowed for a high rate of coordination in a game where, ex-ante, one might have expected a low rate of coordination. Hence, in this context, the choice of Grand Central Station has particular salience; such a salient choice may also be termed as a *focal point*.

Crawford et al. (2008) use level- k models to examine 2×2 coordination games with symmetric and asymmetric payoffs. In asymmetric 2×2 coordination games, individual players may derive a greater payoff from one of the strategies relative to the others, hence, such a strategy is said to exhibit *payoff salience*. However, it is quite likely that another strategy may exhibit *label salience* that conflicts with *payoff salience*. So what outcome do we expect? The main aim is to show that when payoff salience opposes label salience, then even minute asymmetries in the game could significantly reduce the power of focal points.

In a pilot study with unpaid subjects in experiments, the authors first used a 2×2 coordination game called Chicago Skyscrapers. This was similar to the idea in Schelling’s work; each player simultaneously chooses a meeting point with the other player, a common choice is rewarded and no communication is allowed. The difference is that players were given a restricted, easier, choice; each player had to choose between the Sears Tower (a salient Chicago landmark) and a less salient building across the road, the AT&T building. The payoff matrices are shown in Figure 13.10. Player 1 is the row player and player 2 is the column player. The choices made by the subjects are shown in brackets next to each strategy.

Symmetric game			Slightly asymmetric game			Moderately asymmetric game		
	Sears (90%)	AT&T (10%)		Sears (58%)	AT&T (42%)		Sears (47%)	AT&T (53%)
Sears (90%)	100, 100	0, 0	Sears (61%)	100, 101	0, 0	Sears (61%)	100, 110	0, 0
AT&T (10%)	0, 0	100, 100	AT&T (39%)	0, 0	100, 100	AT&T (39%)	0, 0	110, 100

Figure 13.10 The non-incentivized Chicago skyscrapers game in Crawford et al. (2008). With permission from American Economic Association.

Symmetric game			Slight asymmetries		
	X (76%)	Y		X (28%)	Y
X (76%)	5, 5	0, 0	X (78%)	5, 5.1	0, 0
Y	0, 0	5, 5	Y	0, 0	5.1, 5

Moderate asymmetries			Large asymmetries		
	X (61%)	Y		X (60%)	Y
X (33%)	5, 6	0, 0	X (36%)	5, 10	0, 0
Y	0, 0	6, 5	Y	0, 0	10, 5

Figure 13.11 Results with incentivized experiments in Crawford et al. (2008). With permission from American Economic Association.

For the symmetric game, the results are consistent with Schelling's insights; 90% of the row and the column players choose the salient strategy, Sears Tower. Thus, in 81% of the cases, successful coordination on the Sears Tower takes place. However, the power of focal points becomes weaker as the degree of asymmetry in the game increases, which also increases the tension between payoff and label salience. In the slightly asymmetric and the moderately asymmetric games, coordination is achieved on the Sears Tower, respectively, in 35.4% and 23.5% of the cases.

The authors then conducted incentivized experiments in which Sears and AT&T were replaced, respectively, with the alphabets X and Y. The expectation is that X is more salient than Y. As the authors say (p. 1446): "We anticipated that X would be more salient than Y, because X marks the spot and X precedes Y in the alphabet." The nature of the results was identical to the Chicago Skyscrapers game; Figure 13.11 which also shows the percentage of row and column players choosing the action X, shown in brackets.

The results suggest the following conclusions. First, the choices made in the symmetric game suggest that X is a salient action, but to a lesser extent as compared to the Sears Tower in the Chicago Skyscrapers game. Second, under slight payoff asymmetries most subjects prefer the payoff salient action of their opponent. However, as the degree of asymmetry increased, subjects switched to favoring their own payoff salient action. This is reflected in the coordination rates across games. Let the coordination rate in a game be the joint play of the strategy profiles (X, X)

and (Y, Y). In the symmetric game, the coordination rate is $(0.76)^2 + (0.24)^2 = 64\%$. However, in the asymmetric games, the coordination rates are lower but increase as the degree of asymmetry successively increases (respectively, 38%, 46%, and 47%). This pattern of coordination is the reverse from the Chicago skyscrapers games (where it fell rather than increased).

These stylized facts cannot be explained using the QRE because there is no concept of the salience of a payoff label in that theory. Crawford et al. (2008) give an explanation in terms of the level- k model. They make the plausible assumption that the behavior of type $L0$ players may differ in the symmetric and the asymmetric games, in the following manner.

A1 If the coordination game is symmetric, then type $L0$ players choose the “label salient action” with probability $p \geq 0.5$.

In the slightly asymmetric game in Figure 13.11 the payoff salient actions are Y for the row player and X for the column player. The behavior of type $L0$ players with respect to their payoff salient actions is given in the next assumption.

A2 If the coordination game is asymmetric, then type $L0$ players choose the “payoff salient action” with probability $p \geq 0.5$. Using the slightly asymmetric and the moderately asymmetric games, define

$$\underline{p} = \frac{5.1}{5.1 + 5} = 0.505; \quad \bar{p} = \frac{6}{6 + 5} = 0.545. \quad (13.32)$$

Assumption A2 requires that

$$\underline{p} < p < \bar{p}, \quad (13.33)$$

i.e., the payoff bias of type $L0$ players in the asymmetric games is quite modest.

We first show that Schelling’s insight about focal points holds perfectly in the symmetric coordination game.

Proposition 13.27 *Under assumption A1, all row and column players choose X in the symmetric coordination game.*

Proof: Consider the symmetric coordination game shown in Figure 13.11. Using assumption A1, all row and column players who are type $L0$ choose the label salient action X with a probability $p > 0.5$. All type $L1$ players choose a best response to the play of type $L0$ players. Their expected payoff from choosing X is greater than that from choosing Y (since $5p > 5(1-p)$), so they choose X with probability 1. Hence, by an analogous argument, all types Lk , $k \geq 2$ of row and column players also play X with probability 1. ■

Proposition 13.28 *Consider the slightly asymmetric game in Figure 13.11 and suppose that assumption A2 holds. Row and column players of type $L1$ choose their partner’s payoff salient actions, respectively X, Y. Row and column players of type $L2$ choose their own payoff salient action, respectively Y, X.*

Proof: From assumption A2, we know that row and column players of type $L0$ play their payoff salient actions (Y for row player and X for column player) with probability p that is given in (13.33). Thus, the expected payoff from playing X for row players of type $L1$ is relatively greater

because $5p > 5.1(1 - p) \Leftrightarrow p > \underline{p}$, which holds from (13.33). An analogous argument shows that the expected payoff of column players of type $L1$ from playing their partner's payoff salient action, Y , is relatively greater. Players of type $L2$ play their best replies, conditional on the behavior of type $L1$ players. Row players of type $L2$ know that column players of type $L1$ play Y to which the best response is their own payoff salient action, Y . Similarly, column players of type $L2$ know that row players of type $L1$ play X , to which the best response is their own payoff salient action X . ■

The results for the games with moderate and large asymmetries are similar, so, we derive the results for the game with moderate asymmetries only.

Proposition 13.29 *Consider the moderately asymmetric game under assumption A2. Players of type $L1$ play their own payoff salient action (row players play Y and column players play X). Players of type $L2$ play the payoff salient action of their opponents (row players play X , column players play Y).*

Proof: From assumption A2, players of type $L0$ play their payoff salient actions with probability p given in (13.33). Conditional on these beliefs, players of type $L1$ play as follows. Row players of type $L1$ prefer to play Y because $5p < 6(1 - p) \Leftrightarrow p < \bar{p}$, which is true from (13.33). By the same reasoning, column players of type $L1$ optimally play X . Since row players of type $L1$ play Y , the best reply of column players of type $L2$ is to play Y . By an analogous reasoning, the best reply of a row player of type $L2$ is X . ■

Crawford et al. (2008), show that under assumptions A1, A2, the qualitative results of the level- k model match their experimental data. They also show that for plausible distributions of the frequencies of the various types, the model can also closely replicate the quantitative findings from the experimental data.

13.4.2 *Coordination and communication with boundedly rational players*

There has been an interest in the role that communication can play in affecting the degree of coordination in games where coordination failures are a possibility.²⁹ The development of language is regarded by biologists and anthropologists as one of the critical features of human evolution. Perhaps it allowed primitive man to coordinate individual actions, in important collective actions such as hunting.

Two essential components of human communication and language are the ability to listen and speak. In the context of the coordination games typically considered in economics, *listening* often involves an attempt to figure out the intentions of the other players and *speaking* typically involves an attempt to influence the beliefs/actions of others. The literature on cheap talk games (e.g., Crawford and Sobel (1982); Farrell (1987, 1988)) makes progress in this direction.

Farrell (1987), for instance, focuses on problems of coordination in the Battle of Sexes (BOS) game. The robustness of these results was checked in a more extended class of games by Rabin (1994). A central concern of this endeavor was to check if coordination could be increased in the BOS game by *preplay communication*. Preplay communication took the form of cheap talk, i.e., messages that were non-binding and also costless. One or both players could engage in cheap talk; this is respectively known as *one-way* and *two-way* communication. While the language-content of the underlying messages was not specified, which might be objected to by anthropologists and

²⁹ See for instance, Ellingsen and Östling (2010) and Crawford (2007), Crawford (2003).

	H	D
H	0, 0	a, 1
D	1, a	0, 0

Figure 13.12 Battle of Sexes game in Crawford (2007); $a > 1$.

sociologists, it was taken as given that a preexisting common language made the messages unambiguous. This literature considers specific game theoretic equilibria and also identifies the set of rationalizable strategies. An important finding is that rationalizable preplay communication does not ensure either coordination or a Pareto-efficient equilibrium, although such communication may enhance coordination.

Following Crawford (2007) we now apply the level- k model to the Battle of Sexes (BOS) game shown in Figure 13.12. There are two strategies for each player, H and D. The respective payoffs of the row and the column players are shown in Figure 13.12, where $a > 1$.

THE OUTCOME UNDER CLASSICAL GAME THEORY

The unique symmetric equilibrium for the BOS game is in mixed strategies. Each player plays H with probability $p = a/(1 + a)$. Coordination occurs when either of the two outcomes, (H,D) or (D,H) are played. Given the equilibrium mixed strategy profile, these two outcomes occur with an identical probability of $p(1 - p)$, hence, the expected coordination rate, in equilibrium, is $2p(1 - p)$. Substituting for p gives the expected coordination rate,

$$\rho(a) = \frac{2a}{(1 + a)^2}. \quad (13.34)$$

A simple calculation gives rise to the following remark.

Remark 13.2 *The coordination rate, $\rho(a)$, is maximized at $a = 1$ and the maximum coordination rate is given by $\rho_{\max} = 0.5$. Since, $\lim_{a \rightarrow \infty} \rho(a) = 0$ the payoff from each pure strategy in the mixed strategy profile is $a/(1 + a)$.*

THE OUTCOME UNDER A LEVEL-K MODEL WITH NO COMMUNICATION

Crawford (2007) assumes that there are four types of level- k players $L1, L2, L3, L4$. The anchoring type $L0$ is not physically present, however, it exists in the mind of type $L1$ as a model of what other players are like. Type $L1$ believes that type $L0$ players randomize uniformly over their pure strategies.

Proposition 13.30 *Consider the BOS game in Figure 13.12. The matrix of equilibrium actions in a level- k model with no communication is shown in Table 13.2.*

Proof: Conditional on the behavior of type $L0$ players, type $L1$ players receive a higher payoff from playing H rather than D (because $a > 1$). Conditional on the strategy of type $L1$ players, type $L2$ players play their best response, D. Analogously, conditional on the strategy of type $L2$ players, type $L3$ players choose H. Similarly, type $L4$ players play their best response, D, to the strategy

Table 13.2 Equilibrium actions in a level- k model without communication.

Types	L1	L2	L3	L4
L1	H, H	H, D	H, H	H, D
L2	D, H	D, D	D, H	D, D
L3	H, H	H, D	H, H	H, D
L4	D, H	D, D	D, H	D, D

of type $L3$ players. The resulting strategy profiles of types $L1, L2, L3, L4$ in pairwise interaction are shown in Table 13.2. For instance, the strategy profile in the (3,2)th cell of the matrix, (H,D), shows that if a type $L3$ row player plays against a type $L2$ column player, then they play their respective optimal strategies H and D. ■

Recall our assumption in this section that type $L0$ players exist only in the minds of type $L1$ players as a model of other players. The extent of cooperation, thus, depends on the frequency of types $L1, L2, L3, L4$. From Proposition 13.30, the two cooperative outcomes (H,D) or (D,H) are possible only if a type $L1$ or $L3$ player plays against an opponent of type $L2$ or $L4$ (i.e., an odd numbered type plays against an even numbered type of player). Denote the composite frequency of types $L1$ and $L3$ by λ and the composite frequency of types $L2$ and $L4$ by $1 - \lambda$. The probability that $L1$ or $L3$ row players meet $L2$ or $L4$ column players is $\lambda(1 - \lambda)$. Similarly the probability of the complementary event that $L2$ or $L4$ row players meet $L1$ or $L3$ column players is also $\lambda(1 - \lambda)$. Hence, the total probability of cooperation is

$$\hat{\rho}(\lambda) = 2\lambda(1 - \lambda). \quad (13.35)$$

Remark 13.3 $\hat{\rho}$ is maximized at $\lambda = 0.5$, and the maximum value of coordination is $\hat{\rho}_{\max} = 0.5$. This gives the maximum coordination rate under a level- k model.

Using Remarks 13.2 and 13.3, we can compare the coordination rates, respectively, in the mixed strategy Nash equilibrium and in the level- k model. Since $\rho(1) = \rho_{\max} = 0.5$, $a > 1$, and $\rho'(a) < 0$, it follows that $\rho(a) < 0.5$ for all $a > 1$. Thus, $\hat{\rho}(0.5) = \hat{\rho}_{\max} > \rho(a)$ for all $a > 1$. Thus, $\lambda = 0.5$ is a sufficient condition for the coordination rate in a level- k model to be relatively greater than in the mixed strategy Nash equilibrium. Furthermore, since, $\rho'(a) < 0$, $\lim_{a \rightarrow \infty} \rho(a) = 0$ and $\hat{\rho}$ is independent of a , thus, as the degree of asymmetry in the game increases, the gaps in the coordination rates under the level- k model and the mixed strategy Nash equilibrium continue to widen. Players using the level- k model do not coordinate on the basis of risk dominance or payoff dominance. Coordination, when it happens, is entirely accidental. Since level- k players use a non-equilibrium model, ex-post, they are often surprised by the outcome.

THE OUTCOME UNDER A LEVEL- k MODEL WITH TWO-WAY COMMUNICATION

Suppose that there is a preexisting language that is common knowledge. In this language suppose that the message h sent by any player means “I intend to play H,” while the message d means “I intend to play D.” Let r denote the random message $0.5h + 0.5d$. There are two stages. In the first stage, players exchange messages that are chosen simultaneously. In the second stage, players choose actions, H, D, conditional on having received the message in the first stage. Let R denote the random strategy, $0.5H + 0.5D$. There could conceivably be many rounds of communication

Table 13.3 Optimal choice of messages and actions in the two-stage game with communication.

Types (messages)	L_1 (r)	L_2 (h)	L_3 (d)	L_4 (h)
L_1 (r)	R, R	D, H	H, D	D, H
L_2 (h)	H, D	H, H	H, D	H, H
L_3 (d)	D, H	D, H	D, D	D, H
L_4 (h)	H, D	H, H	H, D	H, H

followed by the second stage choice of actions, as, for instance, in Farrell (1987). Here, we allow for only one round of communication followed by the second stage choice of actions.

We make the following assumptions in the model with communication:

- C1: Type L_0 players choose each of the available actions with equal probability (random choice). Having observed the outcome of this randomization, they always send a truthful message (h if they realized action is H and d if the realized action is D). Such players completely ignore the message received by other players when choosing their actions.
- C2: As is standard in level- k models, in stage 2, players of each type choose their best reply, conditional on beliefs about the strategies of lower types, formed after observing the first stage messages. Players do not, however, revise their beliefs about the types of other players even if the message is unexpected.
- C3: When the expected payoff from truthtelling and lying is identical, then subjects tell the truth, i.e., they have a preference for truthtelling.

Proposition 13.31 *The matrix of optimal actions and messages in a level- k model with one round of communications is shown in Table 13.3. The optimal messages (either h or d) for each type are shown immediately following the type designation and the actions can be read off from the individual cells in the matrix.*

Proof: Type L_1 players (first row and first column) believe that the opponents are exclusively type L_0 players. Given assumption C1, type L_0 players are equally likely to choose the action H or D, and their messages truthfully reflect their actions in stage 2. However, when they choose their message, type L_1 players do not know the message chosen by type L_0 , which is equally likely to be h or d. Type L_1 players also know that their message will be completely ignored by type L_0 players, so truthful and untruthful messages give the same payoff. Thus, since they are equally likely to play H or D subsequently, assumption C3 (preference for truthtelling) ensures that they send a truthful message, i.e., message r.

In the second stage, believing that the messages are sent by truthful type L_0 players, type L_1 players best respond by playing H if the message is d and D if the message is h. Thus, when a type L_1 player plays against an opponent who is a type L_1 player, all four outcomes, i.e., HH, HD, DH, DD are equally possible; this is reflected in the (1,1)th cell in Table 13.3. In all other entries in row 1 and column 1 in Table 13.3, type L_1 plays a best response to the observed message of the opponent (H if message is d and D if the message is h).

Type L_2 players believe that all other players are of type L_1 who will randomize equally between the two messages. They also believe (correctly) that type L_1 players consider the messages they receive to be a truthful representation of the play of the opponent in stage 2. Since $a > 1$, type L_2 players always choose message h and action H. Type L_1 players will choose action D after observing the message h. Hence, type L_2 players always achieve their desired payoff of $a > 1$ against type

$L1$ players (but not necessarily against other types of players). This explains the optimal message, h , and the strategy choice of such players: H in the second row and H in the second column in Table 13.3.

Consider now type $L3$ players who believe that all other players are of type $L2$. We know that type $L2$ players always announce the message h and play the strategy H . The best response of type $L3$ players is to announce d (since $C3$ holds) and play action D . Analogously, type $L4$ players truthfully announce h and play the action H . This explains all the elements in Table 13.3. ■

Comparing Propositions 13.30, 13.31 we can draw out the implications of communication.

1. Other than type $L1$, communication alters the actions of all other types. In the absence of communication, type $L1$ has a unique best reply, namely, H . However, in the presence of communication, type $L1$ always best responds to the observed first stage message, which is believed to be sent by a truthful type $L0$ opponent. Type $L2$ effectively plays the role of a Stackelberg leader to the perceived type $L1$ opponent. Thus, type $L2$ sends message h , and plays H ensuring a payoff of $a > 1$. Contrast this with the play of D by type $L2$ in the absence of communication.
2. Consider the pairing in which a type $L1$ player encounters a type $L1$ opponent. In the absence of communication, the outcome is (H, H) , which gives each player a payoff of zero. However, in the presence of communication, all four outcomes are equally possible (see Table 13.3). Hence, at least half the times, these pairings end up with a coordinated outcome, although it is purely by chance.
3. Consider the pairing in which a type $L1$ player encounters a type $L3$ opponent. In the absence of communication, the outcome is (H, H) , and each player gets a payoff of zero. In the presence of communication, the outcome is one of the two coordinated outcomes, i.e., either (D, H) or (H, D) . The secret behind the improved outcome is the listening skills of type $L1$ (who is best responding to a perceived type $L0$ opponent).

Suppose that the proportions of type $L1, L2, L3, L4$ players are, respectively, $\rho_1 = 0.4, \rho_2 = 0.3, \rho_3 = 0.2, \rho_4 = 0.1$; this is consistent with data from similar games. These assumptions imply that the composite frequency of types $L1$ and $L3$ is $\lambda = 0.60$. Thus, using (13.35), the cooperation rate under no communication is $2(.60)(.40) = 0.48$.

Taking account of observations 2, 3, above, we can now compare the overall coordination rates with and without communication. In the pairing $L1, L1$ that takes place with probability $\rho_1^2 = .16$, the coordination rate goes up from 0 to $1/2$, while in the pairing $L1, L3$ or $L3, L1$ (which has probability $2\rho_1\rho_3 = 0.16$) it goes up from zero to one. Hence, the total increase in the probability of coordination is $\frac{1}{2}\rho_1^2 + 2\rho_1\rho_3 = 0.24$, so the new coordination rate is 0.72 . An important aspect of the coordination rates in the level- k model is that they do not depend on how asymmetric the game is, i.e., on the size of the parameter a in the payoff-matrix.

13.4.3 *Risk dominance and payoff dominance in sender-receiver games*

Consider the two-player stag-hunt game shown in Figure 13.13. Two players simultaneously choose actions H and L that have the connotations of “high” and “low” effort levels. Think of action H (hunting stag) as involving high and cooperative effort with a partner, while action L (hunting hare) involves low effort, no cooperation with a partner, but also lower payoffs. Uncoordinated outcomes hurt the player who undertakes action H .

	H	L
H	9, 9	0, 8
L	8, 0	7, 7

Figure 13.13 A stag-hunt game.

There are two pure strategy Nash equilibria in the game: (H, H) and (L, L). The strategy profile (H, H) is the payoff dominant equilibrium, and it is also Pareto optimal, while (L, L) is the risk dominant equilibrium.³⁰ The payoff from choosing H is high for a player if the opponent also chooses H. However, if the opponent chooses L, then choosing H is risky. Hence, in such games there is a basic conflict between the safer strategy, L, which gives lower coordinated payoffs and the riskier strategy, H, which gives higher coordinated payoffs.

In this section we make four assumptions, SH1 to SH4.

SH1: Each player is a level- k player, $L0, L1, L2, \dots$. A player of type Lk believes that all other players are of type $Lk - 1$, $k = 1, 2, \dots$. Type $L0$ chooses uniformly over the two actions in the stag-hunt game, H, L and chooses truthful messages. The actual frequency of type $L0$ players is zero. They exist only in the minds of players of type $L1$ as a model of the type of other players.

THE BENCHMARK CASE OF NO-COMMUNICATION

Consider first the benchmark case of no-communication. Given assumption SH1, type $L1$ players prefer action L over action H because $0.5(8) + 0.5(7) > 0.5(9)$. Type $L2$ players believe that the opponents are of type $L1$, hence, they choose their best reply, L. Proceeding in an analogous manner, all higher types also choose L. Since the frequency of type $L0$ players is zero (assumption SH1), in any pairing of players, the risk dominant equilibrium (L, L) is chosen. This provides firmer foundations for the concept of a risk dominant equilibrium.

Psychologists such as Colman et al. (2014) raise two concerns with the results under no-communication. First, the prediction (L, L) fails to explain why in experimental games there is so much coordination on the outcome (H, H). They cite Colman and Stirk (1998), who find that 79% of players play the strategy H in the stag-hunt game. Second, Colman et al. worry about the asymmetry in the theory because players “never credit their co-players with the same level of strategic sophistication as themselves.” We shall see below that the outcome (H,H) is predicted to arise in level- k models once we allow for the possibility of communication. Further, the empirical evidence does provide good support for level- k models in a range of other games; we shall review some of the evidence below.

THE POSSIBILITY OF ONE-WAY COMMUNICATION

What outcome can we expect if we introduce the possibility of preplay communication in which players can simultaneously send costless messages? We shall only consider one-way communication below.

SH2: There is some preexisting language in which the messages h and l , respectively, are commonly known to mean that the player intends to play the strategies H and L.

³⁰ These concepts are discussed in the appendix on game theory.

One possibility, suggested by Farrell (1988), is that each player can send a message h which is self-committing. However, such a message is not incentive compatible. A player who intends to play L may also have an incentive to send the message, h . Communication can be one-way or two-way. The evidence on the role of communication in enhancing cooperation is mixed. For instance, Cooper et al. (1992) find that two-way communication is conducive for the play of (H, H) , but one-way communication is ineffective. Charness (2000), on the other hand, finds that one-way communication is effective in promoting the play of (H, H) .

Under one-way communication, in the preplay round, a sender chooses a message from the set $\{h, l\}$ that is observed by the receiver. The receiver does not choose to send any message. This is followed by the simultaneous choice of actions from the set $\{H, L\}$.

SH3: When indifferent between two actions, players of all types have a slight preference to tell the truth.

SH4: If any player of type L_k , $k > 1$, receives an unexpected message, then it assumes that the opponent is a truthful player of type L_0 .³¹

Proposition 13.32 *Consider the stag-hunt game in Figure 13.13 with one-way communication and suppose that assumptions SH1–SH4 hold.*

- (a) *The outcome is the payoff dominant equilibrium, (H, H) , in the following two cases.*
 - (i) *When two players of type L_k , $k \geq 2$ play the stag-hunt game.*
 - (ii) *When a type L_1 receiver meets a sender of any type L_k , $k \geq 2$.*
- (b) *When a type L_1 sender plays against a receiver of type L_k , $k \geq 1$ the risk dominant equilibrium (L, L) is played.*

Proof: We need to consider the strategy of each type of player, separately in the role of a sender and a receiver. From assumption SH1, a type L_0 sender randomizes among H , L , and truthfully chooses the message that corresponds to the action intended. Type L_0 players are unaffected by the message of the sender.

Consider a type L_1 sender. He believes the receiver is of type L_0 . Since type L_0 players randomize uniformly over actions, ex-ante type L_1 players have a preference to play L over H . Since type L_0 receivers are unaffected by the message and players have a slight preference to tell the truth, type L_1 senders send the honest message, l , and plays L in the second stage. Type L_1 receivers believe that they are facing an honest type L_0 sender, so responds to the message h by playing H and to the message l by playing L .

Type L_2 senders believe that the receiver is a type L_1 player. Given the strategy of type L_1 receivers, it is optimal for a type L_2 sender to send the message h and play the action H ; this results in the payoff dominant equilibrium (H, H) . Type L_2 receivers expect a type L_1 sender to send the message l and play L . The best reply of a type L_2 receiver is, therefore, L ; this results in the risk dominant equilibrium (L, L) . What happens if a type L_2 receiver receives a counterfactual message h (given that the message l was expected)? Given assumption SH4, the type L_2 receiver believes that the message is sent by a truthful type L_0 player. The best response to such a message is to play H .

Type L_3 senders will send message h and play H because they believe that when type L_2 receivers observe this counterfactual message, they will respond with H . Type L_3 receivers know that type L_2 senders will send the message h and play the action H to which their best reply is H .

³¹ This assumption may seem strong, but it is consistent with the “non-equilibrium in beliefs” feature of level- k models. In any case, the reasonableness of this or similar assumption is eventually an empirical matter. For the moment, this is the most pedagogically convenient assumption.

Suppose that a type $L3$ receiver receives the unexpected message l . Given assumption SH4, the receiver then believes that it faces a truthful type $L0$ sender. Thus, the best response to this message is L . All higher types in their roles as senders and receivers will likewise play just as type $L3$ players do.

Hence, in any pairwise matching of types $L2$ or higher, the equilibrium play is the payoff dominant equilibrium (H, H) . ■

A SUMMARY

Coordination, when it occurs in a level- k model, is entirely accidental. When there are multiple equilibria, traditional concepts such as payoff dominance and risk dominance are not useful to choose among the equilibria in level- k models. The relative frequency of the various types determines the fraction of times that a particular equilibrium is played. Some other alternatives such as global games also do not rely on the traditional selection principles such as payoff dominance and risk dominance. In global games (see Morris and Shin, 1998), players observe privately some perturbations to their payoffs. The perturbed game is dominance solvable and leads to a unique outcome. In many models, such as models of bank runs, where global games choose the risk dominant outcome, so too do level- k models. However level- k models have better behavioral foundations and in more complex games, the two models have different predictions; see Crawford et al. (2013, pp. 40–41).³²

13.4.4 *The market entry game*

Consider the notion of *k-rationalizability* in Bernheim (1984) and Pearce (1984). The definition of *k-rationalizability* can be given iteratively. A strategy for a player is 1-rationalizable if it is a best response to some strategy profile of the opponents. A strategy is 2-rationalizable if it is a best response to a profile of strategies of the other players that are 1-rationalizable. For any k , a *k-rationalizable* strategy of a player is a best response to $k - 1$ -rationalizable strategies of the opponents. In two-player games, a *k-rationalizable* strategy survives k rounds of iterated deletion of strictly dominated strategies. A Nash equilibrium is *k-rationalizable* but the converse is not true. The cognitive requirements on the part of players to play *k-rationalizable* strategies are much more modest as compared to a Nash equilibrium, hence, such strategies are behaviorally more appealing.³³ Level- k models respect *k-rationalizability*.

Suppose that we restrict ourselves to dominance solvable two-player games in which iterated deletion of dominated strategies gives rise to a unique equilibrium strategy profile. We give a heuristic proof that level- k players will play such an equilibrium strategy profile. Suppose that type $L0$ players randomize over the set of pure strategies. Type $L1$ players believe that the opponent is of type $L0$ and each strategy of the opponent has a positive probability of being played. Since type $L1$ players always plays a best reply they will never play a dominated strategy (1-rationalizable strategy). Now consider a type $L2$ player who imagines the opponent to be of type $L1$ who has deleted his dominated strategies. Since type $L2$ also always plays a best reply, he too will never play any of his dominated strategies (2-rationalizable strategy). Proceeding in this manner, if there are k types, then the outcome is *k-rationalizable*. For high enough k , one obtains the unique equilibrium strategy profile.

³² A key implication of global games is that multiple equilibria arise on account of the common knowledge of payoff functions. However, Heinemann et al. (2004) find only very small differences in the behavior of their subjects in experiments for the private information and the common information conditions.

³³ For further discussion of these issues, see Crawford et al. (2013).

In level- k models, each type Lk , $k \geq 1$ acts in a manner similar to a Stackelberg leader to type $Lk - 1$. An important application of this idea is to the *market entry game* originally due to Selten and Güth (1982). Suppose that there is a market with a level of demand equal to $d > 0$ integer units. There are N potential entrants, each capable of supplying 1 unit. The entrants decide to enter simultaneously. If d or fewer entrants enter, so supply is no greater than demand, then each entrant earns a payoff of 1. If more than d entrants enter, then each entrant gets a payoff of zero. Firms that decide not to enter receive a certain payoff of 0.5. The payoffs for this game are typically set up as follows; see Sundali et al. (1995). Let π_i be the payoff to the i th entrant. Let $k \geq 0$ be the outside option and let $n \leq N$ be the number of entrants. Then

$$\pi_i = \begin{cases} k & \text{if firm stays out} \\ k + r(d - n) & \text{if firm enters and } n < d, \\ k & \text{if firm enters and } n \geq d \end{cases}$$

where $k \geq 0, r > 0$ are the two exogenous parameters.

The efficient outcome is for exactly d firms to enter. It is easily checked that this is a Nash equilibrium.³⁴ Can firms, who choose their entry decisions simultaneously, achieve an efficient outcome? Kahneman (1988) ran early experiments with the market entry game. Despite being unable to communicate or signal, close to the efficient number of firms make the entry decision. This finding amazed Kahneman and he commented that to a psychologist “it looks like magic.” This “magical” finding was replicated by other researchers, in particular, after the subjects were allowed a few rounds of learning. The result also survived when several aspects of the underlying game were varied.³⁵

The empirical results are consistent with a Nash equilibrium. However, a CH model applied to the market entry game also makes the same prediction. Suppose that with $\tau = 1.7$, the respective fractions of the players of levels, $L0, L1, L2$ is as given in (13.30), i.e., $f(0) = 0.18, f(1) = 0.31, f(2) = 0.26$, and so on. Players of type $L0$ randomize uniformly between entering and not entering. Players of type $L1$ correctly anticipate that half of the level $L0$ players will enter and they believe that all players are level $L0$. So if $0.5N < d$ they enter, otherwise they stay out. Level $L2$ players believe that the fractions of type $L0$ and $L1$ players is given by (13.31). Hence, they will enter if $N \left(\frac{g_2(0)}{2} + g_2(1) \right) < d$ otherwise they stay out. This analysis can be carried out until it is no longer profitable for a player of some level, Lk , to enter. For a large enough value of k , these responses smooth out, so the efficient outcome is achieved.

Crawford et al. (2013) suggest that the BOS game shown in Figure 13.12 can be considered to be a type of market entry game with suitable reinterpretation; see Figure 13.14. Figure 13.14 represents a two-player market entry game that has capacity $d = 1$. If both players enter they get a payoff of zero. If both players stay out, they get their reservation payoff of zero. If one player enters and the other does not, then the entrant gets a payoff of $a > 1$.³⁶ What entry decision should both firms make, if they choose their decisions simultaneously? The answer is given above in Section

³⁴ If less than d firms enter then it is a dominant strategy for an entrant to enter and if d or more firms enter then it is a dominant strategy to stay out.

³⁵ See, for instance, the findings of Sundali et al. (1995) who varied the size of the market, d . Further research varied the size of the outside option (Rapoport et al., 1998), allowed for asymmetric costs of entry (Rapoport et al., 2002), risky outside options (Rapoport et al., 2002), and asymmetric information about payoffs (Rapoport and Erev, 1998), among others. However, allowing for some learning, the magical result survives.

³⁶ There is a slight difficulty, though, of interpreting the payoffs from the strategy profile (Out, In); the player who stays out gets 1. Typically one expects that the outside option is unaffected by the decisions of the rival.

	In	Out
In	0, 0	a, 1
Out	1, a	0, 0

Figure 13.14 A two-player market entry game ($a > 1$).

Source: Crawford et al. (2013) with permission from the American Economic Association.

13.4.2 under various assumptions on the absence or presence of communication between the players.

Goldfarb and Yang (2009) and Goldfarb and Xiao (2011) conduct field studies of the market entry game. In Goldfarb and Yang (2009) managers must decide on technology adoption (which is the entry decision).³⁷ Level-0 managers assume that they are monopolists and maximize profits. Each of the higher types Lk , $k \geq 1$ assumes that all opponents are of lower types. The types are estimated assuming that they come from the Poisson CH model. The estimated value of τ is 2.67 (by contrast the estimated value of τ for the experimental games in Camerer et al. (2004a) typically lies between 1 and 2). There is heterogeneity in the estimated levels of managers. Greater rationality (higher k in the CH model) is associated with more education and larger cities; more rational managers also survive in the market longer. Following the US Telecommunications Act of 1996, Goldfarb and Xiao (2011) find that the CH model performs better than a noisy Nash model in 1998 as compared to 2002 when both models perform almost equally well. This is consistent with the agenda of level- k and CH models to fit better the evidence for new games and in the early rounds of play.

13.4.5 Empirical evidence on level- k models

Costa-Gomes et al. (2001) tested for the presence of level- k players with types $L0, L1, L2, \dots$ using MOUSELAB experiments that are described in the previous chapter. Costa-Gomes et al. (2001) consider four non-strategic types and five strategic types. *Altruists* maximize the sum of their and the opponent's payoffs. *Pessimists* make maximin decisions that maximize their payoffs over the minimum payoffs over all the choices of the opponents. *Optimists* make maximax decisions, i.e., maximize their payoffs over the maximum payoffs that arise by varying the opponents' actions. *Naive* players are $L1$ players who assume that the opponent chooses uniformly randomly over all possible actions. In addition to type $L2$, the four other strategic types are as follows.

1. The "Equilibrium type" is a player who satisfies the assumptions of traditional game theory in terms of cognition and rationality.
2. Type D1 (or Dominance 1) players engage in one round of iterated deletion of pure strategies and then play a best response assuming that the opponent will play each of the remaining strategies using a uniform probability distribution.
3. Type D2 (or Dominance 2) players engage in two rounds of iterated deletion of dominated strategies and then best respond by assuming a uniform probability distribution over the opponents' remaining strategies.

³⁷ There are four possible choices that relate to the adoption of two possible technologies, A and B. (1) Adopt no technology. (2) Adopt technology A only. (3) Adopt technology B only. (4) Adopt both A and B.

4. The “Sophisticated type” is a player who best responds to the (possibly) mixed strategies of other players, correctly forecasting the frequencies of the types of other players.

Costa-Gomes et al. (2001) build on the MOUSELAB experiments of Johnson et al. (2002) by exposing subjects in experiments to 18 normal form games that allow for various levels of iterated dominance. The strategy spaces in each game are fairly low (the games are either 2×2 , or 2×3 , or 2×4) and the experimental design minimizes learning effects and repeated game effects. Each game has a unique pure strategy equilibrium. The MOUSELAB interface for one of the games is shown in Figure 13.15. In order to minimize framing effects, (1) each player is treated as a row player and referred to as “You,” while the opponent is described as “S/He,” and (2) strategies of the row player are labeled as #,* while those of the column player are labeled as &, @. There are eight boxes corresponding to the payoffs (or points) for each player. The boxes in the first two columns show the payoffs of the row player, while the boxes in the last two columns show the payoffs of the column player. Moving the cursor to a box immediately reveals the contents. For instance, in Figure 13.15, the subject has moved the cursor to a box that reveals an own payoff of 42 when he plays # and the opponent plays @.

The design advantage of Costa-Gomes et al. (2001) over Johnson et al. (2002) is that individual searches can vary in three important dimensions.³⁸ Up–down searches in own payoffs and left–right searches in the opponent’s payoffs are suggestive of game theoretic reasoning. However, frequent transitions from own payoffs to the opponent’s payoffs are suggestive of social prefer-

	S/He: &	s/He: @ :	s/He &	s/He: @ :
You: #		42 +		
You: *				

	YOUR POINTS	HER/HIS POINTS
Choose one	You: #	You: *

Figure 13.15 MOUSELAB interface in one of the games in Costa-Gomes et al. (2001). With permission of The Econometric Society.

³⁸ The work of Johnson et al. (2002) is described in the previous chapter.

ences. A baseline treatment was followed by a trained subjects treatment in which subjects were trained in identifying equilibrium outcomes. The findings are as follows:

1. There are systematic departures from equilibrium. Equilibrium strategies are more likely to be played when the game can be solved by one or two rounds of iterated dominance relative to cases where three or more rounds of iterated dominance are required.³⁹ Quantal response models would find it difficult to explain these findings.
2. Focusing on decisions alone (i.e., ignoring search and lookup data for the moment), there are high estimated frequencies of types *L2*, Naive (which is type *L1* here), Optimistic, and *D1*. The strategic types make up more than 70% of the types but the presence of Equilibrium and Sophisticated types is low. Type *D1* makes up about 20% of the sample.
3. The information searches of subjects are even more heterogeneous than their decisions (even when errors in decisions are taken into account). When search data is taken into account, along with data on actual decisions, the econometric results suggest that the most important types are Naive (i.e., *L1*) and *L2*; these make up about 67–89% of the population. Type *D1* players account for 0–20% of the subjects.

Brocas et al., (2014) consider two-player betting games that have three states and two-sided private information. Players observe one or two of the three states and must decide whether to bet or not. If at least one player chooses not to bet, then they each get a known payoff. If both bet, then they get a payoff that depends on the realized state. The payoffs are hidden in opaque boxes and players must click the boxes to reveal the payoffs. Following each round of the game, subjects are given feedback, so learning opportunities are present. Two different game situations, based on the degree of cognitive difficulty, are used—a simple situation and a more complex situation. The incidence of Nash equilibrium play in the simple and the complex games is, respectively, about 50% and 25%.

These results are in contrast to those from bargaining games in Section 12.5, where moving to the most cognitively challenging bargaining situation, bargaining under two-sided asymmetric information, improved the proportion of Nash choices. The pattern of lookups, using the MOUSELAB software, turns out empirically to be a good predictor about whether subjects will play Nash. There is substantial heterogeneity in choices that enables the separation of subject choices into four clusters. Three of these clusters enable the detection of play that is consistent with level-1, 2, and 3 players. A majority of subjects in these clusters do play in a manner that is consistent with the predictions of level-*k* models.

Costa-Gomes and Crawford (2006) consider 16 dominance solvable two-player normal form games of simultaneous moves called guessing games. The aim is to separate the various types of players more cleanly. Denote the two players by *i, j*, their respective guesses by x_i, x_j , their respective targets by $t_i > 0, t_j > 0$, and the respective domains of real numbers over which choices are made by $[a_i, b_i], [a_j, b_j]$. If a player makes a choice above the upper limit, then his choice is adjusted downwards to the upper limit. Conversely, if a player makes a choice below the lower limit, his choice is adjusted upwards to the lower limit. Denote the respective adjusted guesses of the players by $\hat{x}_i, \hat{x}_j$. The players make their guesses simultaneously following which the payoffs are realized.

Let $e_i = |\hat{x}_i - t_i \hat{x}_j|$ be the absolute value of how close player *i*'s guess is to the target times player *j*'s guess. The payoff of player *i* increases as $e_i \rightarrow 0$ and is highest at $e_i = 0$. Analogously,

³⁹ In games which can be solved by just two rounds of iterated deletion of strategies, the types *L2*, *D1*, and Equilibrium behave in an identical manner and so types can be confounded. This problem is addressed by considering a range of games with different structures.

the payoff of player j increases as $e_j = |\hat{x}_j - t_j \hat{x}_i|$ becomes closer to zero and is highest at $e_j = 0$. The games are dominance solvable in 3–52 rounds. The targets and limits of players vary independently from one game to the other.⁴⁰ In all games that are considered, $t_i t_j \neq 1$. The next proposition shows that in the unique equilibrium, if $t_i t_j < 1$, then the guesses of the players are determined by their lower limits a_i, a_j , while if $t_i t_j > 1$, then they are determined by their upper limit b_i, b_j .

Proposition 13.33

(i) Let $t_i t_j < 1$. Then in the unique equilibrium, the guesses of the two players are determined as follows.

$$\hat{x}_i = \begin{cases} a_i & \text{if } t_i a_j \leq a_i \\ \min \{t_i a_j, b_i\} & \text{if } t_i a_j > a_i \end{cases}; \hat{x}_j = \begin{cases} a_j & \text{if } t_j a_i \leq a_j \\ \min \{t_j a_i, b_j\} & \text{if } t_j a_i > a_j \end{cases}. \quad (13.36)$$

(ii) Let $t_i t_j > 1$. Then in the unique equilibrium, the guesses of the two players are determined as follows.

$$\hat{x}_i = \begin{cases} b_i & \text{if } t_i b_j \geq b_i \\ \max \{t_i b_j, a_i\} & \text{if } t_i b_j < b_i \end{cases}; \hat{x}_j = \begin{cases} b_j & \text{if } t_j b_i \geq b_j \\ \max \{t_j b_i, a_j\} & \text{if } t_j b_i < b_j \end{cases}.$$

Proof: (a) The proof involves checking that the various strategy profiles are a Nash equilibrium under the specified conditions. For part (i) we need to consider the following four cases. (1) $t_i a_j \leq a_i$ and $t_j a_i \leq a_j$. (2) $t_i a_j \leq a_i$ and $t_j a_i > a_j$. (3) $t_j a_i \leq a_j$ and $t_i a_j > a_i$. (4) $t_i a_j > a_i$ and $t_j a_i > a_j$. We only prove cases (1) and (2); the other two cases are similar, so they are omitted.

Consider case (1), $t_i a_j \leq a_i$ and $t_j a_i \leq a_j$. In this case, from (13.36), the claimed Nash equilibrium is $(\hat{x}_i, \hat{x}_j) = (a_i, a_j)$. To see that this is the case, conditional on $\hat{x}_j = a_j$, we have $e_i = |\hat{x}_i - t_i a_j|$. But $t_i a_j \leq a_i$, hence, the smallest possible value of e_i (and the highest payoff for player i) occurs when $\hat{x}_i = a_i$. Similarly, conditional on $\hat{x}_i = a_i$ we have that $e_j = |\hat{x}_j - t_j a_i|$ and $t_j a_i \leq a_j$. Thus, player j achieves his highest payoff by guessing $\hat{x}_j = a_j$.

Now consider case (2), $t_i a_j \leq a_i$ and $t_j a_i > a_j$. The claimed Nash equilibrium in this case is $(\hat{x}_i, \hat{x}_j) = (a_i, \min \{t_j a_i, b_j\})$. Conditional on the guess $\hat{x}_j = \min \{t_j a_i, b_j\}$ we have $e_i = |\hat{x}_i - t_i \min \{t_j a_i, b_j\}|$. First suppose that $\min \{t_j a_i, b_j\} = t_j a_i$, so $e_i = |\hat{x}_i - t_i t_j a_i|$. Since $t_i t_j < 1$ it follows that the minimum possible value of e_i (and highest payoff of player i) arises when $\hat{x}_i = a_i$, as claimed. If $\min \{t_j a_i, b_j\} = b_j$, then $b_j \leq t_j a_i$, hence, the lowest guess that player i can choose that is closest to $t_i \hat{x}_j$ is still $\hat{x}_i = a_i$. Now let us check that the guess of player j , $\hat{x}_j = \min \{t_j a_i, b_j\}$ is optimal given the guess of player i , $\hat{x}_i = a_i$. In this case, $e_j = |\hat{x}_j - t_j a_i|$ but $t_j a_i > a_j$. If $t_j a_i > b_j$, then the closest guess of player j to the guess of player i is $\hat{x}_j = b_j$. If on the other hand, $a_j < t_j a_i < b_j$, then the closest guess of player j is $\hat{x}_j = t_j a_i$. This completes the proof of part (i).

The proof of part (ii) is analogous to part (i), hence, it is omitted. ■

Subjects who make guesses consistent with Proposition 13.33 are known as Equilibrium types. Sophisticated types maximize their payoffs with the correct distribution of the other players in the sample.

Now consider the guesses of level- k players. Type $L0$ players make a uniform random guess over their intervals $[a, b]$. Type $L1$ players assume that all other players are of type $L0$.

⁴⁰ In the actual experiments, possible targets considered are 0.5, 0.7, 1.3, and 1.5 while the possible limits considered are [100, 500], [100, 900], [300, 500], and [300, 900].

Proposition 13.34 (Costa-Gomes and Crawford, 2006): Suppose the point payoff function of player i is $e_i = g(x_i - t_i x_j)$ where g is symmetric, continuous, almost everywhere differentiable function that is weakly decreasing in $|x_i - t_i x_j|$, where $\hat{x}_i$ is a random variable uniformly distributed on $[a, b]$. Then, for any player with a continuous, almost everywhere differentiable von Neumann–Morgenstern utility function u , his expected-utility maximizing choice is $x_i^* = t_i \left(\frac{a+b}{2} \right)$. If $t_i \left(\frac{a+b}{2} \right) > b$, then $\hat{x}_i = b$ and if $t_i \left(\frac{a+b}{2} \right) < a$, then $\hat{x}_i = a$.

Given the result in Proposition 13.34, the remaining guesses of higher-level players are straightforward. Player i of type $L2$ believes that others are of type $L1$, hence, $e_i = \left| \hat{x}_i - t_i \frac{a+b}{2} \right|$, which is minimized by a guess of $\hat{x}_i = t_i \frac{a+b}{2}$. If this guess falls outside some interval $[c, d]$ over which the player is allowed to guess then the guess is adjusted to the lower or the upper interval, whichever is nearest. Analogously, player i of type $L3$ makes the guess $\hat{x}_i = t_i^2 \frac{a+b}{2}$ and so on.

There were a total of 88 subjects in the main treatment. The guesses of 43 subjects conformed (within a very small margin of error) with the guesses of level- k types in 10 out of the 16 games. Of these the guesses of 12 subjects matched $L2$ very closely; 72% of the guesses were exact, which is difficult to ascribe to chance because there were 200–800 possible exact guesses in each game. The experimental design builds in risk neutrality and self-interested behavior, hence, the presence of level- k types is not the result of Equilibrium types making guesses that are confounded by issues of risk aversion or social preferences.

Based on the decisions made by subjects, the econometric analysis shows up the presence of only types $L1$, $L2$, $L3$, and Equilibrium. The breakdown was: 20 $L1$ types, 12 $L2$ types, 3 $L3$ types, and 8 Equilibrium types. The data reveals the complete absence of the Dominance types (e.g., types $D1$, $D2$ in the study of Costa-Gomes et al., 2001) as well as the Sophisticated types. The data is also hard to reconcile with QRE.

A MOUSELAB interface was used to elicit the search and lookup patterns of subjects. Subjects were presented with a six box display. Among the boxes, they could search for their own and their partner's limits and targets and then enter a guess. The maintained hypothesis is that each subject belongs to one of the types: Level- k , Dominance, or Equilibrium. The search pattern was then matched to types. Consider, for instance, a player i of type $L1$ whose ideal guess is $t_i (a + b) / 2$. Hence, the lookups must involve elements of the set $\{[a_j, b_j], t_i\}$ in any order. We leave it to the reader to construct the predicted lookup patterns of types $L2$, $L3$, $D1$, $D2$, Equilibrium, and Sophisticated.⁴¹

The searches of types $L1$, $L2$, and Equilibrium, were very similar in the baseline and the robot conditions. In the robot condition, subjects play against a robot in order to switch off their social preferences. The robots are preprogrammed to play a particular decision rule. Unlike the results in Johnson et al. (2002), explained in the previous chapter, the robot treatment does not nudge behavior towards equilibrium. In conjunction, the search and decision data confirm that a significant fraction of players are of types $L1$, $L2$, Equilibrium, or apparent combinations of Equilibrium and $L3$ types. Irrationality or social preferences can be ruled out as factors contributing to the departure from equilibrium behavior of subjects. In several cases, when the type of the subject is not evident from decisions (or guesses) alone, the search data on lookups is able to clarify the type of the subject.⁴²

⁴¹ The answers can be found in Table 4, page 1751 of Costa-Gomes and Crawford (2006).

⁴² For instance, the guesses of subject "309" seem to identify the subject as clearly type $L2$. However, when the search data was examined, it was entirely inconsistent with the pattern of lookups for a type $L2$.

Georganas et al. (2012) consider the stability of level- k types across two kinds of two-player games—undercutting games and guessing games. The guessing game is as in Costa-Gomes and Crawford (2006). The undercutting game is characterized by two integers p, q . Two players choose integers from the identical strategy sets $s_i = \{1, 2, \dots, p, \dots, q\}$, $i = 1, 2$. Any player $j = 1, 2$ pays player $i \neq j$ an amount \$10 if $s_i = p < s_j$ or if $1 + s_i = s_j \leq p$. If player i expects that the opponent will choose an integer greater than p , then the best response is $s_i = p$. If the expectation is that the opponent will choose an integer less than or equal to p , then the best response is $s_i = s_j - 1$, i.e., undercut by 1 the guess of the other player. If both players choose the number 1, then both earn a payoff of one. If both choose the highest number q , then they both lose 11 (this rules out mixed strategy Nash equilibria). If one player chooses 1 and the other chooses q , then the player choosing 1 loses 11. In all other cases, players earn a payoff of zero. In the unique Nash equilibrium, both players choose the number 1. The subjects in experiments are shown payoff tables corresponding to this game. The players play four undercutting games, UG1-UG4, and six guessing games, GG5-GG10.

For pooled data in the undercutting games, the percentage of types $L0$, $L1$, $L2$, $L3$, and Equilibrium (those who play the Nash equilibrium) are, respectively, 4.31, 28.45, 26.72, 5.17, 35.34. The corresponding percentages for the pooled data from the guessing games are 1.72, 50, 10.34, 10.34, and 27.59. These findings are somewhat similar to those in Costa-Gomes and Crawford (2006) except that the proportion of Equilibrium players in the undercutting games is significantly higher. In the four undercutting games, the distribution of types is similar. However, the distribution of types in the six guessing games is quite variable. For instance, the percentage of $L0$ players varies from 0.68% in GG6 to 43.10% in GG7; the percentage of $L1$ players goes up from 37.07% in GG7 to 70.69% in GG5. There is similar variability in the proportion that play Nash.

Another difference is that while Costa-Gomes and Crawford (2006) find that 48.9% of the choices are exact matches with the predictions of level- k models, this percentage is only 14.22 in the case of Georganas et al. (2012). Thus, one needs to fit the behavior of subjects by assuming that they play noisy best responses. The authors use five psychometric measures: An IQ quiz, an eye gaze test, memory quiz, the Cognitive Reflection test (CRT), and a one-player Takeover Game. But correlations in scores obtained by players in the psychometric measures are low and these measures are not able to predict the types of players in the level- k model.

The main argument in Georganas et al. (2012) is that the variability of types in the guessing game and stability in the undercutting games suggests that the level- k model may be a heuristic that is triggered in some games and not in others. They are unable to identify what activates the trigger to use a level- k model, identifying a possible limitation in the use of the level- k model. What accounts for the difference in their results from Costa-Gomes and Crawford (2006)? They argue that the source is the difference in instructions. The Costa-Gomes and Crawford (2006) instructions are extremely detailed and prior to the experiments, subjects were required to calculate their best responses to hypothetical questions and to hypothetical strategies of the opponent. This may have primed them to play an exact best response in the guessing games. By contrast, the instructions in Georganas et al. (2012) are minimal and subjects are not primed to play any best responses.

13.4.6 *Do non-strategic players exist?*

The estimates from Costa-Gomes et al. (2001) and other studies such as Stahl and Wilson (1994, 1995) and Costa-Gomes et al. (2006) show that type $L0$ players exist mainly as a model

that players of type $L1$ have about their opponents. However, several recent papers also make an intriguing case for the presence of non-strategic players. We now consider some of these papers.

Costa-Gomes and Weizsäcker (2008) consider 14 normal form two-player games. They record the actions of players in these games and also elicit their beliefs about other players. The beliefs are elicited using incentive compatible mechanisms. The behavior of subjects is largely inconsistent in the sense that the actions taken by the players are not consistent with their stated beliefs. One possible resolution to the problem is that players do not optimize. In order to account for this problem, a noisy best response model is used, however, the inconsistency remains. A second possibility is that players are not able to optimize due to cognitive limitations. To test this possibility, players were provided with calculators, yet the inconsistency remained. Eight behavioral models are used to fit the data. It appears that the play of subjects is consistent with level $L1$ players who expect their opponents to randomize over their actions. However, the elicited beliefs reveal that they expect their opponents to be of type $L1$ as well. This is problematic because then they should optimally be choosing strategies that are consistent with type $L2$ players. The authors rule out possible explanations of their results based on risk aversion and social preferences. In a similar vein, Danz et al. (2012) find that for repeated games, when no feedback is given, there is very little consistency between beliefs and actions. With feedback, the consistency improves but is not perfect.

Agranov et al. (2012) consider the presence of non-strategic players, defined as those who choose a dominated action in a p -beauty contest game in which $p = \frac{2}{3}$. They also check for the portability of strategic thinking by then asking subjects to play a Monty Hall problem. The authors implement a choice process from Caplin et al. (2011) in which subjects have a chance to reflect on their choices for a fixed length of time. The pattern of their choices is then used to classify them into alternative types.

Subjects in experiments were told that the p -beauty contest with $p = \frac{2}{3}$ was played with eight undergraduate students. If the subject can come up with a choice that is better (closer to $\frac{2}{3}$ of the average of the choice of the eight students), then they win a prize of \$10. Importantly, the choice of the subjects in experiments does not alter the original choices of the eight students. Two sub-treatments are run in which subjects are given, respectively, 180 and 30 seconds to make a choice.

In the *strategic choice process treatment* (SCP), the two sub-treatments of 180 and 30 seconds were run but with one important difference. Subjects could make as many choices as they wanted to during the respective time intervals. One of the seconds in each case is then chosen at random and the prevailing choice of the subject at that second is picked to reward the subject. In the $\frac{2}{3}$ -beauty contest, where subjects pick a number from $[0, 100]$, picking a number greater than 67 is a dominated strategy. *Naive players* are defined as those who choose a number greater than 67 at least once after 30 seconds or more in the experiments. Players who are not naive are classified as *strategic players*.

In the SCP experiments, 43.3% of the players were classified as naive players. If one used simply the final choice of the subjects, then 5% of the subjects are classified as naive. The authors argue that using final choices, thus, underestimates the proportion of naive players. The average choice of the naive players was 50 using either the final choices or the intermediate choice data. This is consistent with the assumed choice of level $L0$ players in level- k models. Over the relevant time period that subjects make choices in the SCP experiments, naive players change their choices more frequently, in a statistically significant sense, as compared to strategic players. Strategic players also converge much faster to their eventual choices. Another indicator of strategic reasoning is to see if players alter their guesses, over time, in steps of $\frac{2}{3}$ times their previous

choices (within some reasonable margin, chosen to be 5% by the authors). Strategic players make a larger number of such strategic choices as compared to naive players.

The portability of these results is tested using the Monty Hall game. In this game, subjects are shown three closed doors and told that there is a \$5 reward behind one of the three doors and nothing behind the other two doors. Each subject is then invited to choose one closed door. Having chosen the door, the experimenter reveals to the subject that he knows which of the doors hides the \$5 prize. The experimenter then opens one of the other doors that does not have the prize. Conditional on the move of the experimenter, should the subject keep to his original choice of the door or choose the other door? It can be shown that the subject can increase the probability of winning by altering his choice.⁴³ A probit regression is used in which the dependent variable equals 1 if the subject altered his choice and 0 otherwise. The main finding is that naive subjects are less likely to switch doors, while among strategic subjects, those who end with lower final choices (relative to intermediate choices) are more likely to switch doors. Thus, there seems to be an element of portability of strategic behavior across games in these experiments.

De Sousa et al. (2013) consider the possibility of non-strategic players among a pool of 270 chess players. In phase 1, the subjects played ten p-beauty contest games with $p = \frac{2}{3}$ and $p = \frac{4}{3}$; the guesses were chosen from the interval $[0, 100]$. Subjects played these games against other subjects classified into types A, B, C, or D based on their Elo chess rankings. A random type was also introduced; such a type chose numbers at random. Five games each were played with values of $p = \frac{2}{3}$ and $\frac{4}{3}$ against each of these five possible opponents.

In phase 2, subjects played a new game, the 11–20 game against another chess player of unspecified Elo ranking. In this game, due to Arad and Rubinstein (2012), two players must simultaneously request a number between 11 and 20 shekels. Both players receive their requests. In addition, if a player requests exactly one shekel less than the other player, he receives an extra 20 shekels. There is a unique, symmetric, mixed strategy Nash equilibrium in which players choose numbers 11–14 with probability zero, numbers 15 and 16 with probability $\frac{1}{4}$ and numbers 17, 18, 19, and 20 with respective probabilities $\frac{4}{20}$, $\frac{3}{20}$, $\frac{2}{20}$, $\frac{1}{20}$. In phase 3, subjects were given the chance to participate in a p-beauty contest game with $p = \frac{2}{3}$ against all other participants.

In the p-beauty contest with guesses allowed in the interval $[0, 100]$ the Nash prediction with $p = \frac{2}{3}$ is a guess of zero, while with $p = \frac{4}{3}$ it is 100. The authors define non-strategic players as those who choose a higher guess under $p = \frac{2}{3}$ as compared to under $p = \frac{4}{3}$. The number of times that players engage in such non-strategic choices when playing against the same opponent allows the authors to construct a pairwise rationalizable actions index (PRAI) that takes integer values between 0 and 5. A value of 0 for the index indicates that players consistently make non-strategic choices (i.e., relatively lower numbers chosen under $p = \frac{4}{3}$ as compared to under $p = \frac{2}{3}$), while a value of 5 indicates that they never make such choices. Intermediate values of the index, 2, 3, and 4 indicate decreasing non-strategic choices. Individuals with a PRAI index of 0–3 are classified as non-strategic and those with 4–5 are classified as strategic.

The non-strategic players do not react to changes in p from $\frac{2}{3}$ to $\frac{4}{3}$ and their mean play is close to a choice of 50.⁴⁴ On the other hand, strategic players show significant changes in guesses as p changes. Their average guess is 36.12 when $p = \frac{2}{3}$ and 70.09 when $p = \frac{4}{3}$. In each case, their

⁴³ Ex-ante, the probability that the chosen door hides the \$5 prize is $\frac{1}{3}$ and the probability that the other two doors contain the hidden prize is $\frac{2}{3}$. Conditional on the experimenter opening one of the doors and showing the absence of the prize, the remaining door now has a $\frac{2}{3}$ probability of containing the prize. This is greater than the $\frac{1}{3}$ probability associated with the player's chosen door.

⁴⁴ The average choice is 48.98 when $p = \frac{2}{3}$ and 51.25 when $p = \frac{4}{3}$.

behavior is roughly consistent with level $L1$ players who best respond to beliefs that other players are level $L0$ players who choose uniformly randomly. Strategic players earn more in the game relative to non-strategic players. Non-strategic players do not spend any less time contemplating about their guesses as compared to strategic players. Thus, the difference in guesses probably lies in the different manner in which they process the relevant information. Interestingly, when p changes, non-strategic players react in terms of the time spent on making their guesses, however, their mean guesses do not change. Changing the stakes has no effect on these results. Higher Elo rankings of chess players are not significantly correlated with their position on the PRAI index.

In phase 2, the authors added four comprehension questions to the 11–20 game. Non-strategic players made more mistakes in answering these questions relative to strategic players. Players typically fail to play the equilibrium strategies in the 11–20 game (as in Arad and Rubinstein, 2012). However, players who are ranked higher on the PRAI index in the p-beauty contest game make choices that are closer to equilibrium suggesting some portability of results.

Burchardi and Penczynski (2014) consider a p-beauty contest played between two teams of two players each. There are two stages. In the first stage, each player conveys a message to his team member along with a written justification that has no time limits or size limits. The number suggested in this stage is termed the *suggested decision*. This stage allows players to influence the team member's decision. In the second stage, players simultaneously make choices for the team—such choices are known as *final decisions*. The final decision of each player is chosen with probability one half and termed as the team *decision*; this is the basis on which the team's reward is decided. The suggested decision reflects an individual's reasoning and corresponds to choices in the existing p-beauty contest experiments. This is the basis for the statistical analysis. The final decision, by contrast, does not reflect individual decision making.

A total of 84 subjects played three rounds of a p-beauty contest with $p = \frac{2}{3}$. Unlike many other studies, students do not play any test rounds that involve strategic considerations. This ensures that students are not already primed to respond in a particular manner. The messages sent by the players in the first stage are then analyzed to classify them into various types of players. Since the messages are written in free-form text, two trained research assistants are employed to classify the messages into categories such as level- k players, Equilibrium players, and Sophisticated players.⁴⁵ Within level- k players, the research assistants are asked to state the lower and upper bounds on levels. This classification is rechecked by three more research assistants. Only when there is sufficient consensus among the research assistants on the response of a subject is the data on the subject included into the statistical tests.

The results are as follows. The mean of the suggested decisions in the first stage is 43.93 that is slightly higher than Nagel's (1995) figure of 36.73. For 50 out of the 84 subjects for which the lower and upper bounds coincide, there were 17 players of level $L0$, 26 players of level $L1$, 6 players of level $L2$, and 1 player of level $L3$. For 20 out of the remaining subjects for which the upper and lower bounds differ, the bounds are nevertheless between players of level $L0$ and $L2$. The behavior of one player each is consistent with the Equilibrium and the Sophisticated types. While 20% of the players are unambiguously of type $L0$, the possibility that a further 24 out of the 84 players are of type $L0$ cannot be excluded. In taking their final decisions, lower level- k players are influenced by the first stage messages sent by higher level- k players.⁴⁶

⁴⁵ Equilibrium players play in conformity with a Nash equilibrium and suggest a choice of zero, while Sophisticated players may figure out the equilibrium choice but prefer to play a non-equilibrium action, possibly because they believe that others are less than fully rational in the Nash sense.

⁴⁶ This particular result appears in a companion paper by Penczynski (2012).

The hypothesis that the level $L0$ players make uniform random choices is rejected using a Kolgomorov–Smirnov test. A structural level- k model fitted to the data reveals the following proportions of types. 37% are type $L0$, 47% are type $L1$, 15% are type $L2$, and 1% are type $L3$. The authors hypothesize that the relatively larger percentage of type $L0$ players in their sample, relative to many other studies (e.g., Camerer et al., 2004a) is due to the different methodologies employed by the studies and the nature of the pre-experimental instructions, particularly on account of the priming for strategic thinking that is implicit in the test questions.

13.5 Psychological game theory

In this section we consider psychological game theory. Consider an extensive form game in which the set of players is given by $N = \{1, 2, \dots, n\}$. Let Z be the set of terminal nodes of the game. In classical game theory, the utility function of player $i \in N$ is given by the mapping $u_i : Z \rightarrow \mathbb{R}$. In particular, the utility function does not depend on the beliefs of players about the choices and beliefs of other players. However, in a range of situations, it appears that the beliefs of players about the choices and the beliefs of others directly impact their utility.

Example 13.8 (*Guilt*): John frequently visits cities A and B , and he typically uses a taxi to get around. In city A , tipping a taxi driver is considered insulting, while in city B it is the norm to tip a percentage of the fare. Suppose that it is common knowledge that if taxi drivers do not receive a tip, they quietly drive away. In city A , John gives no tip, and feels no remorse from not giving it. However, in city B , the taxi driver expects John to give him a tip (taxi driver's first order belief) and John believes that the taxi driver expects a tip from him (John's second order belief). Based on his second order belief, John cannot bear the guilt of letting the taxi driver down and not paying the tip. Thus, he tips every time he takes a taxi in city B . Clearly, John does not have a utility function of the form $u_i : Z \rightarrow \mathbb{R}$. His utility seems to depend on his beliefs about the beliefs of others.

Guilt is just one of the emotions that creates a potential departure from the classical specification of utility, $u_i : Z \rightarrow \mathbb{R}$. However, players may derive additional psychological utility or disutility from a range of other emotions relating to kindness, anger, surprise, malice, joy, and hope, that can be captured by the beliefs in the game (Elster, 1998).

Can one not simply adjust the material payoffs in classical game theory to take account of such psychological utilities? As Geanakoplos et al. (1989) note in their pioneering paper, this solution will not work because *beliefs themselves are endogenous*. Hence, a new approach, *psychological game theory*, is needed to analyze such situations. Heuristically, if the set B captures the first and higher order beliefs of all players in the game, then in psychological game theory, the utility function takes the form $u_i : Z \times B \rightarrow \mathbb{R}$.⁴⁷ In equilibrium, one requires the actions of players to be consistent with beliefs of first, second, and higher orders and for the beliefs themselves to be consistent in a manner to be specified below. Most applications that we consider truncate beliefs at order 2. Clearly, explicitly incorporating higher order beliefs will be even more cognitively challenging for the players.

We now give further motivating examples in which we only consider first and second order beliefs, followed by a more formal development of the relevant theory.

⁴⁷ Gilboa and Schmeidler (1988) is an early attempt to model hierarchies of initial beliefs of players but they were mainly interested in issues of common knowledge in a manner that does not play a central role in modern psychological game theory.




Figure 13.16 The bravery game.

Source: Reprinted from *Games and Economic Behavior*, 1(1): 60–79. Geanakoplos, J., Pearce, D., and Stacchetti, E., “Psychological games and sequential rationality.” Copyright © 1989, with permission from Elsevier.

Example 13.9 (Geanakoplos et al., 1989: *The Bravery Game*) Consider the extensive form game is shown in Figure 13.16. There are two players, 1 and 2. Player 1’s strategy can take one of two possible values from the set {bold, timid}. Suppose that player 1 plays a mixed strategy in which he is bold with probability p and timid with probability $1 - p$. Player 2 is passive and does not choose any actions in the game. The monetary payoffs of the two players from the actions bold and timid, respectively, are $(2, 2)$ and $(3, 1)$.

A “psychological game” specifies, in addition to a classical extensive form game, the beliefs of players and the utility function is given by $u_i : Z \times B \rightarrow \mathbb{R}$. We only consider the first and second order beliefs. Since there is only one decision node in the game, the only relevant beliefs are the initial beliefs. In games with multiple sequential nodes, we need to take account of the revision in beliefs of the players. Let b_2^1 be the first order belief of player 2 about the probability, p , with which player 1 will play the strategy “bold.” Throughout, we will use the superscript on b for the order of beliefs and the subscript for the player. Let b_1^2 be the second order belief of player 1 about b_2^1 . Thus, beliefs are constructed in a hierarchical manner. Player 2 cares about his belief about how bold player 1 is (i.e., b_2^1). Player 1 cares about his belief about what player 2 thinks of his boldness (i.e., b_1^2). Interpret the payoffs at the terminal nodes that are inclusive of psychological beliefs as utilities.

Player 2 likes player 1 to be bold. At the upper branch where player 1 is revealed to be bold, player 2 receives a utility of $2(1 + b_2^1)$. The stronger player 2’s initial beliefs are that player 1 is bold (higher b_2^1), the greater is his utility from being vindicated. At the lower branch, where player 1 is revealed to be timid, the utility of player 2 drops to $1 - b_2^1$; a higher ex-ante confidence in the boldness of player 1, now hurts player 2 even more. Player 1 prefers to be timid rather than bold, so irrespective of his action, his utility is decreasing in his belief in the confidence that player 2 has in his boldness, b_1^2 .

Note that utilities at the terminal nodes are functions of beliefs, which are themselves endogenous. Hence, we cannot simply apply the backward induction argument of classical game theory. In this simple game, a “psychological equilibrium” arises when (i) conditional on the beliefs, players maximize their utilities, and (ii) strategies and beliefs coincide, i.e.,

$$p = b_2^1 = b_1^2.$$

The bravery game has the following three psychological equilibria.

- (a) $p = b_2^1 = b_1^2 = 1$. In this equilibrium, player 1 plays the strategy bold with probability 1 and the first and second order beliefs reflects this. Given $b_2^1 = b_1^2 = 1$, the strategy

- “bold” strictly dominates “timid,” so any deviation from $p = 1$ reduces the utility of player 1. The utilities of players 1 and 2 are, respectively, 1 and 4.
- (b) $p = b_2^1 = b_1^2 = 0$. In this equilibrium, player 1 plays the strategy timid with probability 1. When $b_2^1 = b_1^2 = 0$, the utility of both players corresponds to the material payoffs only. Hence, player 1 chooses his materially dominant strategy, timid. The respective utilities of the two players are 3 and 1.
- (c) $p = b_2^1 = b_1^2 = 0.5$. Player 1 randomizes equally between his two pure strategies and his utilities from both pure strategies are identical. The expected utility of player 1 is 1.5, while that of player 2 is 1.75.

The utility of player 1 is highest in case (b) where player 2 does not expect player 1 to be bold. As the expectation of player 2 about the boldness of player 1 successively increases, in cases (c) and (a), the utility of player 1 falls. Thus, player 1 benefits most if player 2 holds the most pessimistic expectations about his boldness.

Battigalli and Dufwenberg (2007) consider the role of guilt in the formation of beliefs. Consider the following simple example.

Example 13.10 (Dufwenberg, 2008: Guilt in psychological games.) Consider the two-player sequential move game in Figure 13.17. Player 1 derives utility only from material payoffs but player 2 derives utility from psychological payoffs as well. Suppose that player 2 plays r with probability p . Player 1 believes that player 2 will play r with probability b_1^1 (player 1’s first order beliefs) and the beliefs of player 2 about b_1^1 are given by b_2^2 (player 2’s second order beliefs).

Let us interpret the payoffs of the players at the terminal nodes as utilities. The utility of player 2 from the action profile (d, l) suggest that he is guilt-averse. The higher is b_2^2 , i.e., his belief that player 1 expects him to play r , the lower is the guilt-ridden utility of player 2 from playing l . In a psychological equilibrium of this game,⁴⁸ we require that (i) conditional on their beliefs, players maximize their utility comprising of material and psychological payoffs, and (ii) there is equality of the actual probability with which an action is played and the first and second order beliefs. So, we require that




Figure 13.17 A guilt game.

Source: Dufwenberg (2008). Setten Durlauf, *The New Palgrave Dictionary of Economics*, published 2008, Palgrave Macmillan UK, reproduced with permission of Palgrave Macmillan.

⁴⁸ A formal definition of a psychological equilibrium is given below.

$$b_1^1 = b_2^2 = p. \quad (13.37)$$

In the absence of “guilt aversion” (i.e., if we substitute 3 for $3(1 - b_2^2)$ in Figure 13.17) in the subgame beginning with player 2’s information set, l is the preferred strategy. However, the presence of guilt aversion could induce player 2 to play r instead.

One equilibrium is as follows. Let $b_1^1 = b_2^2 = 1$, player 2 plays r with probability $p = 1$ and player 1 plays d with probability 1. Given the psychological payoffs, this outcome is subgame perfect and satisfies (13.37). See Battigalli and Dufwenberg (2009) for discussion of how this equilibrium may furthermore be justified via a “psychological forward induction” argument.

Example 13.11 Charness and Dufwenberg (2006) conducted trust game experiments in which preplay communication was allowed. For instance, the second mover in a trust game, the trustee, could, in the preplay communication round, send a costless message to the first mover that she will reciprocate the first mover’s (trustor’s) trust. In traditional game theory, this is simply cheap talk and will not affect the outcome. However, Charness and Dufwenberg find that such preplay communication increases the likelihood of trust and trustworthiness. Guilt aversion provides one possible rationalization of these results. Once having made a promise, the trustee feels guilty if he lets down the expectations of the trustor. This model and the results are considered in greater detail in Section 13.5.5.

These results suggest that preplay communication can be a useful coordination device to narrow down the set of equilibria. Not only is the model useful to explain reciprocity in the gift exchange game, it can potentially explain behavior in diverse contexts such as collusion in oligopolistic markets (again through the effectiveness of preplay communication) and tipping behavior in restaurants. The model is thus useful in the domain of positive reciprocity, but does not attempt to explain why individuals hurt each other in negative reciprocity games.

13.5.1 Psychological Nash equilibria in single-stage games

Many factors that are salient in psychological games, for instance, guilt and reciprocity, require an extensive form game in which players move in sequence. However, the resulting general framework based on a refinement of sequential equilibrium for multi-stage games with conditional belief hierarchy may be fairly complicated for readers not already initiated in this area.⁴⁹ Hence, following Geanakoplos et al. (1989), there is pedagogical merit in considering first, a single-stage simultaneous move game.

Let $N = \{1, 2, \dots, n\}$ be the set of players. For any player $i \in N$, the pure strategy set is given by $S_i \subseteq \mathbb{R}$. The set of pure strategy profiles in the game is given by $S = \times_i S_i$. The pure strategy profiles of players other than player i are given by $S_{-i} = \times_{j \neq i} S_j$. The set of all possible mixed strategies of player $i \in N$ is given by $\Sigma_i = \Delta S_i$. The set of all mixed strategy profiles of all players is given by $\Sigma = \times_i \Sigma_i$. The mixed strategy profiles of players other than i belong to the set $\Sigma_{-i} = \times_{j \neq i} \Sigma_j$.

Let B_i^k be the set of k th order beliefs of player i . The beliefs are constructed recursively in a hierarchical manner, as follows. The set of first order beliefs of player i is given by

$$B_i^1 = \Delta S_{-i}, i \in N. \quad (13.38)$$

⁴⁹ See Battigalli and Dufwenberg (2009), that we consider below, for the most general framework.

First order beliefs are simply probability distributions over the pure strategy profiles of others. Define $B_{-i}^1 = \times_{j \neq i} B_j^1$ as the set of first order beliefs of all players other than player i . Then the set of second order beliefs of player i is given by

$$B_i^2 = \Delta(S_{-i} \times B_{-i}^1), i \in N. \quad (13.39)$$

From (13.39), second order beliefs are distributions over the opponents' strategy choices and the opponents' first order beliefs. Proceeding analogously, the set of third order beliefs of player i is given by

$$B_i^3 = \Delta(S_{-i} \times B_{-i}^1 \times B_{-i}^2), i \in N,$$

where $B_{-i}^2 = \times_{j \neq i} B_j^2$ is the set of second order beliefs of players other than player i . Thus, third order beliefs allow for a correlation between strategies, first order beliefs, and second order beliefs of others. In general, beliefs of order $k > 1$ are given by

$$B_i^k = \Delta(S_{-i} \times B_{-i}^1 \times B_{-i}^2 \times \dots \times B_{-i}^{k-1}). \quad (13.40)$$

The hierarchy of beliefs of any player $i \in N$ is denoted by $B_i = (B_i^1, B_i^2, B_i^3, \dots) \in \times_{j=1}^{\infty} B_i^j$. The profile of the hierarchy of beliefs of all players is then given by $B = (B_1, B_2, \dots, B_n) \in \times_{i=1}^n B_i$. In most applications that we consider, it suffices to take into account only beliefs of order $k = 2$. Given the cognitive limitations of humans, it would be surprising if beliefs of order 4 or above played a major role in human decisions, although this is entirely an empirical matter.

Using (13.39), we can recover from $\Delta(S_{-i} \times B_{-i}^1)$, the first order beliefs of the other players by taking the marginal distribution with respect to S_{-i} . If these beliefs coincide with the first order beliefs given in (13.38), then the first and second order beliefs are said to be *coherent*. In general, the beliefs of a player are coherent if the marginals computed from the k th order beliefs coincide with the $k - 1$ order beliefs for all $k > 1$.

In what follows, beliefs are assumed to be coherent for each player $i \in N$.⁵⁰ We also assume mutual knowledge of coherency, so no player believes that the beliefs of any other player are incoherent.

Example 13.12 Consider a two-player, simultaneous move game in normal form in which the strategy sets of the row player and the column player are, respectively, $S_1 = \{U, D\}$ and $S_2 = \{L, R\}$. We only consider beliefs up to order $k = 2$. Let the first order beliefs of the players be given by

$$b_1^1 = \{\sigma_2(L) = q, \sigma_2(R) = 1 - q\}, b_2^1 = \{\sigma_1(U) = p, \sigma_1(D) = 1 - p\}, \quad (13.41)$$

where $p, q \in [0, 1]$, $b_1^1 \in B_1^1 = \Delta S_2$ and $b_2^1 \in B_2^1 = \Delta S_1$ and $\sigma_i(x)$ is the probability with which player $i = 1, 2$ plays the pure strategy x . Suppose that the second order beliefs of player 1, $b_1^2 \in B_1^2$, defined in (13.39), are given as follows.

1. With probability $r_1 \in [0, 1]$, player 2 plays L and his first order beliefs are that player 1 plays U with probability $p_1 \in [0, 1]$ and D with probability $1 - p_1$.

⁵⁰ We avoid introducing extra notation for coherent beliefs, say, $\bar{B}_i$. But if this is confusing, the reader may, henceforth, wish to use the notation B_i for the belief hierarchy of player i and the notation $\bar{B}_i$ for the coherent belief hierarchy of player i .

2. With probability $r_2 \in [0, 1]$, player 2 plays L and his first order beliefs are that player 1 plays U with probability $p_2 \in [0, 1]$ and D with probability $1 - p_2$.
 3. With probability $1 - r_1 - r_2 \in [0, 1]$, player 2 plays R and his first order beliefs are that player 1 plays U with probability $p_3 \in [0, 1]$ and D with probability $1 - p_3$.
- The beliefs of player 1 must be coherent in the following sense. From (13.41), the first order beliefs of player 1 are that player 2 plays L with probability q . We can also recover the probability that player 1 assigns to L from the marginals of his second order beliefs. This probability is $r_1 + r_2$. Hence, coherency requires that $q = r_1 + r_2$.

The utility function of player $i \in N$ is given by

$$\bar{u}_i : B_i \times S \rightarrow \mathbb{R}. \quad (13.42)$$

Denote the vector of utility functions by $u = (u_1, u_2, \dots, u_n)$. Under any mixed strategy profile, $\sigma \in \Sigma$, player $i \in N$ maximizes his expected utility function, $u_i : B_i \times \Sigma \rightarrow \mathbb{R}$, such that

$$u_i(b_i, \sigma) = \sum_{s \in S} (\sigma_1(s_1) \times \sigma_2(s_2) \times \dots \times \sigma_n(s_n)) \bar{u}_i(b_i, s), \quad b_i \in B_i. \quad (13.43)$$

Definition 13.10 (Normal form psychological game): A normal form psychological game is given by $G_P = (N, S, u)$, where the utility function of any player $i \in N$ is given in (13.43).

A mixed strategy profile $\sigma^* = (\sigma_i^*, \sigma_{-i}^*) \in \Sigma$ is a psychologically consistent profile of strategies if it is consistent with the beliefs of players of all orders. For instance, let $N = 2$. Then (σ_1^*, σ_2^*) is a psychologically consistent profile of strategies if $b_1^1 = \sigma_2^*$, $b_2^1 = \sigma_1^*$, $b_1^2 = \sigma_1^*$, $b_2^2 = \sigma_2^* = \dots$. More generally, $(\sigma_i^*, \sigma_{-i}^*)$ is a psychologically consistent profile of strategies if for all i , $b_i^1 = \sigma_{-i}^*$, $b_{-i}^1 = \sigma_i^*$, $b_i^2 = \sigma_i^*$, $b_{-i}^2 = \sigma_{-i}^* = \dots$.

Definition 13.11 (Psychological Nash equilibrium): A psychological Nash equilibrium of a normal form psychological game, $G_P = (N, S, u)$, is a pair (b^*, σ^*) such that:

1. Given b^* , the strategy profile σ^* is a psychologically consistent profile of strategies.
2. $u_i(b_i^*, (\sigma_i^*, \sigma_{-i}^*)) \geq u_i(b_i^*, (\sigma_i, \sigma_{-i}^*))$ for all $i \in N$ and for all $\sigma_i \in \Sigma_i$.

Proposition 13.35 (Geanakoplos et al. 1989): Let G_P be a normal form psychological game. If for all $i \in N$, $u_i : B_i \times \Sigma \rightarrow \mathbb{R}$ is continuous and $B_i \times \Sigma$ is compact, then a psychological Nash equilibrium exists for the game G_P .

The proof of Proposition 13.35 is similar to the existence of a Nash equilibrium in normal form games, so it is omitted.⁵¹

Players may also derive utility from their beliefs about the beliefs of other players about others, not just beliefs about them.⁵² The next example illustrates such a case.

Example 13.13 [Geanakoplos et al. 1989] Suppose that player 1 takes a stranger on a date. Nature (player 0) chooses two possible types of strangers, player 2 and player 3, with equal probability. Player 2 accepts the proposal of player 1 with probability p and player 3 accepts the proposal with probability r . The payoffs at the terminal nodes are in the order of players 1, 2, and 3, respectively. Player 3 only cares for material payoffs and in the event that she is

⁵¹ The interested reader can see the proof in Geanakoplos et al. (1989, pp. 68–9).

⁵² Notice that for any player, $i \in N$ such beliefs also belong to the set B_i .




Figure 13.18 The confidence game in Geanakoplos et al. (1989).

Source: Reprinted from *Games and Economic Behavior*, 1(1): 60–79. Geanakoplos, J., Pearce, D., and Stacchetti, E., “Psychological games and sequential rationality.” Copyright © 1989, with permission from Elsevier.

proposed, she receives a payoff of 1 from rejecting and 0 from accepting, so she always rejects. Players 1 and 2 care for material and psychological payoffs.

Let the first order beliefs of player 1 about p and r be, respectively, b_{12}^1 and b_{13}^1 (in general, we will denote the k th order beliefs of player i about player $j \neq i$ by b_{ij}^k). Player 1 derives higher utility if he is confident of being accepted (perhaps it adds to his self-esteem); the total probability of this event is $b_{12}^1 + b_{13}^1$. However, he also suffers higher disutility if he is more confident and is rejected (perhaps on account of reference point effects). If Player 2 wishes to accept the proposal, she derives higher utility from a more confident player 1 (perhaps due to a social signaling role of partners). Hence, she needs to form beliefs about the unobservable beliefs of player 1, $b_{12}^1 + b_{13}^1$. Let b_{21}^2 be her beliefs about b_{12}^1 and let c_{21}^2 be her beliefs about b_{13}^1 . Then, the beliefs of player 2 about the confidence of player 1 are $b_{21}^2 + c_{21}^2$; these directly enter the utility function of player 2. The corresponding psychological game is shown in Figure 13.18. One of the exercises asks you to find all the psychological Nash equilibria for this game.

Geanakoplos et al. (1989) also extend this framework to extensive form games. However, even in extensive form games, the framework continues to employ the initial beliefs of players.⁵³ This is restrictive because as a sequential game unfolds, the beliefs of players change in response to the observed actions of others. Battigalli and Dufwenberg (2009) extend this framework in several directions that we now consider.

13.5.2 Psychological sequential Nash equilibria in multi-stage games

Following Battigalli and Dufwenberg (2009), consider a finite *multi-stage game*, $\Gamma = \langle N, H \rangle$, with observable actions, where the set of players is $N = \{1, 2, \dots, n\}$, $n \geq 2$, and H is the set of all possible histories of the game.⁵⁴ We assume that there are no moves by nature, no chance moves and there is complete information. The set of stages is given by $t = 0, 1, 2, \dots, \tau + 1$,

⁵³ Geanakoplos et al. (1989, p. 70) are fully aware of the limitations of their approach.

⁵⁴ The extensive form game Γ (see the Appendix on game theory) specifies a finite set of players with finite strategies, a specification of the set of nodes, the players who move at each node, the strategies available to players at each node, the information partition of the game, and the mapping of final nodes to payoffs.

where $\tau + 1$ is the terminal stage in which no player chooses any action (this corresponds to the terminal nodes in an extensive form game). In any stage t , all players simultaneously choose actions from the set of their available actions and the action profile in stage t is given by the vector $a^t = (a_1^t, a_2^t, \dots, a_n^t)$, where a_i^t is the action chosen by player $i \in N$. The profile a^t is immediately observed by all players. It is possible that a subset of the players choose no action in stage t , i.e., they are passive. For instance, the Stackelberg game is a two-stage game in which, in stage $t = 0$, the leader chooses a quantity that is immediately observed by both players but the follower is passive. In stage $t = 1$, the follower chooses a quantity but the leader is passive.

Histories are defined as follows. At the start of play, in stage $t = 0$, we have the null history, h^0 , and players choose the profile of actions a^0 . At the beginning of stage $t = 1$, the history is $h^1 = \{a^0\}$ and players choose the action profile a^1 . At the beginning of stage $t = 2$, the history is $h^2 = \{a^0, a^1\}$, and players choose the action profile a^2 , and so on. In particular, in stage t , the history is $h^t = \{a^0, a^1, \dots, a^{t-1}\}$, and players choose the action profile a^t , having observed the history h^t . Denote by H^t , the set of all possible stage t histories of the game.

The set of actions available to player $i \in N$ at history $h^t \in H^t$ is denoted by $A_i(h^t)$. We assume that $A_i(h^t)$ is non-empty, although it may be a singleton; the interpretation being that such a player is passive at that history. In the terminal history $h^{\tau+1}$ all players are passive, so $A_i(h^t) = \{\text{Do nothing}\}$ for all $i \in N$. The (updated) pure strategy for player $i \in N$ conditional on history h^t having been reached is denoted by $s_i^t : H^t \rightarrow A_i(H^t)$, so $s_i^t = s_i(h^t) \in A_i(h^t)$. A pure strategy for player $i \in N$ for the entire game, s_i , is a pure strategy choice conditional on each possible history having been reached, i.e., $s_i = (s_i^0, s_i^1, \dots, s_i^\tau)$.⁵⁵

Each player moves at most once at each history of the game. Hence, the set of all actions of player i for the game is $A_i = \cup_{h^t \in H^\tau} A_i(h^t)$. In short, a pure strategy for player i for the entire game is given by $s_i : H^\tau \rightarrow A_i$. Denote the set of pure strategies of player i for the entire game by S_i . We may define $S = \prod_{i \in N} S_i$ as the set of strategy profiles of all players for the entire game and $S_{-i} = \prod_{j \neq i} S_j$. Denote by $S_i(h^t)$ the set of strategies that allow history h^t to materialize. Hence, we can define $S(h^t) = \prod_{i \in N} S_i(h^t)$ and $S_{-i}(h^t) = \prod_{j \neq i} S_j(h^t)$, respectively, as the strategy profiles of all players, and the strategy profiles of players other than i , that allow history h^t .

In classical game theory, the utility of players is given by the function $u_i : H^{\tau+1} \rightarrow \mathbb{R}$. However, such a utility function does not explain the outcomes in Examples 13.8, 13.9, and 13.10 considered above.

We now define the behavior strategies of any player $i \in N$. At any information set given the history $h^t \in H^\tau$, the probability distribution over the stage t actions of player i is given by $\Delta A_i(h^t)$. A behavior strategy for player i , Φ_i , specifies the mixed strategy of the player at each history of the game, so it is an element of the set

$$\Phi_i = \Delta A_i(h^0) \times \Delta A_i(h^1) \times \dots \times \Delta A_i(h^\tau) = \prod_{h^t \in H^\tau} \Delta A_i(h^t). \quad (13.44)$$

Note that in contrast to behavior strategies, the set of mixed strategies of player i is given by $\Sigma_i = \Delta S_i$.⁵⁶ Define the behavior strategies of players other than player i by $\Phi_{-i} = \prod_{j \neq i} \Phi_j$.

⁵⁵ Just to be absolutely clear, the (updated) pure strategy, s_i^t , and the pure strategy for the entire game, s_i , are identical up to history h^t .

⁵⁶ Kuhn (1953) showed that in games of perfect recall, mixed and behavior strategies are equivalent; see Chapter 6 in Maschler et al. (2013).




Figure 13.19 A simple example of a sequential move psychological game.

Source: Reprinted from *Journal of Economic Theory* 144(1): 1–35. Battigali, P. and Dufwenberg, M. “Dynamic psychological games.” Copyright © 2009, with permission from Elsevier.

Suppose that some history $h^t \in H^t$ has materialized. Denote by $\Phi_i(h^t)$, the behavior strategy of player i that assigns probability 1 to all the actions of player i that allow history h^t to materialize. Let $\Phi_{-i}(h^t) = \prod_{j \neq i} \Phi_j(h^t)$.

We now define for any player $i \in N$, first order beliefs, second order beliefs, and so on. These are the analogues of (13.38) to (13.40), however, there is one main difference. In a multi-stage game, we allow for the beliefs of players to evolve as play progresses through the various stages. The hierarchies of beliefs can be revised, which is a major difference from the framework of Geanakoplos et al. (1989).⁵⁷

The most general case of multistage games may be pedagogically too complex for some readers, although this is the case considered in Battigalli and Dufwenberg (2009). For this reason, we first consider the issues in a simple example. This is followed by a heuristic discussion of how the insights from this example can be extended to general finite multistage games. Finally, following Dufwenberg and Kirchsteiger (2004), we consider an application of the general framework to issues of reciprocity.

Example 13.14 (Battigalli and Dufwenberg, 2009): Consider the two-player, two-stage game shown in Figure 13.19. Each player has two actions at each decision node, $t = 0, 1$. In each stage, only one player can move, while the other player is passive. Ann’s set of actions at the two histories are $A_A(h^0) = \{\text{Trust, Don't trust}\}$, $A_A(h^1) = \{\text{Do nothing}\}$. Bob’s action sets are $A_B(h^0) = \{\text{Do nothing}\}$ and $A_B(h^1) = \{\text{Share, Grab}\}$.

The set of behavior strategies of Ann, Φ_A , specifies all possible probabilities with which the action Trust is played at history h^0 . Likewise the set of behavior strategies of Bob specifies all possible probabilities with which the action Share is played at history h^1 . The strategy profiles of players leading to the three terminal nodes are {Don't trust, Do nothing}, {Trust, Grab}, and {Trust, Share}. The utility functions of the players map each of the terminal nodes to real numbers that we call utilities, as shown in Figure 13.19. The payoffs in the classical case from the strategy profile {Trust, Grab} are (0, 4) and the utilities from this strategy profile in the psychological game, as we see below, are (0, $4 - 5\beta$)

⁵⁷ To be sure, even in the classical framework, one allows for Bayesian updating of beliefs as play proceeds in a multistage game. Common knowledge of rationality, typically assumed in the classical framework, allows all higher order beliefs to be revised as well (the infinite regress case). However, in applications of psychological games, beliefs are typically truncated at some small integer value (typically between 2 and 3). In contrast when we consider the framework of psychological games for sequential games below, we shall allow not only for standard Bayesian belief revision, but also for revision of beliefs that directly enter into the utility functions of players.

Consider first the classical case that corresponds to $\beta = 0$ in Figure 13.19. The unique subgame perfect outcome is for Bob to play Grab at history h^1 and for Ann to play Don't trust at history h^0 . Hence, the utilities of the players are given by (1,1).

Let us now apply psychological game theory and restrict beliefs up to a maximum of order 2 for each player. We shall assume that players have utility functions as in (13.42), so they derive utility from their own beliefs only. At history h^0 , the set of first order beliefs of Ann, $B_A^1(h^0)$, must specify all possible beliefs about all possible strategies of Bob at nodes where Bob is not passive. A particular first order belief of Ann is that at history $h^1 = \text{Trust}$, Bob will play the action Share with probability α and the action Grab with probability $1 - \alpha$. We write this as

$$b_A^1(h^0) = \Pr(\text{Share} \mid \text{Trust}) = \alpha.$$

Letting $0 \leq \alpha \leq 1$ allows us to list all members of the set $B_A^1(h^0)$. One may specify Ann's second order beliefs but these are redundant in this game.

Now suppose that Bob is asked to make a choice between Share and Grab. Bob must assign probability 1 that he is at the history $h^1 = \text{Trust}$ (this is exactly as in Bayesian updating).⁵⁸ We now need to specify the belief hierarchy of Bob at history $h^1 = \text{Trust}$. Bob's first order beliefs about Ann $b_B^1(h^1 = \text{Trust})$ are irrelevant because Ann is passive at this history. The second order beliefs of Bob, $b_B^2(h^1 = \text{Trust})$, are probability distributions over sets of the form $\{\text{Trust}, \text{Don't}\} \times B_A^1(h^0)$. Let β be the probability assigned by Bob to the first order belief of Ann, $b_A^1(h^0) = \alpha$, when Bob observes that Ann has played the action Trust.

It remains to specify the utilities of the players. Recall that in (13.42), in a static game, we specified the utility of player i as $\bar{u}_i : B_i \times S \rightarrow \mathbb{R}$. However, in multistage games, we need to specify the beliefs of players that enter into the utility functions of players at each history of the game. A natural extension is given by

$$\bar{u}_i(h^k) : S(h^k) \times B_i(h^k) \rightarrow \mathbb{R}, h^k \in H, \quad (13.45)$$

where $B_i(h^k)$ is the belief hierarchy of player i at history h^k . In particular, we are interested in the role of "guilt aversion" in this example. This is modeled by allowing β to directly enter the payoff of Bob, following the strategy profile (Trust, Grab). Bob feels guilty from letting down the expectations of Ann, hence, he gets disutility from choosing the action Grab. In particular, we assume that Bob's utility from Grab is $4 - 5\beta$, thus, his disutility is increasing in his belief that Ann expects him to choose the action Share. This explains the utilities depicted in Figure 13.19 corresponding to each of the terminal nodes.

This completes the description of the psychological game in Figure 13.19. In addition to the usual extensive form game $\langle N, H \rangle$, we have now specified the belief hierarchy of each player at each node (although we have ignored beliefs that are not relevant) and the utilities of players at each terminal node, consistent with (13.45). As in static games (see Example 13.12), we require that beliefs be coherent in multistage games, although in the simple setup of this example this is fairly trivial.

How do we find the psychological equilibrium? We require that each player maximize his utility at each history of the game, after assigning probability 1 to the strategy profile that has led to that history, conditional on beliefs at that history and expectations of the strategy profiles in subsequent histories. Further, beliefs must be coherent in the sense of Example 13.12

⁵⁸ We may at this stage specify any first order beliefs of Bob at history h^0 . They do not affect the calculations.

(trivially satisfied in this example) and the strategies and beliefs of players must be consistent with each other. This is made precise in a manner that is analogous to a sequential equilibrium.

Let us now find the psychological equilibria of this game. Ann must choose her action at history h^0 to maximize her expected utility given her beliefs. If she chooses not to trust, she gets 1. If she chooses to trust, she gets 2α , so she should trust if $\alpha > \frac{1}{2}$. What action should Bob choose at history $h^1 = \text{Trust}$? Bob has no subsequent histories to worry about at history $h^1 = \text{Trust}$. So he chooses the action Share if $2 > 4 - 5\beta$, or if $\beta > \frac{2}{5}$. Here is one sequential psychological Nash equilibrium for this game: Ann plays Trust with probability 1, Bob plays Share with probability 1, and $\alpha = \beta = 1$. This satisfies the following requirements.

- (i) Ann's choice of action is optimal, given her belief, $\alpha = 1$, and the equilibrium choice of Bob in the next stage.
- (ii) Bob's choice is optimal, given the preceding choice of Ann, and his beliefs as captured by $\beta = 1$.
- (iii) The first order beliefs of Ann are derived from the equilibrium actions. The second order beliefs of Bob are consistent with the first order beliefs of Ann.

Example 13.15 (Example 13.14 continued) : Now suppose that instead of utilities of the form in (13.45), we have the following general utility function proposed by Battigalli and Dufwenberg (2009)

$$\bar{u}_i(h^k) : S(h^k) \times B_i(h^k) \times \prod_{j \neq i} (B_j(h^k) \times S_j(h^k)) \rightarrow \mathbb{R}. \quad (13.46)$$

In (13.46), the beliefs of other players directly enter the utility function of the player. In models of "social conformity," the utility functions of players also reflect the beliefs of others. As another example, consider a doctor who feels concerned about a patient's anxiety about her disease. The doctor has superior information about the disease and the patient's health. Yet, in the absence of mandatory disclosure, the doctor might be unwilling to reveal the true extent of the problem to the patient because his concern for the patient's anxiety following disclosure may reduce his own utility.

In the context of Example 13.14, one may redraw Figure 13.19 as shown in Figure 13.20.

In Figure 13.20, we have replaced β by α , where α and β are defined in Example 13.14. Since Bob maximizes his expected utility at history $h^1 = \text{Trust}$, the expected utility from the action Grab is $E[4 - 5\alpha \mid h^1] = 4 - 5\beta$. This leads to the same outcome as in Example 13.14.




Figure 13.20 Dependence of utilities on beliefs of others.

Source: Reprinted from *Journal of Economic Theory* 144(1): 1–35. Battigalli, P. and Dufwenberg, M. "Dynamic psychological games." Copyright © 2009, with permission from Elsevier.

However, one needs beliefs of one order lower, which may simplify the problem. In most cases that we consider below, we will be concerned with utility functions of the form (13.45) rather than (13.46). ■

Example 13.14 gives a blueprint of how one might proceed in any finite multistage game when players have the utility function in (13.45). We give a heuristic discussion of the general case, followed by a concrete application to reciprocity in Section 13.5.3.⁵⁹ Consider a multistage game $\Gamma = \langle N, H \rangle$. The belief hierarchies are constructed in a manner similar to that for static games in (13.38) to (13.40). The main difference is that we need to consider belief hierarchies at each history of the game and these beliefs are about the history-contingent strategies of others. At each possible history of the game, $h^k \in H$, let $B_i(h^k) = \{b_i^1(h^k), b_i^2(h^k), \dots, b_i^l(h^k)\}$ be the set of belief hierarchies of player i , up to level l . Example 13.14 shows how the first and second order beliefs are constructed in a simple game. Let $B_i = (B_i(h^k))_{h^k \in H}$ be the set of belief hierarchies of player i for all possible histories and let $B = (B_i)_{i \in N}$ be the set of belief hierarchies for all players for the entire game.

The first order beliefs at any history are probability distributions over the strategies of the other players. The second order beliefs are probability distributions over the strategies of others and the first order beliefs of other players and so on. This hierarchy in beliefs is *conditional* in the sense that it conditions on the history of play up to that point and on the strategies and beliefs of others (as in (13.40)). At any history, players must assign probability 1 to all strategy profiles that have led to that particular history. For instance, in Figure 13.19, if Bob is asked to move, then he must assign probability 1 to the action Trust having been played by Ann. We require beliefs to be coherent in the sense of Example 13.12. With these components in place, we can define a psychological game as $\Gamma = \langle N, H, (\bar{u}_i(h^k))_{i \in N, h^k \in H} \rangle$, where $\bar{u}_i(h^k)$ is defined in (13.45) (or more generally in (13.46)).

A behavior strategy for player i , Φ_i , is defined in (13.44). Denote by $\Phi = (\Phi_i)_{i \in N}$ the profile of behavior strategies of all players for the game. An *assessment*, defined as (Φ, B) , is said to be consistent if, for each history, the first order beliefs in B are derived from Φ and the higher order beliefs place probability 1 on lower order beliefs (see Example 13.14 for an illustration). Finally, a consistent assessment (Φ, B) is a sequential equilibrium of the psychological game $\Gamma = \langle N, H, (\bar{u}_i(h^k))_{i \in N, h^k \in H} \rangle$ if at each history $h^k \in H$, the pure strategies in the support of the behavior strategies of the players maximize their utility functions.

A NOTE ON THE COMPARISON WITH GAMES OF ASYMMETRIC INFORMATION

We describe the basics of games of asymmetric information in the Appendix. Let $N = \{1, 2, \dots, n\}$ be the set of players. In such games, the utility function of player $i \in N$ is given by $u_i : Z \times \Theta \rightarrow \mathbb{R}$ where Θ is a vector of exogenously given characteristics of the players such as tastes and ability. Each possible characteristic may be considered as the type of a player that is drawn by nature from a commonly known probability distribution. Players observe the type that nature draws for them but do not observe the types drawn for the other players (they do, however, typically have common priors about the distribution of types of other players). One may then have a hierarchy of beliefs about the exogenous types of players. These beliefs play an important role in the determination of an equilibrium outcome in the game. However, the payoffs at each terminal node are exogenous. By contrast, in psychological games, the hierarchy of beliefs determine not only the

⁵⁹ The formal framework requires the use of conditional belief hierarchies formulated in Battigalli and Siniscalchi (1999); see Battigalli and Dufwenberg (2009) for the details.

outcome of the game but also determines, endogenously, the material payoffs at each end node. Hence, these are quite different classes of games.

The assumption of common knowledge of the utility function in psychological game theory is a strong one. There is merit in developing a more general framework that incorporates asymmetric information about, say, the utility functions of other players, within psychological games. For an attempt in this direction, in the context of trust games, see Attanasi et al. (2015).

13.5.3 Fairness and reciprocity in extensive form games

In this section, we give a concrete application of the framework proposed in Section 13.5.2 to issues of reciprocity. The seminal application to reciprocity, using the machinery described in Section 13.5.1 for one-stage games is due to Rabin (1993). However, many important economic applications can be modeled in a satisfactory manner only in sequential games. In order to adequately explore the role of reciprocity in such cases, Rabin's model needs to be extended. When players move in sequence, they can observe the actions of others and infer the kindness (or unkindness) of their opponents prior to choosing their own actions. Furthermore, after observing the actions of others, players can update their belief hierarchy, which is not possible in normal form games. In this section, we consider the work of Dufwenberg and Kirchsteiger (2004) who extend Rabin's framework to sequential form games.

Consider the sequential prisoner's dilemma game in panel A of Figure 13.21. In the normal form corresponding to panel A (not shown), the strategy profile (C, c) is a psychological Nash equilibrium in the single-stage model in Rabin (1993). However, a consideration of the extensive form makes it difficult to sustain (C, c) as an equilibrium. Even if player 2 initially held a belief that player 1 is kind, once player 1 plays D , such a belief is difficult to entertain. By playing D , player 1 ensures that player 2 gets a lower payoff for each of her strategy choices, relative to the case where player 1 plays C .

In the extensive-form model proposed by Dufwenberg and Kirchsteiger (2004), the beliefs about players evolve over the course of an extensive form game, allowing player 2 to interpret the choice of D by player 1 as unkind. With some minor modifications, the Dufwenberg and




Figure 13.21 Sequential prisoner's dilemma game.

Source: Reprinted from *Games and Economic Behavior* 47: 268–98. Dufwenberg, M. and Kirchsteiger, G., "A theory of sequential reciprocity." Copyright © 2004, with permission from Elsevier.

Kirchsteiger (2004) model becomes a special case of Battigalli and Dufwenberg (2009). Falk and Fischbacher (2006) also consider issues of reciprocity but extend this framework further to allow concerns for inequity-aversion (see Section 13.5.4).

Consider a *multi-stage game*, Γ , where the set of players is $N = \{1, 2, \dots, n\}$, $n \geq 2$. There are no moves by nature and *each player only moves once* in any stage. Let Φ_i be the set of behavior strategies of player $i \in N$ for the game (see (13.44)). We denote a typical element of Φ_i by ϕ_i . A behavior strategy specifies the (possibly mixed) strategies that players play at each possible history of the multistage game when they are asked to make a choice. An example of a behavior strategy in the game in Panel A of Figure 13.21 is that player 1 plays C and player 2 plays c following C and d following D . Let $\Phi = \times_i \Phi_i$ be the set of behavior strategy profiles for the game, with typical element given by ϕ .

Dufwenberg and Kirchsteiger (2004) restrict attention to the case where players play pure strategies only. Hence, the interpretation of any particular behavior strategy of the opponents is that players believe that they are facing a population of pure strategy playing opponents, whose mixture is such that the population proportions are consistent with the behavior strategy.

We will only take account of beliefs of order 1 and 2 for each player. In order to capture the notion of history dependence of strategies and beliefs in the extensive form, we define, for any player $i \in N$ and at any history $h \in H$ in the game the following three entities. $\phi_i(h)$ is the behavior strategy of player i ; $B_i^1(h) = \left(B_{ij}^1(h)\right)_{j \neq i}$ and $B_i^2(h) = \left(B_{ijk}^2(h)\right)_{j \neq i, j \neq k}$ are, respectively, first and second order beliefs of the player. Notice that second order beliefs have elements of the form $B_{ijk}^2(h)$, i.e., the beliefs of player i about the first order beliefs of player $j \neq i$ about the strategy of player $k \neq j$.

Consider, for instance, the game in Panel A in Figure 13.21. Suppose that, at history h^0 , the first order belief of player 2 about the action of player 1, $b_{21}^1(h^0) \in B_{21}^1(h^0)$, is that player 1 will play C with probability p and D with probability $1 - p$. Consider the history of the game given by $h^1 = \{D\}$. In this case, consistent with Bayes' rule, the first order belief of player 2 about player 1, $b_{21}^1(h^1)$, specifies that player 1 plays C with probability 0 and D with probability 1. When player 2 observes the history $h^1 = \{D\}$, player 2 assumes that player 1 has deliberately played D , i.e., playing D is not a mistake on the part of player 1. This assumption allows player 2 to confidently retaliate against player 1.

We can now define the payoffs and utilities of players. The material payoffs of players, used in classical game theory, at any terminal history $z \in Z$ are defined by the mappings $\pi_i(z) : \Phi(z) \rightarrow \mathbb{R}$, $i \in N$. In contrast, in psychological games, players derive utility not at the final nodes that the game reaches but, at any history $h \in H$, (via $\Phi(h)$) from their own first and second order beliefs. Suppose that the utilities of players in psychological games at any history $h \in H$ are given analogously to (13.45), (13.46) by

$$u_i(h) : \Phi(h) \times \prod_{j \neq i} \left(B_{ij}^1(h) \times \prod_{k \neq j} B_{ijk}^2(h) \right) \rightarrow \mathbb{R}, i \in N. \quad (13.47)$$

Consider Panel B in Figure 13.21. In this case, if player 1 plays the strategy U , then both players get the inefficient payoff of -1000 each. Suppose that the first order beliefs of player 1 at history h^0 specify that player 2 will play c at the left node and d at the right node. By his actions, player 1 can now give player 2 one of the payoffs in the set $\{-1000, 0, 1\}$. Relative to the payoff of -1000 , a payoff of 0 seems high. Should player 2, on observing the history $h^1 = \{D\}$, infer that player 1 is being kind to him? Dufwenberg and Kirchsteiger (2004) view the inference of kindness by player 2, in this case, as perverse. They argue that player 2 should have the same beliefs about

the kindness of player 1 in both panels in Figure 13.21. This gives rise to the notion of efficient strategies in the next definition.

Definition 13.12 (*Efficient strategies*): A strategy profile is efficient if, given the history of play, and the subsequent choices of players, there does not exist another strategy profile that gives at least one player a higher payoff and no player a lower payoff. Denote the set of efficient strategies of player $i \in N$ at history h by $E_i(h)$.

Clearly, the choice of U by player 1 in Panel B of Figure 13.21 cannot be part of an efficient strategy profile.

Definition 13.13 Consider some history of the game, $h \in H$. Suppose that the first order belief of player i about player $j \neq i$ is $b_{ij}^1(h)$. Denote the vector of such beliefs by $B_i^1(h) = (b_{ij}^1(h))_{j \neq i}$ with typical element given by $b_i^1(h)$.

At any history h , if player i plays the behavior strategy $\phi_i(h)$ and believes that other players will play the vector of strategies $b_i^1(h)$, then player i believes that the material payoff of player $j \neq i$ is $\pi_j(\phi_i(h), b_i^1(h))$. We use these payoffs to define equitable payoffs, next.

Definition 13.14 (*Equitable payoff*): Consider some history of the game, $h \in H$. The equitable payoff of player j at node h , from the perspective of player i , is defined as

$$\bar{\pi}_j(b_i^1(h)) = \frac{1}{2} \left[\begin{array}{l} \{\max \pi_j(\phi_i(h), b_i^1(h)) : \phi_i(h) \in \Phi_i(h)\} \\ + \{\min \pi_j(\phi_i(h), b_i^1(h)) : \phi_i(h) \in E_i(h)\} \end{array} \right]. \quad (13.48)$$

In (13.48), note that we focus on different sets in calculating the maximum and minimum payoffs. In particular, for the reasons described above, when we calculate the minimum payoff of the player, we restrict attention to efficient strategies only.

We now define the kindness functions. In each case, departures of payoffs from equitable payoffs are used to measure kindness of players.

Definition 13.15 (*Kindness of player i to player j*): Suppose that at some history of the game, $h \in H$, player i chooses the behavior strategy $\phi_i(h)$ and believes that the strategy profile of other players is $b_i^1(h)$. Then player i perceives that his kindness towards player j is given by $k_{ij} : \Phi_i(h) \times B_i^1(h) \rightarrow \mathbb{R}$, where $\Phi_i(h)$ is the set of behavior strategies of player i at history h and $B_i^1(h)$ is the set of his first order beliefs about the behavior strategies of all other players. In particular,

$$k_{ij}(\phi_i(h), b_i^1(h)) = \pi_j(\phi_i(h), b_i^1(h)) - \bar{\pi}_j(b_i^1(h)); \phi_i(h) \in \Phi_i(h), b_i^1(h) \in B_i^1(h). \quad (13.49)$$

Thus, player i is kind to player j at history $h \in H$ if his actions ensure a payoff to j greater than his equitable payoff (in this case, $k_{ij} > 0$). Analogously, player i is neutral to player j at history $h \in H$ if $k_{ij} = 0$ and unkind if $k_{ij} < 0$.

We now define the kindness that player i perceives from the actions of player $j \neq i$, which we term as *reciprocity*, denoted by $\tilde{k}_{ij}$. At some history of the game, $h \in H$, when player i attempts to infer the reciprocity of player j he must infer what player j thinks about the actions of other players, $b_j^1(h)$, which is not observable. As a proxy, player i uses his second order beliefs, $b_{ij}^2(h)$, about player's j first order beliefs $b_j^1(h)$; $b_{ij}^2(h)$ is an element of the set $B_{ij}^2(h) = \Pi_{k \neq j} B_{ijk}^2(h)$, and

$$b_{ij}^2(h) = (b_{ij1}^2(h), \dots, b_{ijj-1}^2(h), b_{ijj+1}^2(h), \dots, b_{ijn}^2(h)). \quad (13.50)$$

Definition 13.16 (Reciprocity of player j to player i): Consider some history of the game, $h \in H$. Fix the first order beliefs of player i about player j , $b_{ij}^1(h)$. We now vary the second order beliefs of player i about player j 's beliefs about other players $k \neq j$, i.e., the vector b_{ij}^2 . The belief of player i about the kindness of player j towards him is given by a function, $\tilde{k}_{ij} : B_{ij}^1(h) \times B_{ij}^2(h) \rightarrow \mathbb{R}$ such that

$$\tilde{k}_{ij}(b_{ij}^1(h), b_{ij}^2(h)) = \pi_i(b_{ij}^1(h), b_{ij}^2(h)) - \bar{\pi}_i(b_{ij}^2(h)), \quad b_{ij}^1(h) \in B_{ij}^1(h), b_{ij}^2(h) \in B_{ij}^2(h), \quad (13.51)$$

where $\bar{\pi}_i$ is the equitable payoff of player i . $\pi_i(b_{ij}^1(h), b_{ij}^2(h))$ is the payoff of player i when he believes that player j plays $b_{ij}^1(h)$, and his (player i 's) beliefs about the belief of player j about the strategies of all non- j players is $b_{ij}^2(h)$, as defined in (13.50). (To be absolutely sure, $b_i^2(h) = (b_{i1}^2(h), \dots, b_{ii-1}^2(h), b_{ii+1}^2(h), \dots, b_{in}^2(h))$).

Recall that $B_{ij}^1(h)$ is the set of all first order beliefs of player i about player j at history h and $\Phi_j(h)$ is the set of all possible behavior strategies of player j at history h , hence, $B_{ij}^1(h) = \Phi_j(h)$. Similarly the set of all possible second order beliefs of player i about j is given by $B_{ij}^2(h)$; $b_{ij}^2(h)$, defined in (13.50) is a typical member of this set. $B_j^1(h)$ is the set of first order beliefs of player j about other players. In equilibrium these sets are equal, $B_{ij}^2(h) = B_j^1(h)$. Thus, comparing (13.49), (13.51) we get that mathematically, the function $\tilde{k}_{ij}$ is equivalent to k_{ji} ; for future reference we write this as

$$\tilde{k}_{ij} = k_{ji}. \quad (13.52)$$

(13.52) facilitates the computation of the function $\tilde{k}_{ij}$. Rather than computing the reciprocity of player j towards i , we can simply compute, using (13.49), the kindness function of player j towards i ; this is a simpler computation because it only requires first order beliefs.

Definition 13.17 (Utility function of players): Consider some history of the game, $h \in H$. The utility function of player $i \in N$ in (13.47) is operationalized in the following form

$$u_i(\phi(h), b_i^1(h), \{b_{ij}^2\}_{j \neq i}) = \pi_i(\phi(h)) + \sum_{j \neq i} \lambda_{ij} k_{ij}(\phi_i(h), b_i^1(h)) \tilde{k}_{ij}(b_{ij}^1(h), b_{ij}^2), \quad (13.53)$$

where $\lambda_{ij} \geq 0$, $\phi(h) \in \Phi(h)$, $b_i^1(h) \in B_i^1(h)$, $b_{ij}^1(h) \in B_{ij}^1(h)$, $b_{ij}^2 \in B_{ij}^2$.

The first term on the RHS of (13.53) gives the material payoff of player i , which depends on the behavior strategies chosen by all players. The second term gives the psychological payoffs that capture the interaction of kindness and reciprocity terms summed across all the other players. For each player $j \neq i$, the parameter λ_{ij} gives the *salience of the psychological payoff* with respect to that player. Each of these terms is made up of the product $k_{ij} \tilde{k}_{ij}$, i.e., the kindness of i towards j (k_{ij}) and the reciprocity of player j towards i ($\tilde{k}_{ij}$). Using (13.52) we can also write $k_{ij} \tilde{k}_{ij} = k_{ij} k_{ji}$. Thus, the psychological payoffs can be enhanced if player i responds to kindness with kindness and unkindness with unkindness; this is the essence of reciprocity payoffs.

Definition 13.18 (Psychological game with reciprocity incentives): The game $\Gamma^* = (\Gamma, (u_i)_{i \in N})$ where u_i is defined in (13.53) is known as a psychological game with reciprocity incentives.

Recall that a behavior strategy for a player, ϕ_i , is a description of the mixed strategy that player i chooses at each history $h \in H$. Let us pick a particular history $h \in H$. Denote by $\Phi_i^h \subseteq \Phi_i$,

the subset of behavior strategies such that Φ_i^h and Φ_i coincide for all histories preceding history h and succeeding history h (so they potentially differ only at history h). However, in Φ_i^h , we are free to choose the behavior strategy of player i at history h .

Definition 13.19 (Sequential reciprocity equilibrium, SRE): A strategy profile $\phi^* = (\phi_1^*, \phi_2^*, \dots, \phi_n^*)$ is a sequential reciprocity equilibrium if:

- (a) $\phi_i^* \in \arg \max_{\phi_i \in \Phi_i^h} u_i(\phi(h), b_i^1(h), \{b_{ij}^2\}_{j \neq i})$ for all $h \in H$ and for all $i \in N$, where u_i is defined in (13.53).
- (b) $b_i^1 = \phi_{-i}^*$ for all $i \in N$, where ϕ_{-i}^* is the vector of optimal behavior strategies of players other than i .
- (c) $b_{ij}^2 = (\phi_1^*, \dots, \phi_{j-1}^*, \phi_{j+1}^*, \dots, \phi_n^*)$ for all i and j such that $i \neq j$, where $b_{ij}^2(h)$ is defined in (13.51).

When $\lambda_{ij} = 0$ for all $i, j \in N$, then the psychological payoffs are irrelevant, so Definition 13.19 simply corresponds to the definition of a subgame perfect Nash equilibrium. However, when $\lambda_{ij} > 0$, the outcomes from the two concepts will typically differ. The reason is that psychological payoffs depend on the history of the game through the role of intentions that are revealed from the choices made by other players. However, in a subgame perfect equilibrium, as one solves the game backwards, one does not take account of such intentions (except updating of beliefs, consistent with Bayes' rule). Thus, backward induction is not suitable to solve the multi-stage game used in this section.

The proof of the existence of a SRE may appear to be difficult. However, from Definition 13.19(a), each individual optimizes at a particular $h \in H$, holding fixed the strategies at all other $h' \in H$ (hence, the restriction $\phi_i \in \Phi_i^h$). This observation simplifies the proof and an application of the Kakutani fixed point theorem to best reply correspondences ensures that a SRE exists; see Dufwenberg and Kirchsteiger (2004, p. 280).

The framework outlined above differs from Rabin (1993) in the definition of the efficient set $E_i(h)$, as well as in the definitions of the kindness functions. The kindness functions here are defined as differences from equitable payoffs, while in Rabin (1993) these are ratios that are normalized by a factor (difference between maximum and minimum payoffs) that restricts the ratios to lie between 0 and 1. Further, the sets from which the individual components on the RHS in (13.49), (13.51) are drawn are different in Rabin's framework.⁶⁰ We have also slightly modified the payoff functions in Dufwenberg and Kirchsteiger (2004) in this section in order to make them comparable to the framework developed in Section 13.5.2.

As an illustration of a SRE, consider the next example that we also use later to examine the evidence on reciprocity.

Example 13.16 (Dufwenberg et al., 2011) Consider a three-player public goods game, $N = \{1, 2, 3\}$, in which each of the three players is given an endowment of 20 tokens. The players simultaneously choose contributions $s_i \in S_i = [0, 20]$; we focus only on degenerate behavior strategies in this example. The total amount of the public good produced is $G = \frac{3}{2}(s_1 + s_2 + s_3)$. Each player gets a return of $\frac{1}{3}G = \frac{1}{2}(s_1 + s_2 + s_3)$ from the public good. The strategy profiles for the game are $s = (s_1, s_2, s_3) \in S = [0, 20]^3$. In the absence of any reciprocity considerations, the material payoffs of the players are given by $\pi_i : S \rightarrow \mathbb{R}$ such that

⁶⁰ For a discussion of the differences between the two frameworks and more on the rationale for the framework discussed in this section see Dufwenberg and Kirchsteiger (2004, Section 5).

$$\pi_i(s) = (20 - s_i) + \frac{1}{2} (s_1 + s_2 + s_3), \quad i \in N. \quad (13.54)$$

In a Nash equilibrium, it is optimal for each player to free-ride, so $s_i^N = 0$, $i = 1, 2, 3$. Now introduce the possibility of reciprocity. The utility function of each player has the form $u_i : S \times B \rightarrow \mathbb{R}$, where B is the set of all possible first and second order beliefs of each player about choices and beliefs. Denote by b_{ij}^1 , $i \neq j$, the first order beliefs of player i about the contributions, s_j , of another player j . Denote by b_{ijk}^2 the second order beliefs of player i about the first order beliefs of player j about the contributions of player k , b_{ijk}^1 , $j \neq k$. Using (13.53), and for all i , letting $\lambda_{ij} = \lambda_i$, $i \neq j$, the utility function of player $i \in N$ is

$$u_i = (20 - s_i) + \frac{1}{2} (s_1 + s_2 + s_3) + \lambda_i (k_{ij}\tilde{k}_{ij} + k_{ik}\tilde{k}_{ik}), \quad i \neq j \neq k. \quad (13.55)$$

The kindness of player i to players j and k is particularly easy to calculate in public goods games, due to the linearity and independence of contributions. Using (13.48), (13.54), we can compute the equitable payoff of player j , from the perspective of player i . Conditional on his beliefs about the contributions of others, player i can ensure the highest payoff for player j by contributing the entire endowment of 20, and the least payoff by contributing nothing. Thus,

$$\begin{aligned} \bar{\pi}_j(b_{ij}^1, b_{ik}^1) &= \frac{1}{2} \left[\left((20 - b_{ij}^1) + \frac{1}{2} (20 + b_{ij}^1 + b_{ik}^1) \right) + \left((20 - b_{ij}^1) + \frac{1}{2} (0 + b_{ij}^1 + b_{ik}^1) \right) \right], \\ &\Rightarrow \bar{\pi}_j(b_{ij}^1, b_{ik}^1) = (20 - b_{ij}^1) + \frac{1}{2} (b_{ij}^1 + b_{ik}^1) + \frac{1}{2} 10. \end{aligned}$$

The kindness of player i to player j , k_{ij} , can now be calculated by using (13.49)

$$k_{ij}(s_i, b_{ij}^1, b_{ik}^1) = \left[(20 - b_{ij}^1) + \frac{1}{2} (s_i + b_{ij}^1 + b_{ik}^1) \right] - \left[(20 - b_{ij}^1) + \frac{1}{2} (b_{ij}^1 + b_{ik}^1) + \frac{1}{2} 10 \right]. \quad (13.56)$$

The first term in (13.56) is the payoff of player j , $\pi_j(s_i, b_i^1(h))$, where $b_i^1(h) = (b_{ij}^1, b_{ik}^1)$. Varying his own contributions, s_i , player i can change player j 's payoff; this results in variations in k_{ij} . Simplifying (13.56), we get

$$\Rightarrow k_{ij}(s_i, b_{ij}^1, b_{ik}^1) = \frac{1}{2} (s_i - 10).$$

Let us now calculate the reciprocity that player i perceives that players j and k give to him, respectively, $\tilde{k}_{ij}$ and $\tilde{k}_{ik}$. Using (13.52), $\tilde{k}_{ij} = k_{ji}(s_j, b_{ji}^1, b_{jk}^1) = \frac{1}{2} (s_j - 10)$. But player i does not know s_j , which is chosen simultaneously. Hence, he substitutes his first order belief about s_j , b_{ij}^1 , so, $\tilde{k}_{ij} = \frac{1}{2} (b_{ij}^1 - 10)$. In an analogous manner, $\tilde{k}_{ik} = \frac{1}{2} (b_{ik}^1 - 10)$. Substituting these intermediate calculations in (13.55) we get

$$\begin{aligned} u_i &= (20 - s_i) + \frac{1}{2} (s_i + s_j + s_k) + \frac{\lambda_i}{4} \left[(s_i - 10) (b_{ij}^1 - 10) + (s_i - 10) (b_{ik}^1 - 10) \right]. \\ &\Rightarrow u_i = \left[(20 - s_i) + \frac{1}{2} (s_i + s_j + s_k) \right] + \frac{\lambda_i}{4} \left[(s_i - 10) (b_{ij}^1 + b_{ik}^1 - 20) \right]. \end{aligned} \quad (13.57)$$

From (13.57), $\frac{\partial u_i}{\partial s_i} = -\frac{1}{2} + \frac{\lambda_i}{4} (b_{ij}^1 + b_{ik}^1 - 20)$. Thus, if $b_{ij}^1 + b_{ik}^1 - 20 \leq 0$, then u_i is maximized by choosing $s_i = 0$. On the other hand, if $b_{ij}^1 + b_{ik}^1 - 20 > 0$ and λ_i is high enough,

then it may be optimal to choose $s_i = 20$, otherwise if λ_i is low, then $s_i = 0$ continues to be optimal.

Since λ_i is unobservable, Dufwenberg et al. (2011) test a weaker prediction of the theory, namely, that either $b_{ij}^1 + b_{ik}^1 = 20$ and $s_i = 0$ or there is positive correlation between s_i and $b_{ij}^1 + b_{ik}^1$. They consider a 2×2 factorial design that varies the “label frame” and the “valence frame.” The label frame simply changes the labels of the game, while the valence frame potentially alters the reference point of the participants. Variations in the label frame are achieved through the treatments: Neutral and Community. In the Neutral treatment, the instructions speak of “the experiment,” while in the Community treatment, the instructions speak of “the community experiment.” Both frames are then interacted with two further treatments: Give and Take. The Give treatment is as described in the public goods experiment above. In the Take treatment, subjects are endowed with the resources and then allowed to take from it (so s_i has the interpretation of the amount that player i withdraws from the common resource). The monetary payoffs and the equilibria in the two games are identical.

Subjects first make their contributions in a one-shot game. Then they are asked to guess the average contributions of others $(b_{ij}^1 + b_{ik}^1)/2$ and their guesses of the sum of guesses of others (second order beliefs). Subjects are paid for the accuracy of their guesses. Here we give the results on first order beliefs only. Section 13.5.5 below, on guilt aversion, considers the results for second order beliefs. Figure 13.22 shows a plot of average levels of s_i versus average levels of $b_{ij}^1 + b_{ik}^1$ from various treatments (the dotted line is the 45° trend line). Contributions and first order beliefs are positively correlated, which is consistent with the presence of reciprocity in explaining the actions of players. An OLS regression confirms the statistical significance of first order beliefs in explaining contributions by the players.

As a second application of a SRE, consider the sequential prisoner’s dilemma game in Panel A of Figure 13.21. Since there are only two players, relabel the parameter λ_{ij} in (13.53) by λ_i . The results for this game are summarized in the next proposition.

Proposition 13.36 Consider the sequential prisoner’s dilemma game in Panel A of Figure 13.21. In a sequential reciprocity equilibrium:

- (i) If player 1 chooses D, then player 2 chooses d.
- (ii) If player 1 chooses C, then the following hold. If $\lambda_2 > 1$, player 2 chooses c and if $\lambda_2 < 0.5$, player 2 chooses d. If $0.5 < \lambda_2 < 1$, then player 2 chooses c with probability $p = \frac{2\lambda_2 - 1}{\lambda_2}$.
- (iii) If $\lambda_2 < 0.5$, then player 2 chooses d irrespective of the action of player 1.
- (iv) If $\lambda_2 > 1$, then there are three possible outcomes. (a) Irrespective of the magnitude of λ_1 , player 1 chooses C. (b) $\lambda_1 > 1$ and player 1 chooses D. (c) $\lambda_1 > 1$ and player 1 chooses C with probability $q = \frac{\lambda_1 - 1}{2\lambda_1}$.

Proof: We only prove parts (i) and (ii). The proofs of parts (iii), (iv) are analogous, so they are omitted.⁶¹

(i) The choice of D by player 1 gives to player 2 a lower payoff for every strategy that 2 could choose. Hence, such a choice by player 1 is interpreted as unkind by player 2. Player 2 must then choose d on both counts—such a choice maximizes the material payoffs and improves the psychological payoffs by reciprocating the unkindness of player 1.

⁶¹ Readers interested in the proof may consult Dufwenberg and Kirchsteiger (2005, pp. 293–5).




Figure 13.22 Correlation between contributions and first order beliefs in a public goods game.

Source: Reprinted from *Games and Economic Behavior* 73: 459–78. Dufwenberg, M., Gächter, S., and Hennig-Schmidt, H., “The framing of games and the psychology of play.” Copyright 2011, with permission from Elsevier.

(ii) Suppose that player 2 wishes to calculate his kindness to player 1 when player 1 chooses C . The equitable payoff of player 1 in this case is $\frac{1}{2}(1 + (-1)) = 0$. Thus, if 2 plays c , then his kindness to 1 is $1 - 0 = 1 > 0$ and if he plays d , then his kindness is $-1 - 0 = -1 < 0$. Let us now calculate $\hat{k}_{21}$, given in (13.51), i.e., player 2's perception of 1's kindness towards him. b_{21}^2 gives the second order belief of player 2 about player 1's first order belief about the probability with which player 2 will play c . Hence, player 2 knows that the payoff that player 1 intended to give him by playing C is $b_{21}^2 + (1 - b_{21}^2)2$. Recall the definition of the equitable payoff in (13.48), and from part (i) the fact that if 1 plays D , player 2 responds with d . Hence, the equitable payoff for player 2 is $\frac{1}{2}(b_{21}^2 + (1 - b_{21}^2)2 + 0)$. Thus, the kindness of player 1 towards 2 as perceived by player 2 (the analog of (13.51)) is $b_{21}^2 + (1 - b_{21}^2)2 - \frac{1}{2}(b_{21}^2 + (1 - b_{21}^2)2 + 0) = 1 - \frac{b_{21}^2}{2}$.

Using the definition of the utility function in (13.53), if 1 plays C , player 2's utility from playing c is $1 + \lambda_2 \times 1 \times \left(1 - \frac{b_{21}^2}{2}\right)$. But when 2 plays d his kindness (see above) is -1 , hence, his payoff from d is $2 + \lambda_2 \times (-1) \times \left(1 - \frac{b_{21}^2}{2}\right)$. Thus player 2 prefers to play c rather than d if $\lambda_2(2 - b_{21}^2) > 1$. In equilibrium (Definition 13.19), the second order beliefs of players must equal the first order beliefs, which, in turn, are consistent with the actual choices of strategies. Thus, if c is to be an equilibrium following the choice of C by player 1, then $b_{21}^2 = 1$. In this case $\lambda_2(2 - b_{21}^2) > 1$

implies $\lambda_2 > 1$. However, if 2 plays d following the play of C , then it must be that $\lambda_2 (2 - b_{21}^2) < 1$, which implies that $\lambda_2 < 0.5$. For intermediate values, $0.5 < \lambda_2 < 1$, player 2 cannot play c or d as part of a SRE, which suggests that there is “randomized choice.”⁶² Indifference between the payoffs of player 2 from playing c and d give $\lambda_2 (2 - b_{21}^2) = 1$. The second order beliefs must hold in equilibrium, hence, b_{21}^2 must equal the probability, p , with which 2 plays c which gives $p = \frac{2\lambda_2 - 1}{\lambda_2}$. ■

13.5.4 Fairness, reciprocity, and inequity-aversion in extensive form games

We now extend the model in Section 13.5.3 to allow for a combination of *intention* and *outcome* based reciprocity. The resulting model can be fairly complex. We now give a simplified and intuitive treatment of the Falk and Fischbacher (2006) model. The empirical evidence in Falk et al. (2008) supports the predictions in Falk and Fischbacher (2006) over those in Dufwenberg and Kirchsteiger (2004).

Consider a finite, two-player, extensive form game of complete and perfect information.⁶³ Suppose that T is a finite set of nodes, and the set of terminal nodes is Z . Let the set of nodes where player $i = 1, 2$ have the move denoted by T_i with generic element denoted by t_i . Let $A(t_i)$ be the set of actions available to player i at node t_i . We list the strategy of player 1 first, followed by that of player 2.

The material payoffs of the players are given by $\pi_i : Z \rightarrow \mathbb{R}$. Let $\Delta A(t_i)$ be the set of probability distributions over the set of actions of player i at node t_i . The behavior strategies for player $i = 1, 2$ belong to the set $\Phi_i = \Pi_{t_i \in T_i} \Delta A(t_i)$, i.e., probability distributions over each of player i 's decision nodes. Let a generic element of Φ_i be denoted by ϕ_i . Suppose that the strategies chosen by the players are $\phi_i \in \Phi_i$, $i = 1, 2$. Then, the expected payoffs of the two players are given, respectively, by $\pi_1(\phi_1, \phi_2)$, $\pi_2(\phi_1, \phi_2)$. Denote the expected payoff of player i , conditional on being at node t_i by $\pi_i(\phi_1, \phi_2 | t_i)$. The first order beliefs of player i about the behavior strategy of player j , are denoted by b_i^1 . Second order beliefs of player i , are denoted by b_i^2 , and reflect beliefs about b_j^1 , $j \neq i$, where $b_i^2 \in \Phi_i$ and $b_i^1 \in \Phi_j$. These beliefs are node-dependent, i.e., we allow for the updating of beliefs as we progress through the game.

At any node, the kindness term, $\tilde{k}_j$, measures the perceived kindness of player i that is received from player $j \neq i$.⁶⁴ It depends on (1) *outcomes*, measured by an *outcome term* η_j , and (2) *intentions*, measured by an *intention factor* $\vartheta_j \in [0, 1]$, where $\vartheta_j = 1$ denotes full intent of player j in inducing a given outcome term for player i . Values of $\vartheta_j < 1$ denote less than full intent and lower values of ϑ_j denote lower intent.

Definition 13.20 *The kindness perceived by player i from player j at node t , is given by*

$$\tilde{k}_j(b_i^2, b_i^1 | t) = \vartheta_j(b_i^2, b_i^1 | t) \eta_j(b_i^2, b_i^1 | t), \quad (13.58)$$

where b_i^1 and b_i^2 are first and second order beliefs of player i at node t .

⁶² Recall our comment above that players in this model always choose pure strategies. Mixed strategies here have the interpretation of randomized choices in which the proportions of the population playing the different pure strategies are such that the opponent appears to be playing a mixed strategy that mirrors the population proportions playing the pure strategies.

⁶³ For the extension to finite n -player extensive form games, see the Appendix in Falk and Fischbacher (2006).

⁶⁴ We omit double subscripts on the kindness term because there are only two players and to differentiate the kindness function in this model, which is different to those in the previous sections.

Since ϑ_j is non-negative, the sign of the kindness term is determined entirely by the sign of the outcome term. The outcome term, η_j , at node $t \in T$ is given by

$$\eta_j(b_i^2, b_i^1 | t) = \pi_i(b_i^2, b_i^1 | t) - \pi_j(b_i^2, b_i^1 | t). \quad (13.59)$$

From Definition 13.20, when player i wishes to compute the kindness of player j toward him, the relevant beliefs are his first and second order beliefs, b_i^2, b_i^1 , as in the models of Rabin and Dufwenberg–Kirchsteiger. The payoff of player i , at node t , given the first and second order beliefs, is given by $\pi_i(b_i^2, b_i^1 | t)$ and the payoff of the opponent, $j \neq i$, is $\pi_j(b_i^2, b_i^1 | t)$. Player i takes as his reference payoff, the payoff of the opponent. In the Rabin and Dufwenberg–Kirchsteiger models, the relevant analogue was some notion of an equitable payoff. This is the sense in which the Falk–Fischbacher model incorporates inequity aversion. Outcomes are advantageous for player i when $\eta_j > 0$ and disadvantageous when $\eta_j < 0$. Thus, player i perceives player j to be kind if $\eta_j > 0$ and unkind if $\eta_j < 0$.

The outcome terms, η_1 and η_2 , do not provide information on the kindness or unkindness of intentions. For instance, they do not reveal if the other player had access to a fairer alternative. We now turn to beliefs about intentions.

Suppose that player i wishes to infer the intentions of player j . First, he would like to infer the set of payoffs, Π_i , that are available to both players. But these payoffs depend on the strategy chosen by player j , $\phi_j \in \Phi_j$, and the beliefs of player j about the strategy of player i , b_j^1 . Since b_j^1 is unobserved by player i he uses his beliefs about b_j^1 , i.e., $b_i^2 \in \Phi_i$, where Φ_i is the set of behavior strategies of player i . Denote by $\Pi_i(b_i^2, \phi_j | t)$, the relevant set of payoffs for both players (from the perspective of player i) at node $t \in T$.

$$\Pi_i(b_i^2, \phi_j | t) = \{\pi_i(b_i^2, \phi_j | t), \pi_j(b_i^2, \phi_j | t) | \phi_j \in \Phi_j\}. \quad (13.60)$$

Suppose that player j induces a payoff distribution $(\pi_i, \pi_j) \in \Pi_i(b_i^2, \phi_j | t)$ by a choice of his action $\phi_j \in \Phi_j$. How intentional is the payoff distribution (π_i, π_j) chosen by player j ? To answer this question, we define the *intention factor* $\vartheta_j(b_i^2, b_i^1 | t)$. Recall that if $\vartheta_j = 1$, then the choice is fully intentional. For less than fully intentional choices, ϑ_j is less than 1. Formally, letting $\epsilon_i \in [0, 1]$ be a *pure outcome concern parameter*,

$$\vartheta_j(b_i^2, b_i^1 | t) = \begin{cases} 1 & \text{if } \pi_i \geq \pi_j \text{ and } \exists \tilde{\pi}_i \in \Pi_i(b_i^2, \phi_j | t) : \tilde{\pi}_i < \pi_i, \\ \epsilon_i & \text{if } \pi_i \geq \pi_j \text{ and } \forall \tilde{\pi}_i \in \Pi_i(b_i^2, \phi_j | t) : \tilde{\pi}_i \geq \pi_i, \\ 1 & \text{if } \pi_i < \pi_j \text{ and } \exists \tilde{\pi}_i \in \Pi_i(b_i^2, \phi_j | t) : \tilde{\pi}_i > \pi_i, \\ \epsilon_i & \text{if } \pi_i < \pi_j \text{ and } \forall \tilde{\pi}_i \in \Pi_i(b_i^2, \phi_j | t) : \tilde{\pi}_i \leq \pi_i. \end{cases} \quad (13.61)$$

The interpretation of (13.61) is as follows. In the first two rows, player j treats player i kindly by the outcome criterion, i.e., $\pi_i \geq \pi_j$. However, the action of player j bears a different interpretation depending on the alternatives available. In the first row, an alternative exists that gives a worse payoff to player i ($\tilde{\pi}_i < \pi_i$). Thus, player j could have been unkind but chose to be intentionally kind, so $\vartheta_j = 1$. In the second row, player j induces a greater payoff for player i but player j had no alternative that would have given player i a lower payoff. Hence, the kindness of player j is untested in this case. So player i assigns a lower value to the intentional kindness factor, i.e., $\vartheta_j = \epsilon_i \leq 1$.

In the last two rows of (13.61), player i is at a disadvantage relative to player j (i.e., $\pi_i < \pi_j$). So is this unkindness by player j to player i , intentional? In the third row, player j could have improved

player i 's payoff in the alternative ($\tilde{\pi}_i > \pi_i$) but choose not to, hence, the unkindness of player j is fully intentional, so $\vartheta_j = 1$. By contrast, in row 4, player i is at a disadvantage but none of the alternatives can give player i a higher payoff, hence, player i cannot infer fully if player j wanted to treat him in an unkind manner, so $\vartheta_j = \epsilon_i \leq 1$.

Suppose $\epsilon_i = 1$. Then, from (13.61), ϑ_j always equals 1. Hence, from the kindness function (13.58) we get that $\tilde{k}_j(b_i^2, b_i^1 | t) = \eta_j(b_i^2, b_i^1 | t)$, thus, the players care only about the outcome, i.e., whether it is advantageous ($\eta_j > 0$) or disadvantageous ($\eta_j < 0$). On the other hand, if $\epsilon_i = 0$, then players assign weight only to fully intentional outcomes (rows 1 and 3 in (13.61)).

Example 13.17 Suppose that

$$\Pi_i = \{(10, 3), (8, 10), (6, 5)\},$$

where Π_i is defined in (13.60); the first number in any payoff pair represents the payoff of player i and the second number, the payoff of player j . If player j chooses the outcome (10, 3), then the first row of (13.61) applies, so $\vartheta_j = 1$. If player j chooses the outcome (8, 10), then the third row of (13.61) applies, so $\vartheta_j = 1$. Finally, if player j chooses the outcome (6, 5), then the second row of (13.61) applies, so $\vartheta_j = \epsilon_i$.

Definition 13.21 Pick an arbitrary node t . Let $z \in Z$ be a terminal node that follows node t . Denote by $\lambda(t, z)$ the unique node that is the immediate successor of node t on the path to node z .

Definition 13.22 (Reciprocation term r_i): The reciprocation term of player $i = 1, 2$ at node t is given by

$$r_i(t, z, b_i^2, b_i^1) = \pi_j(b_i^2, b_i^1 | \lambda(t, z)) - \pi_j(b_i^2, b_i^1 | t), j \neq i. \quad (13.62)$$

The reciprocation term of player i , given in (13.62), shows how player i alters the payoff of player j when he has the move at node $\lambda(t, z)$. For any given $\pi_j(b_i^2, b_i^1 | t)$, player i can choose to either reward (positive reciprocation, $r_i > 0$) or punish (negative reciprocation, $r_i < 0$) player j by choosing the size of $\pi_j(b_i^2, b_i^1 | \lambda(t, z))$ through his own actions.

The utility function in the psychological game for player $i = 1, 2$ is defined as

$$u_i(z, b_i^2, b_i^1) = \pi_i(z) + \mu_i \sum_{t'=t}^z \tilde{k}_j(b_i^2, b_i^1 | t') r_i(t', z, b_i^2, b_i^1). \quad (13.63)$$

Since $\pi_i : Z \rightarrow \mathbb{R}$, the material payoffs of any player are defined at each of the end nodes only. $\mu_i \geq 0$ is the *reciprocity parameter*; it is the relative weight placed by a player on the *psychological payoff* (second term), relative to the material payoffs (first term). While the material payoffs only depend on reaching a particular end node, the psychological payoffs depend on play along the entire path leading up to the end node. At each node where player i has a move, he computes the kindness of player j (the term $\tilde{k}_j$) as well as his own reciprocation to the kindness of j (the term r_i). So when player j is being kind ($\tilde{k}_j > 0$), positive reciprocation by player i ($r_i > 0$) improves the psychological payoff of player i .

Falk and Fischbacher (2006) call a subgame perfect psychological Nash equilibrium a *reciprocity equilibrium*. Heuristically, for any subgame that starts with some node $t \in T$, we can use (13.63) to compute the possible utilities of the players that players maximize. This allows us to determine the subgame perfect equilibria, which with the additional restriction of the consistency of actual actions with first and second order beliefs, constitutes a *reciprocity equilibrium*. For different parameter values, Falk and Fischbacher show that their model is consistent with

results from several kinds of experimental games such as the ultimatum game, the gift exchange game, the dictator game, the public goods game, and the prisoner's dilemma game.

As an application of this framework, consider an ultimatum game. Suppose that the proposer, P , who is given one dollar can offer an amount $0 \leq c \leq 1$ to a responder, R . Suppose that, conditional on a choice of c , R accepts the offer with some probability $p(c) \geq 0$. If R accepts, the payoffs of P and R are $1 - c$ and c respectively. Otherwise, each player gets zero each.

In any subgame where the history is c , let P 's first order beliefs about p be given by $b_P^1 = p'$ and R 's second order belief about p , or beliefs about b_P^1 , be given by $b_R^2 = p''$. Let the intention factor of P be $\vartheta_P(c)$. Using (13.63), we can write R 's utility from accepting P 's offer, U_R^a , as.

$$U_R^a = c + \mu_R \vartheta_P(c) p'' [c - (1 - c)] [(1 - c) - p'' (1 - c)]. \quad (13.64)$$

If R accepts, her material payoff is c (first term). The second term is the psychological payoff $\mu_R \tilde{k}_{PR}$. Using (13.62), the reciprocation term is $r_R = (1 - c) - p'' (1 - c)$; this is the difference between P 's payoff at the terminal and the penultimate node (hence, the usage of the probability p''). Using (13.59), $\eta_P = p'' [c - (1 - c)]$ (notice that before R makes his decision we compute the expected payoffs from R 's actions). In an analogous manner, we can derive the payoff of R when she rejects the offer, U_R^r .

$$U_R^r = 0 + \mu_R \vartheta_P(c) p'' [c - (1 - c)] [0 - p'' (1 - c)]. \quad (13.65)$$

Using (13.64), (13.65), the difference in the payoffs from acceptance and rejection is

$$U_R^a - U_R^r = c + \mu_R \vartheta_P(c) p'' (1 - c) (2c - 1). \quad (13.66)$$

A sufficient condition for acceptance to dominate rejection is that the second term on the RHS of (13.66) be positive $\Leftrightarrow c \geq 0.5$. Hence, in this case R always accepts, i.e., $p = 1$; we summarize this result in the next lemma.

Lemma 13.2 *All offers by P such that $c \geq 0.5$ are immediately accepted by R with $p = 1$.*

The next proposition gives the solution to the optimal probability, p , when $c < 0.5$. The proof is left as an exercise to the reader.

Proposition 13.37 *If $\mu_R > 0$, then in the unique reciprocity equilibrium, the acceptance probability of R is given by*

$$p^*(c) = \begin{cases} \min \left\{ \frac{c}{\mu_R(1-c)(1-2c)}, 1 \right\} & \text{if } c < 0.5 \\ 1 & \text{if } c \geq 0.5 \end{cases} \quad (13.67)$$

Proposition 13.37 is illustrated in a diagram in Figure 13.23. It is clear that the acceptance probability is increasing in the offer of P . For a given offer, the probability of acceptance is higher if the weight placed by R on the psychological payoff, μ_R is lower. In the limit as $\mu_R \rightarrow 0$ one gets the acceptance probability to be one, for all positive offers (as in classical game theory).

In the spirit of subgame perfection, given the acceptance probability schedule of R in (13.67), P now chooses the optimal value of c . From (13.67), we can find the smallest offer, c_0 , such that $p(c_0) = 1$; this is shown in Figure 13.23. A simple calculation shows that




Figure 13.23 The effect of c and μ_R on p .

$$c_0 = \frac{1 + 3\mu_R - \sqrt{1 + 6\mu_R + \mu_R^2}}{4\mu_R}. \quad (13.68)$$

Using (13.63), and the consistency condition $p = p' = p''$, we can write the utility of P . Suppose that the initial beliefs of R are that P is going to make some offer $\tilde{c}$. Let c be the actual offer made and observed by R at the node where she makes a decision to accept or reject. In equilibrium, $\tilde{c} = c$. Then P 's utility is given by:

$$U_P = p(c)(1 - c) + \mu_P \vartheta_R p(\tilde{c}) [1 - \tilde{c} - \tilde{c}] [p(c)c - p(\tilde{c})\tilde{c}]. \quad (13.69)$$

Notice that in writing (13.69) we have used (13.62), so that $r_P = [p(c)c - p(\tilde{c})\tilde{c}]$ and from (13.59) we get that $\eta_P = p(\tilde{c}) [1 - \tilde{c} - \tilde{c}]$. We leave the proof of the next proposition as an exercise for the reader.

Proposition 13.38 *The optimal offer by P is given by*

$$c^* = \max \left\{ c_0 = \frac{1 + 3\mu_R - \sqrt{1 + 6\mu_R + \mu_R^2}}{4\mu_R}, 0.5(1 - \mu_P^{-1}) \right\}. \quad (13.70)$$

In (13.70), the first term in the brackets on the RHS depends on μ_R , the weight placed on the psychological payoff by R . Falk and Fischbacher term this as the *reciprocal inclination* of R . The first term is the smallest offer that guarantees acceptance. The responder's reciprocal inclination is crucial for the size of this term. If R cares purely about her material payoffs, i.e., $\mu_R = 0$, then $c_0 = 0$. In the limit as $\mu_R \rightarrow \infty$, $c_0 = 0.5$. In the intermediate ranges of μ_R , c_0 is increasing in μ_R . If, on the other hand, R cares only about material payoffs (and so $c_0 = 0$), then the offer of P depends on the second term in the braces on the RHS of (13.70). This term depends on the concern of P for an equitable outcome, μ_P . If $\mu_P = 0$, then the offer is zero. In general, the second term is increasing in μ_P and as $\mu_P \rightarrow \infty$, then the second term tends to 0.5. Thus, the optimal offer is less than 0.5 in both cases.

13.5.5 Guilt in games

Battigalli and Dufwenberg (2007) consider a general framework for modeling *guilt*. Player i may experience guilt if he lets down player j , based on his second order beliefs (i.e., his beliefs about

the beliefs of player j). Consider Example 13.8 above. Let John and the taxi driver be respectively players 1 and 2. John's utility, u_1 , is given by

$$u_1(t) = y_1 - t - \gamma \max \{0, b_1^2 - t\}, \quad (13.71)$$

where y_1 is John's income, t is the tip that he pays to the taxi driver, and b_1^2 is player 1's beliefs about b_2^1 (i.e., John's belief about the taxi-driver's expectation of the tip). $\gamma \geq 0$ is a measure of the importance of the psychological cost of guilt.⁶⁵ The utility of player 1 from three different levels of tips is given below.

$$\begin{cases} u_1(t) = y_1 - b_1^2 & \text{if } t = b_1^2, \\ u_1(\hat{t}) = y_1 - \hat{t} & \text{if } t = \hat{t} > b_1^2, \\ u_1(\tilde{t}) = y_1 - b_1^2 + (1 - \gamma)(b_1^2 - \tilde{t}) & \text{if } t = \tilde{t} < b_1^2. \end{cases}$$

Since $\hat{t} > b_1^2$, we have $u_1(\hat{t}) < u_1(t = b_1^2)$, thus, John never pays a tip that exceeds his second order beliefs. However, depending on the value of γ we have the three cases $u_1(t = b_1^2) \begin{smallmatrix} \geq \\ \leq \\ = \end{smallmatrix} u_1(\tilde{t})$. If psychological payoffs are salient enough, in the sense $1 < \gamma$, then $u_1(t = b_1^2) > u_1(\tilde{t})$, so it is utility enhancing to pay exactly the tip that John thinks that the taxi driver expects. In other words if the emotion of guilt is strong enough, John always pays a tip that he believes is expected by the taxi driver.

Example 13.16 described contributions in a public goods game in the presence of *reciprocity*. Let us now invoke the role of *guilt* to explain contributions. Here we focus on *simple guilt*, which requires players to experience guilt if they believe that they have let down the expectations of others; see Battigalli and Dufwenberg (2007). For the case of the three-player public goods game in Example 13.16, the objective function of a player $i \in N$ who experiences such feelings of guilt is

$$U_i = \left[(20 - s_i) + \frac{1}{2} (s_1 + s_2 + s_3) \right] - \gamma_i \max \left(0, \frac{b_{ji}^2 + b_{ki}^2}{2} - s_i \right), \quad i \neq j \neq k. \quad (13.72)$$

Thus, players experience guilt if they think that their actual contribution to the public good, s_i , is lower than the average expectations of others about their contributions, $\frac{b_{ji}^1 + b_{ki}^1}{2}$; where b_{ji}^1 is the first order belief of player j about the contribution of player i . Since players do not know what others expect of them, they use their second order beliefs to form beliefs about the beliefs of others, i.e., player i uses $\frac{b_{ji}^2 + b_{ki}^2}{2}$ to guess the value of $\frac{b_{ji}^1 + b_{ki}^1}{2}$; where b_{ji}^2 is the belief of player i about the belief of player j about the contribution of player i . The parameter $\gamma_i > 0$ measures the degree of *guilt aversion* of the player. If $\gamma_i = 0$, then there is no guilt and the player optimally chooses to free ride, $s_i = 0$, as in classical game theory.

Suppose now that $\gamma_i > 0$ and we begin at $s_i = 0$. From the first term in (13.72), the marginal utility of an extra unit of s_i is $-\frac{1}{2}$, and, from the second, the marginal guilt aversion falls by γ_i . Hence, if $\gamma_i < \frac{1}{2}$, then it is still optimal to choose $s_i = 0$. Thus, low levels of guilt aversion do not alter the standard result. Now suppose that $\gamma_i > \frac{1}{2}$. If $\frac{b_{ji}^2 + b_{ki}^2}{2} > s_i$, then the marginal utility of

⁶⁵ For pedagogical simplicity, the utility from guilt is linear in (13.71). However, one needs further empirical evidence to distinguish between alternative functional forms; see Attanasi et al. (2015).

the player can be improved by increasing s_i and if $\frac{b_{iji}^2 + b_{iki}^2}{2} < s_i$, then the marginal utility can be increased by reducing s_i . Hence, if $\gamma_i > \frac{1}{2}$, then it is optimal to choose $s_i = \frac{b_{iji}^2 + b_{iki}^2}{2}$.

Dufwenberg et al. (2011) run experiments to test the predictions of this model. They do not observe γ_i but were able to isolate subjects for whom $s_i = 0$, $\frac{b_{iji}^2 + b_{iki}^2}{2} = 0$ and $\gamma_i < \frac{1}{2}$. For the remaining subjects, they choose to test the weaker hypothesis that the correlation between s_i and $\frac{b_{iji}^2 + b_{iki}^2}{2}$ is positive. The results are shown in Figure 13.24. The diagonal is drawn as a dotted line and the second line in the diagram is the best fitting line between pairs of contributions and second order beliefs. The results are generally consistent with the guilt-aversion model. There is a positive relation between contributions and second order beliefs in all treatments. The filled-in circles show zero contributions when $\frac{b_{iji}^2 + b_{iki}^2}{2} = 0$ and represent the case $s_i = 0$, $\frac{b_{iji}^2 + b_{iki}^2}{2} = 0$ and $\gamma_i < \frac{1}{2}$.

The frames, Give or Take, and Neutral or Community (see description in Example 13.16), have a significant affect on the second order beliefs. Mean contributions are highest in the Give-Neutral treatment, followed by the Give-Community treatment, Take-Neutral treatment and the Take-Community treatment. There are also important differences in the elicited first order beliefs (see Figure 13.22) and second order beliefs (see Figure 13.24) in different treatments. Hence, the framing of games has an important effect on beliefs.





Figure 13.24 The relation between contributions and second order beliefs in the public goods game.

Source: Reprinted from *Games and Economic Behavior* 73: 459–78. Dufwenberg, M., Gächter, S., and Hennig-Schmidt, H. "The framing of games and the psychology of play." Copyright © 2011, with permission from Elsevier.



Panel A: Partnership game in the absence of guilt

Panel B: Partnership game in the presence of guilt

Figure 13.25 A partnership game in the absence and presence of guilt.

Source: Charness and Dufwenberg (2006) with permission of The Econometric Society.

So which of the two explanations, *reciprocity* (see Example 13.16) and *guilt*, explains the level of contributions in the public goods game. First and second order beliefs are highly correlated in each treatment. When a regression is run with contributions as the dependent variable and first and second order beliefs as the independent variables, except for the Take–Community treatment, second order beliefs lose significance. In the Give treatment, the first order beliefs are highly significant, as predicted by the reciprocity hypothesis. On balance, second order beliefs, and, so, guilt aversion is relatively more important in the Take–treatments. However, the high correlation between first and second order beliefs suggests that reciprocity and guilt cannot be neatly separated in the experimental results.

The first study of the relation between beliefs and observed behavior is due to Dufwenberg and Gneezy (2000). In their lost wallet game, player 1 has the opportunity to take money from a lost wallet or pass this decision to player 2 who then decides on an allocation of the money in the wallet between the two players. They find that player 2 passed a higher amount of money to player 1 if he believed that player 1 expects a higher level of transfer. Guerra and Zizzo (2004), Charness and Dufwenberg (2006), and Bacharach et al. (2007) elicit second order beliefs of players. They find that there is a close association between the elicited beliefs and the tendency of players to behave prosocially, a finding that is consistent with the guilt-aversion hypothesis.⁶⁶

Here we consider the work of Charness and Dufwenberg (2006) in more detail. It involves issues of guilt, communication, and promises that tie in nicely with the work on coordination and communication that we extensively analyze in Chapter 12 and other work on promises and ethics that we report in Part 2 of the book.

Consider the game in panel A in Figure 13.25 that is designed to mimic a situation of moral hazard in a principal–agent relation. There are two risk-neutral players who are making a decision to enter into a partnership. Player 1 (principal) and player 2 (agent) each have an outside option of 5. Player 1 can choose not to enter into the partnership by playing the strategy Out, leading to the outside option payoffs (5, 5). Or player 1 can choose to enter into a partnership (the strategy In). If a partnership is entered into, then players sign a contract that specifies a wage–effort tuple

⁶⁶ For a critical review of this evidence, see Attanasi and Nagel (2008).

$(w, e) = (14, 1)$. The wage, w , is contractible but the effort level of player 2 is not observable or verifiable to player 1.

Player 2 can choose an effort level e at a cost $4e$. When player 2 chooses to shirk, $e = 0$ (the strategy Don't Roll), the project fails in the sense that it yields a revenue of 14, so player 1 gets 0 and player 2 gets the promised wage of 14. When player 2 chooses the contractual effort level $e = 1$ (the strategy Roll), there is a cost of effort equal to 4, and the project succeeds with probability $p = \frac{5}{6}$. If the project fails, the total revenue is 14, of which player 2 gets $14 - 4 = 10$, leaving 0 for player 1. In the case of success, the project yields an extra revenue of 12 for a total revenue of 26 (player 2 gets $14 - 4$ and player 1 gets $26 - 14$). Importantly, in the event that the project fails, player 1 cannot tell if it was due to a low effort level on the part of player 2 or high effort but bad luck. Hence, the terminal nodes $(0, 14)$ and $(0, 10)$ are in the same information set for player 1. This is the sense in which there is a moral hazard problem.

Using classical game theory, the unique subgame perfect outcome in this game with risk-neutral players is (Out, Don't) but it is not efficient. Can communication in the form of free form messages from player 2 to player 1, prior to the start of the game, achieve an efficient outcome? Clearly if the players have completely self-regarding preferences (as reflected in the payoffs in panel A), then such communication will be ineffective.

Now introduce the possibility of guilt aversion; see Panel B in Figure 13.25. Suppose that player 2 plays Roll with some probability p . The first order beliefs of player 1 about p are denoted by b_1^1 . The beliefs of player 2 about b_1^1 are denoted by b_2^2 (these are second order beliefs about p). Thus, player 2 believes that player 1 hopes to get the expected payoff: $(1 - b_2^2)0 + b_2^2(\frac{5}{6}12 + \frac{1}{6}0) = 10b_2^2$. Letting the guilt-aversion parameter of player 2 be γ_2 , the guilt-incorporated payoffs of the players are shown in panel B. In particular, the utility of player 2 from Don't Roll is $14 - 10\gamma_2 b_2^2$.

The risk-neutral player 2 chooses Roll, i.e., prefers putting in the contractual effort level $e = 1$ to shirking, $e = 0$, if $14 - 10\gamma_2 b_2^2 < 10$. The higher the second order belief of player 2, b_2^2 , the more likely is this condition to be satisfied and the more likely is player 2 to fulfil the contractual obligations.

In the experiments, prior to the move by player 1, player 2 can send a non-binding promise to player 1 (written as a free-form message) about his intentions, e.g., that he will play the strategy Roll with high probability (or even certainty). Such promises may increase the first order belief of player 1, b_1^1 , as well as the second order belief of player 2, b_2^2 . Thus, promises may, in fact, optimally increase the probability that player 2 plays Roll. Hence, ex-ante communication may be effective in promoting partnerships and achieving the efficient outcome (In, Roll).

As a robustness check, players also play a second game, which is identical to panel-A of Figure 13.25 except that the outside option of both players is 7 each. This is the $(7, 7)$ treatment, while the original one is termed the $(5, 5)$ treatment because if player 1 chooses Out at the initial node, the payoffs are $(5, 5)$. Beliefs of the players, b_1^1 and b_2^2 , are directly elicited at the end of the experiment. Subjects in the role of player 1 were asked about the percentage of player 2s who are likely to choose Roll; the guesses were then rewarded for accuracy. Subjects in the role of player 2s were asked to guess the average guess of player 1s who choose In; these guesses were also rewarded for accuracy.

The results are shown in Figure 13.26. The results in the $(5, 5)$ and the $(7, 7)$ treatments are similar. Consider the $(5, 5)$ treatment. Communication enhances the likelihood of the following three events. (1) Player 1's choice of In (from 56% to 74.5%). (2) Player 2's choice of Roll (from 44% to 67%). (3) The overall play of the strategy (In, Roll) (from 20% to 50%). There is a strong positive correlation between the first order beliefs of player 1, b_1^1 , and the choice of In and the second order beliefs of player 2, b_2^2 , and Roll; this is consistent with the guilt-aversion hypothesis.




Figure 13.26 The effect of communication/promises on partnerships.

Source: Charness and Dufwenberg (2006) with permission of The Econometric Society.

Messages from player 2 to player 1 also significantly increase the first order beliefs of player 1s, b_1^1 (from 41.01% without communication to 59.02% with communication). Since the messages were free-form, messages that contained a statement of intent or promise were particularly effective in raising the level of beliefs, b_1^1 , b_2^2 , and the partnership enhancing actions of the two players, In and Roll. An exercise asks you to consider if these results can be explained by other models.

Ellingsen et al. (2010) implement a more direct method of eliciting beliefs; they find the absence of the guilt-aversion motive in their experiments. This is a potentially challenging finding for theories of guilt aversion that we now describe in more detail, followed by an explanation of these results based on Khlmetzski et al. (2015). The authors consider three games. A dictator game, a trust game, and the partnership game of Charness and Dufwenberg (2006) described in panel A of Figure 13.25. In each case, a player in the game (the dictator in the dictator game, the trustee in the trust game, and player 2 in the partnership game) is provided with *direct information* on the first order beliefs of the other player, before they make their decision.

Furthermore, players are not told that their first order beliefs will be communicated to others, hence, these beliefs are not misstated for strategic reasons. If players knew that their first order beliefs would be communicated to the other player then they could have tried to induce greater guilt on the part of the other player to get a selfishly better payoff. Thus, the second order beliefs of players (beliefs about the first order beliefs of others) are as accurate as possible. One might wonder if this violates the convention in economic experiments that the experimenter should not mislead the subjects. Ellingsen et al. (2010) give a robust defence of their procedure against such a charge.⁶⁷

In the dictator game, dictators were asked to split an amount of 20 Swedish Kroner (SEK) between themselves and a passive responder. Prior to the game, the responders' beliefs about the

⁶⁷ They did not lie to the subjects—rather they simply did not give any information to subjects about what would be done with their first order beliefs. In the exit interviews, none of the subjects complained about being cheated, once the information about the use of first order beliefs was conveyed to them. Obviously there might be externalities for other experimenters if the same subjects participate in other experiments. The authors assign very little probability to such an event.




Figure 13.27 Bubble plots of the relation between induced second order beliefs and contributions in the dictator game (left panel) and the amount sent back by trustees in the trust game (right panel).

Source: Reprinted from *Games and Economic Behavior* 68: 95–107. Ellingsen, T., Johannesson, M., Tjøtta, S., and Torsvik, G., “Testing guilt aversion.” Copyright © 2010, with permission from Elsevier.

likely donations of the dictators were elicited. These beliefs were then conveyed to the dictators before they made their decision. Let us term these as *induced second order beliefs*. If the guilt-aversion hypothesis holds, then one should expect a positive correlation between the amounts offered by the dictator to the receiver and the induced second order beliefs. The results are shown in the left panel of Figure 13.27 in a bubble plot. The size of the bubbles reflects the frequency of a particular combination of beliefs and contributions. The bubbles appear fairly random, and the absence of any statistically significant correlation is confirmed by a calculation of the parametric Pearson correlation coefficient (-0.075) and the non-parametric Spearman correlation coefficient (-0.044). A sensitivity analysis is conducted by removing dictators whose choices may appear unexpected, e.g., offering more than 50%, and those who may not suffer any guilt aversion, i.e., offer nothing. Yet, the correlation is insignificant.

In the trust game, the trustor can send 50 Norwegian Kroner (NOK) to the trustee. The experimenter multiplies the amount sent by 5. The trustee then sends back to the trustor an amount from the set $\{0, 50, 100, 150, 200, 250\}$. After the trustors have made their decisions, their first order beliefs about the modal amount that the trustees are likely to send back is elicited. Trustors are rewarded with NOK 50 if their guess is correct. Before the trustees make their decisions, the first order beliefs of the trustors are conveyed to them—these then become the induced second order beliefs of the trustees. Under the guilt-aversion hypothesis, we expect to see a positive correlation between the induced second order beliefs and the amount sent back by the trustee to the trustor. The results are shown in the right panel of Figure 13.27 in a bubble plot. The distribution of bubbles is almost random, and the non-parametric and parametric correlation coefficients are insignificant. These results are not consistent with the guilt-aversion hypothesis.

The third game, the partnership game (see Figure 13.25), allows a direct comparison with the results of Charness and Dufwenberg (2006). Suppose that player 2 chooses Roll with probability p . Player 1 is asked to guess p and this first order belief is conveyed to player 2 as an induced second order belief before player 2 makes his decision. Player 1 earns a reward of NOK 50 if his guess is within 5 percentage points of the actual p . Under the guilt-aversion hypothesis, we should expect a positive correlation between the induced second order beliefs and p . Yet the correlation is negative, statistically insignificant, and negligible. These results are in contrast to the Charness

and Dufwenberg (2006) results who find support for the guilt-aversion hypothesis with a different experimental design in which second order beliefs were elicited rather than induced in this manner.

The definition of guilt is one-sided in (13.71). In the context of Example 13.8, John experiences *guilt* if his tip is below his perceived expectation of the tip by the taxi driver. However, John experiences no *joy* or *elation* from the surprise of offering a tip that exceeds the expectation of the taxi driver. Khalmetski et al. (2015) extend the utility function in (13.71), in the context of dictator games, to incorporate both guilt and surprise. Their objective is to reconcile findings that support guilt aversion (e.g., Charness and Dufwenberg, 2006) with the results of Ellingsen et al. (2010) that find no support for guilt aversion.

Consider a dictator game in which the dictator has to divide one dollar between himself and the recipient. Denote by t , the amount transferred from the dictator to the receiver. Let us call the dictator and the receiver, respectively, players 1 and 2. Suppose that both players have the common utility function from material payoffs, m_i , $i = 1, 2$ that is strictly increasing and strictly concave ($m'_i > 0$ and $m''_i < 0$). Player 2 has the first order belief, b_2^1 , about the transfer to him, i.e., $b_2^1 = E_2[t]$. This assumption is in the spirit of the Köszegi–Rabin framework that allows for a distribution of reference points (see Part 1 of the book). The second order belief of player 1 about the belief, b_2^1 , is given by, b_1^2 , so $b_1^2 = E_1[b_2^1] = E_1E_2[t]$. Let the densities of b_2^1 and b_1^2 be given, respectively, by f_2^1 and f_1^2 (the corresponding cumulative distribution functions are F_2^1 and F_1^2). In contrast to (13.71), the dictator's utility function is given by

$$U_1(t | b_1^2) = m_1(1 - t) + S(t | b_1^2), \quad (13.73)$$

where

$$S(t | b_1^2) = \alpha \int_0^t (t - x) f_1^2(x) dx - \beta \int_t^1 (x - t) f_1^2(x) dx, \alpha \geq 0, \beta \geq 0. \quad (13.74)$$

Thus, conditional on his second order beliefs, player 1 derives elation from pleasantly surprising player 2 in states of the world where $t > x$, where x is a particular realization of his second order beliefs; this is the first term in (13.74). But he also feels guilty by letting down player 2 in states where $t < x$; this is the second term in (13.74). $\alpha \geq 0$ denotes the propensity of player 1 to make positive surprises (*surprise seeking*) and $\beta \geq 0$, the propensity to avoid negative surprises (*guilt aversion*). In order to make the problem interesting we assume that at least one of α and β is non-zero. The first term in (13.74) is the main difference between (13.71) and (13.73). A second difference is that (13.73) employs a framework that is similar to the Köszegi–Rabin utility function. Thus, players derive utility from the absolute levels of outcomes as well as gain–loss utility when outcomes are compared to all possible eventualities.

Player 1 does not directly observe b_2^1 but he observes an exogenously given informative signal, θ , which is the median of the distribution F_2^1 .⁶⁸ Thus, the conditional distribution of his second order belief is $F_1^2(b_2^1 | \theta)$ and the density is $f_1^2(b_2^1 | \theta)$. Since $t \in [0, 1]$, we assume that $F_1^2(b_2^1 | \theta)$ is continuously differentiable on $[0, 1] \times [0, 1]$. A higher value of the signal, θ , induces first order stochastic dominance in the conditional distribution of beliefs F_1^2 . Thus, we have

$$\theta' > \theta \Rightarrow F_1^2(b_2^1 | \theta') \leq F_1^2(b_2^1 | \theta) \text{ for all } b_2^1 \in [0, 1] \text{ and } \theta', \theta \in [0, 1].$$

⁶⁸ The model can also be solved if we assume that θ is the mean of the distribution F_2^1 .

Given the signal received by player 1 we can replace $S(t \mid b_1^2)$ by $S(t \mid \theta)$ and $U_1(t \mid b_1^2)$ by $U_1(t \mid \theta)$.

The dictator's transfer t to the recipient can now be found as the solution to the following problem.

$$t^*(\theta) \in \arg \max_{t \in [0,1]} U_1(t \mid \theta) = m_1(1-t) + S(t \mid \theta), \quad (13.75)$$

where

$$S(t \mid \theta) = \alpha \int_0^t (t-x) dF_1^2(x \mid \theta) - \beta \int_t^1 (x-t) dF_1^2(x \mid \theta), \alpha \geq 0, \beta \geq 0. \quad (13.76)$$

In what follows below we assume a strictly positive solution, i.e., $t \times (\theta) > 0$.⁶⁹ The intuition for the next proposition, Proposition 13.39, can be seen from (13.76). When the dictator is surprise seeking (relatively high α), then the first term in (13.76) is higher. Hence, there is a greater marginal incentive to create a pleasant surprise ($t^*(\theta) \geq \theta$). A high value of $t^*(\theta)$, in turn, magnifies the first term. The opposite case applies when β is relatively large.

Proposition 13.39 Suppose that there is a strictly positive solution, $t \times (\theta) > 0$, and $m'_i(x) \geq \frac{\alpha+\beta}{2}$.

- (i) If $\alpha > \beta$, then $t^*(\theta) \geq \theta$.
- (ii) If $\alpha < \beta$, then $t^*(\theta) \leq \theta$.

Proof: Applying Leibniz's rule to (13.75) we get

$$\frac{dU_1(t^*(\theta) \mid \theta)}{dt} = -m'_1(1-t^*) + (\alpha - \beta) [F_1^2(t^* \mid \theta)] + \beta \geq 0. \quad (13.77)$$

If $t^*(\theta) = 1$, then $\frac{dU_1(t^*(\theta) \mid \theta)}{dt} > 0$, otherwise if $0 < t^*(\theta) < 1$, then $\frac{dU_1(t^*(\theta) \mid \theta)}{dt} = 0$. Assuming that $m'_i(x) \geq \frac{\alpha+\beta}{2}$, we can rewrite the RHS of (13.77) as

$$(\alpha - \beta) [F_1^2(t^* \mid \theta)] + \beta - \frac{\alpha + \beta}{2} \geq 0,$$

or

$$(\alpha - \beta) \left[F_1^2(t^* \mid \theta) - \frac{1}{2} \right] \geq 0. \quad (13.78)$$

Since θ is assumed to be the median of the distribution $F_1^1, F_1^2(t^* = \theta \mid \theta) = E_1[F_2^1(\theta) \mid \theta] = \frac{1}{2}$. Thus, we have two main cases that satisfy (13.78).

- (a) If $\alpha > \beta$, then $F_1^2(t^* \mid \theta) \geq \frac{1}{2}$, or $F_1^2(t^* \mid \theta) \geq F_1^2(\theta \mid \theta)$. Thus, $t^*(\theta) \geq \theta$.
- (b) If $\alpha < \beta$, then $F_1^2(t^* \mid \theta) \leq \frac{1}{2}$, or $F_1^2(t^* \mid \theta) \leq F_1^2(\theta \mid \theta)$. Thus, $t^*(\theta) \leq \theta$.

A third case arises when $\alpha = \beta$ in which case we can have $t^*(\theta) \geq \theta$. ■

⁶⁹ Khametski et al., (2014) show that if we assume a lower bound on the marginal utility of consumption given by $m'_i(x) \geq \frac{\alpha+\beta}{2}$ for all $x \in [0, 1]$, then (in conjunction with the other assumptions made above) it is always the case that $0 < t^*(\theta)$.

Proposition 13.40 *The comparative static effects of the optimal transfer are as follows.*

- (i) $t^*(\theta)$ is increasing in θ if $\alpha < \beta$.
- (ii) $t^*(\theta)$ is decreasing in θ if $\alpha > \beta$.

Proof: We outline a heuristic proof for the case of a discrete distribution with two outcomes.⁷⁰ Suppose that there are only two values of second order beliefs of player 1 given by $x_L < x_H$ with respective probabilities p and $1 - p$. In order to ensure that both terms on the RHS of (13.76) are non-zero, suppose that the optimal solution $t^*(\theta)$ is such that $x_L < t^*(\theta) < x_H$. Then we can write (13.76) as

$$S(t | \theta) = \alpha p (t^*(\theta) - x_L) - \beta(1 - p) (x_H - t^*(\theta)), \alpha \geq 0, \beta \geq 0. \quad (13.79)$$

Now suppose that player 1 receives a higher signal $\tilde{\theta} > \theta$. Let $\tilde{p}$ be the probability of the outcome x_L under the signal $\tilde{\theta}$. We know that, by assumption, the distribution $F_1^2(b_2^1 | \tilde{\theta})$ first order stochastically dominates the distribution $F_1^2(b_2^1 | \theta)$, hence, $\tilde{p} < p$. From (13.79) we have

$$\frac{dS(t | \theta)}{dt} = \beta + p(\alpha - \beta).$$

Thus,

$$\frac{dS(t | \tilde{\theta})}{dt} - \frac{dS(t | \theta)}{dt} = (\alpha - \beta) (\tilde{p} - p). \quad (13.80)$$

Under both distributions, $F_1^2(b_2^1 | \tilde{\theta})$ and $F_1^2(b_2^1 | \theta)$, the marginal effect of a change in t on the first term in (13.75) is identical. Hence, the effect on the optimal solution t^* when θ changes, is determined entirely by the change in the marginal effect on the second term in (13.75). This is given by (13.80). Since $p > \tilde{p}$, the net marginal effect is determined entirely by the sign of $\alpha - \beta$. If $\beta > \alpha$, then $\frac{dS}{dt}$ goes up after the change in θ , hence, t^* also increases. Conversely, if $\beta < \alpha$, then t^* decreases. ■

Putting Propositions 13.39 and 13.40 together, the prediction of the model is as follows. If $\alpha > \beta$, then $t^*(\theta) \geq \theta$ (i.e., positive surprises) and there is a negative correlation between t^* and the signal θ . However, if $\alpha < \beta$, then $t^*(\theta) \leq \theta$ and there is a positive correlation between t^* and the signal θ . These predictions are confirmed by the empirical tests. If the dictator only had inequity-averse preferences then his transfer decision should not depend on his second order beliefs, nor should we observe the correlations between t^* and the signal θ that are suggested by Proposition 13.40.

The authors follow very closely the experimental design of Ellingsen et al. (2010) except that instead of the *direct-response method*, they use the *strategy method*. The dictator divides 14 euro between himself and the receiver. The receiver is asked to guess the average transfer that dictators will make in the population; this is a measure of θ in the experiments. Before these first order beliefs are revealed to the dictator (as in Ellingsen et al., 2010), the dictator is asked to reveal the transfers that he will make for each level of belief (rounded to the nearest 50 cents) of the receiver. It is assumed that the guess of the receiver is close to the median of her first order beliefs.

⁷⁰ Readers interested in the general proof can consult Khalmetski et al. (2015).

The strategy method allows for a within-subjects design, while earlier studies typically use the between-subjects design. For each dictator–receiver pair, the transfer (as revealed by the strategy method) that corresponded to the actual expectation level of the receiver, was implemented and paid out.

The results are as follows. First, the Ellingsen et al. (2010) results are replicated. Dictators transfer, on average, between 23–24% of their endowments. The average guesses of the recipients about the transfers are between 32–34% of the endowment. The correlation between guesses and transfers is statistically not different from zero. But there are within-subject differences in results that provide a more nuanced picture; 77% of the dictators alter their desired transfers at least once in response to a change in the guess of the recipient (as revealed by the strategy method). The within-subjects correlation between desired transfers and the guesses of the recipient is statistically significant for about 54% of the subjects. This is not consistent with purely outcome based models of inequity-averse social preferences.

When the within-subjects data from the strategy method is used to regress the desired transfers on the guesses of the recipients, the distribution of the coefficients on guesses that are significant at the 5% level is shown in Figure 13.28.

From Proposition 13.40, we expect the coefficient to be negative when $\alpha > \beta$ (surprise seeking dictators) and positive when $\alpha < \beta$ (guilt-averse dictators). Nearly 70% of the coefficients are distributed to the right of zero, so guilt aversion is relatively more important among dictators as compared to surprise-seeking. The distribution is asymmetric around zero (a two-sided sign test rejects the null that the median is zero). The average size of the positive coefficients is 0.58, while the average size of the negative coefficients is 0.53; these differences are statistically insignificant. In 27.2% of the cases, we have $t^* \geq \theta$, suggesting that surprise-seeking is more important than guilt aversion, i.e., $\alpha > \beta$; see Proposition 13.39. Furthermore, for 30% of the dictators who exhibit a significant correlation between transfers and guesses, the correlation is negative. This is consistent with $\alpha > \beta$; see Proposition 13.40. Thus, the authors draw very different conclusions from Ellingsen et al. (2010). Guilt aversion is very important but once one allows for surprise-seeking, which is found to be significant, then overall one does not find a correlation between actions and beliefs.




Figure 13.28 The distribution of the coefficients on guesses that are significant at the 5% level.

Source: Khalmetski et al. (2015).

In a second experiment, Khalmetski et al. (2015) also test for the importance of higher order beliefs, motivated by the idea of “guilt from blame” in Battigalli and Dufwenberg (2007). In the experiment reported above, receivers may not be sure of the intentionality of the transfers made by dictators. For instance, did the dictator really mean to surprise the receiver or did he make a mistake?

The authors consider two treatments. In the private treatment, the receiver is unaware that the dictator observes his first order belief. The dictator knows this, so he also knows that the receiver cannot impute any intentionality to his choice. In the public treatment, it is public knowledge that the dictator is told the first order beliefs of the receiver. Thus, the receiver can infer the intentionality of the dictator’s transfer (third order beliefs) and the dictator may in turn be influenced by his beliefs about the receiver’s third order beliefs. Thus, the fourth order beliefs of the dictator may be important in influencing the share given to the receiver. The experimental results show that average amounts sent in the public treatment are 70% higher than the private treatment, which is consistent with the importance of fourth order beliefs.

13.5.6 Dependence of utility on plans

In the framework of Geanakoplos et al. (1989), players derive utility only from the profile of equilibrium actions. However, the more general framework of Battigalli and Dufwenberg (2009) allows utility to depend on plans that are off the equilibrium path of play. Hence, players may evaluate their utility from the equilibrium strategy profile by taking account of their payoffs from non-equilibrium strategy profiles. This is reminiscent of the idea of stochastic reference points in Köszegi and Rabin (2006) and models of disappointment aversion that we discussed in Part 1 of the book.⁷¹ Here we illustrate the issues in a simple example from Battigalli and Dufwenberg (2009).

Suppose that we extend the trust game in Figure 13.19 by adding an extra stage, as shown in Figure 13.29; the first payoff in any pair of payoffs is Ann’s payoff. Ann can end the game by playing Out at the first node and both players get 0 each, or Ann could play Trust and give the move to Bob. Bob can then play Grab for payoffs (1, 3) (ignore the payoffs in the second row of this terminal node for now) or play Leave to pass the move to Ann. Ann then plays either Split or Reward for respective payoffs (2, 2) and (1, 3).

Now suppose that we consider this game as a psychological game. In particular, we are interested in the psychological utility of Ann. Let z be a terminal node (there are four such nodes in Figure 13.29), B be the set of beliefs for the game and s_A, s_B be respectively, the pure strategies of Ann and Bob (we rule out mixed strategies). Let $\zeta(s_A, s_B)$ be the terminal node that is reached when players play the strategy profile (s_A, s_B) . Suppose that Ann’s utility function at a terminal node z is given by

$$u_A(z, B, s_A) = \pi_A(z) - \theta \left[\sum_{s_B} b_A^1(s_B | h^0) \pi_A(\zeta(s_A, s_B)) - \pi_A(z) \right], \theta > 0. \quad (13.81)$$

The first term on the RHS in (13.81) is the material payoff of Ann at the terminal node z . The second term is gain–loss utility (in a manner similar to the Köszegi–Rabin framework) and θ is a parameter that captures the relative importance of psychological payoffs relative to material payoffs. Based on the initial beliefs of Ann at history h^0 about the strategy of Bob, $b_A^1(s_B | h^0)$, the second term compares each potential material payoff of Ann, $\pi_A(\zeta(s_A, s_B))$, as s_B is

⁷¹ See also Shalev (2000), who considers motivation similar to Köszegi and Rabin in settings with strategic interaction.




Figure 13.29 A psychological game in which utility depends on plans that are off the equilibrium path of play.

Source: Reprinted from *Journal of Economic Theory* 144(1): 1–35. Battigalli, P. and Dufwenberg, M. “Dynamic psychological games.” Copyright © 2009, with permission from Elsevier.

varied, with the actual material payoff at node z , $\pi_A(z)$. In this sense, Ann’s psychological utility depends on strategies that are off the path of play and on her beliefs about the plans of players.

To see the implications of this framework, consider Ann’s plans at history h^0 (where she has a choice between Trust and Out) about what she would do at history h^2 , where she has a choice between Split and Reward. Suppose also that Ann’s initial belief about the actions of Bob at history h^1 are given by $\alpha = b_A^1(\text{Leave} | h^0)$ and $1 - \alpha = b_A^1(\text{Grab} | h^0)$. Consider two possible plans for Ann.

I. Suppose that Ann’s initial strategy is to play (Trust, Split). Then, using (13.81), her initial expectation of the payoff from the strategy pair (Trust, Grab) is given by

$$1 - \theta [(1 - \alpha) 1 + \alpha 2 - 1] = 1 - \theta \alpha. \quad (13.82)$$

Thus, in Figure 13.29, in the second row of payoffs, following the strategy pair (Trust, Grab), we note the psychological payoff of Ann as $1 - \theta \alpha$. Therefore, Ann’s expected utility from the plan (Trust, Split) is

$$(1 - \alpha) (1 - \theta \alpha) + 2\alpha. \quad (13.83)$$

II. Suppose that Ann’s plan is (Trust, Reward). In this case, Ann’s material payoff at history h^2 is 1, which is independent of what Bob does at history h^1 (if Bob plays Grab, Ann still gets 1). Hence, Ann is never disappointed and her ex-ante payoff from the strategy pair (Trust, Reward) is given by 1.

Thus, using (13.82) and (13.83), Ann prefers the strategy (Trust, Split) to the strategy (Trust, Reward) if

$$(1 - \alpha) (1 - \theta \alpha) + 2\alpha > 1. \quad (13.84)$$

Thus, the lower is θ , the more likely is Ann to choose, ex-ante, the plan (Trust, Split). Conversely, if the inequality in (13.84) is reversed, then the ex-ante optimal plan is (Trust, Reward). But (Trust, Reward) is never optimal ex-post because at history h^2 , Ann always prefers to Split ($2 > 1$). Thus, plans can be *dynamically inconsistent*. If Bob can foresee this, he will know that Ann will always Split at history h^2 leading to disappointment for Ann from the strategy profile (Trust, Grab).

Solving the game backward, we now get that Ann will play Out at history h^0 if $0 > 1 - \alpha\theta$ or $\theta > \alpha^{-1}$.

13.6 Correlated equilibrium and social norms

In a series of papers that have culminated in Gintis (2009), Herbert Gintis provides a wide-ranging critique of the existing practice in classical game theory and other social sciences. He argues for a unification of the *behavioral sciences* (economics, anthropology, sociology, psychology, political science, and biology). The behavioral sciences ask the same underlying questions but propose incompatible models of human behavior and sociality. In this section, we give a brief account of some of the central arguments, particularly the shortcomings in classical game theory and the role of a correlated equilibrium.⁷²

Gintis favors retaining the assumption of rationality in economics, i.e., individuals have complete and transitive preferences. The evidence presented in this book in several places (e.g., from preference reversals and from time discounting) shows that individuals do not always make transitive choices. The solution advocated by Gintis is to fix the frame or context within which choices are transitive. Such conditional preferences ensure that intransitivity occurs only when the context or frame changes; for examples, see Sugden (2003) for a decision theoretic framework and al-Nowaihi and Dhami (2013) for work on time preferences.

The main shortcoming of classical game theory is that it provides no foundations for the assumption of *shared mental constructs*, e.g., common priors and common knowledge. The classical definition of a Nash equilibrium does not require a specification of the beliefs of the players. It merely requires that each player plays a best response to the strategy of the other player. Epistemic foundations for a Nash equilibrium give sufficient conditions for the existence of a Nash equilibrium (Aumann and Brandenburger, 1995). These conditions turn out to be extremely stringent. For instance, they require, in games with more than two players, that there be common priors about the state of the world and common knowledge of conjectures of each player in each state. However, there exist no epistemic foundations for common priors or common knowledge of conjectures.

It might be that *social norms* provide one possible source for mutual consistency of beliefs. As Gintis (p. 145) puts it: “It is not *reason* but *humanity* that leads us to believe that the Pareto-superior equilibrium is obvious.” Social norms may allow players to coordinate their conjectures about each other on Pareto optimal outcome. Such social norms may provide the underpinnings to common priors and to mutual and common knowledge among players.

It is sometimes argued that social norms are simply Nash equilibria that are based on *methodological individualism* (social phenomena are to be explained as arising from individual actions alone) and, typically, self-regarding preferences. There are two objections to this view. First, the epistemic foundations of a Nash equilibrium show that the requirements are too stringent relative to the cognitive limitations of humans. Second, most social norms are not Nash equilibria. Such norms do not in fact implement Nash equilibria, but they implement a correlated equilibria (Gintis, 2009, p. 132). Furthermore, Gintis (2009, chapters 10 and 11) argues forcefully that social norms are best explained by sociobiological models of *gene-culture coevolution*.⁷³

⁷² We suppress references to the research papers that the interested reader can find in Gintis (2009).

⁷³ See Section 12.1 in Gintis (2009) for the details. For gene-culture coevolution, see also Chapter 14.

	<i>L</i>	<i>R</i>
<i>U</i>	5, 1	0, 0
<i>D</i>	4, 4	1, 5

Figure 13.30 A coordination game.

Source: Herbert Gintis. *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences* © 2009 Princeton University Press. Reprinted with permission of Princeton University Press.

Consider the following problem of coordination from Gintis (2009, Section 2.11), shown in the normal form game in Figure 13.30.

There are two players, the row player (Robin) and the column player (Colin). Each player has two strategies and there are two pure strategy Nash equilibria (*U*, *L*) and (*D*, *R*). However, *Bayesian rationality*⁷⁴ does not provide any guidance on which of the two Nash equilibria will arise. There is also a mixed strategy Nash equilibrium, in which each player plays the two pure strategies with equal probability; the payoffs from this mixed strategy profile are (2.5, 2.5). The Pareto optimal outcome is (*D*, *L*) but it is not a Nash equilibrium. How can cooperation be sustained between the two players?

Classical game theory is based on *methodological individualism*. By contrast, the epistemic conditions for a Nash equilibrium require common priors and common knowledge, which appear to transcend personal characteristics of individuals. Indeed, these constructs appear to depend on social institutions that help align the beliefs and expectations of individuals.

Classical game theory explains human cooperation through *instrumental rationality arguments* within the setup of infinitely repeated games (or finitely repeated games with some irrationality). In this framework, players have self-regarding preferences but may cooperate because the discounted future monetary payoffs arising from cooperation are greater as compared to non-cooperation. For instance, if perfect public signals of the actions taken by the players are available, then any player who plays non-cooperatively is punished by the other players. If players are *patient enough*, then they are sufficiently deterred by the threat of future punishment, so they cooperate. These results can be extended to imperfect public signals.⁷⁵

Gintis raises three main objections to explaining cooperation in classical game theory. First, many of the results require players to be *too patient*. By contrast, players may have a high degree of present bias in their preferences (see Part 3 of this book). Second, the more realistic case is one of private signals. However, the set of conditions under which cooperation can be attained in repeated games under private signals is quite stringent. In particular, these results typically hold only for near-public signals in which the private signals are ϵ close to public signals.⁷⁶ Gintis (2009, p. 182) expresses reservations about the inability to specify the size of ϵ : “It is difficult to access how critical the information requirements of these folk theorems are because generally the theorem is proved for ‘sufficiently small ϵ ’ with no discussion of the actual order of magnitude involved.” A related objection is the evidence of human cooperation in static and even finitely repeated contexts under full information. Third, players often play mixed strategies in repeated

⁷⁴ *Bayesian rational* individuals satisfy the axioms for Savage’s subjective expected utility model (see Part 1 of the book for details).

⁷⁵ For the case of perfect public signals see Fudenberg and Maskin (1986) and for the case of imperfect public signals see Fudenberg et al. (1994).

⁷⁶ See, for instance, the set of papers in the *Journal of Economic Theory*, 2002, Volume 102.

games, yet the standard theoretical justification for such strategies based on purification theorems does not hold for such games (Gintis, 2009, Chapter 6).

An epistemic game is a normal form game, augmented to include the structure of knowledge in the game. In this framework, there is a state space, Λ , where each element of Λ refers to the pure strategy choices of the players. A knowledge partition specifies, for each player, the information players have about which states have arisen. The epistemic game also specifies in each state, the conjectures of each player and their subjective priors. Readers interested in a formal treatment can consult Chapters 4 and 5 in Gintis (2009).

Gintis argues that the *correlated equilibrium* is a more natural equilibrium criterion as compared to a Nash equilibrium. The most famous result in this regard is due to Aumann (1987) who showed that in an epistemic game, Bayesian rational agents with common subjective priors will play a correlated equilibrium. A correlated equilibrium is a Nash equilibrium of an epistemic game that is augmented with a social planner, or *choreographer*, who moves first, observes a random variable and then gives recommendations to each player to play a particular strategy. The resulting recommendations are a Nash equilibrium in the sense that it is optimal for any player to follow the recommendations, conditional on the other players following their recommendations.

One of the central insights in Gintis is to assign to social norms and conventions, the role of a choreographer in a correlated equilibrium. A *social norm* specifies the responsibilities, expectations, and entitlements that are associated with a particular social role. Let us consider the play of a correlated equilibrium in the example given in Figure 13.30. We will show that the Pareto optimal outcome (D, L) , which is not a Nash equilibrium in the normal form game, can be a part of the correlated equilibrium.

Suppose that there are three states $\Lambda = \{\omega_1, \omega_2, \omega_3\}$. The choreographer (who reflects social norms and conventions) observes the state in Λ but the players do not. The choreographer announces that the three states $\omega_1, \omega_2, \omega_3$ will occur with respective probabilities $\alpha_1, \alpha_2, \alpha_3 = 1 - \alpha_1 - \alpha_2$. Suppose that the actual realization of the state is some $\omega \in \Lambda$. The choreographer then gives recommendations to the players to play their pure strategies in the following manner. If state ω_1 occurs, then Robin is recommended to play U and Colin is recommended to play L . Under state ω_2 the respective recommendations are (D, L) , and under state ω_3 the recommendations are (D, R) . Both players know these rules, i.e., the state-based recommendations and the underlying distribution of the states, $\alpha_1, \alpha_2, \alpha_3$. However, the players (i) hear only their own recommendation, and (ii) do not observe the state.

The recommendations generate a *knowledge partition* of the set Λ . The respective knowledge partitions for Robin and Colin, P_R, P_C , are given by

$$P_R = \{\{\omega_1\}, \{\omega_2, \omega_3\}\}; P_C = \{\{\omega_3\}, \{\omega_1, \omega_2\}\}. \quad (13.85)$$

From (13.85), Robin knows when state ω_1 occurs because only in that state is he given the recommendation to play U . However, if he is given the recommendation to play D , he does not know which of the two states ω_2, ω_3 occurred. P_C can be explained analogously. Consistent with Bayes' rule, the probabilities assigned by Robin to the two states ω_2, ω_3 , when given the recommendation to play D are

$$p_R(\omega_2) = \frac{\alpha_2}{\alpha_2 + \alpha_3}; p_R(\omega_3) = \frac{\alpha_3}{\alpha_2 + \alpha_3}. \quad (13.86)$$

Similarly, when given a recommendation to play L , Colin assigns the following probabilities to the two states ω_1, ω_2

$$p_C(\omega_1) = \frac{\alpha_1}{\alpha_1 + \alpha_2}; p_C(\omega_2) = \frac{\alpha_2}{\alpha_1 + \alpha_2}. \quad (13.87)$$

The recommendations to the players are in pure strategies. However, given the knowledge partitions in (13.85), it is as if each player faces an opponent who plays mixed strategies. For instance, when Robin is given the recommendation to play D , he imagines facing an opponent who plays L with probability $p_R(\omega_2)$ and R with probability $p_R(\omega_3)$. Analogously, when given a recommendation to play L , Colin imagines that he is playing an opponent who plays U with probability $p_C(\omega_1)$ and D with probability $p_C(\omega_2)$.

Let us now check the conditions, which ensure that it is optimal for each player to follow the recommendations of the choreographer. First consider Robin. When given a recommendation to play U he is certain that the state is ω_1 , so he knows for sure that Colin's recommendation is to play L . The best response to L is indeed, U . On the other hand, when given a recommendation to play D , Robin imagines that his opponent is playing a mixed strategy, so D is optimal if

$$4p_R(\omega_2) + p_R(\omega_3) \geq 5p_R(\omega_2) \Leftrightarrow \alpha_1 + 2\alpha_2 \leq 1 \text{ (since } \alpha_3 = 1 - \alpha_1 - \alpha_2 \text{)}. \quad (13.88)$$

Now consider Colin. When given the recommendation to play R , Colin is certain the state is ω_3 and that Robin has been given the recommendation D . Clearly, R is the best response to D . However, when given a recommendation to play L , Colin imagines that Robin plays a mixed strategy. It is optimal for Colin to follow the recommendation L if

$$p_C(\omega_1) + 4p_C(\omega_2) \geq 5p_C(\omega_2) \Leftrightarrow \alpha_2 \leq \alpha_1. \quad (13.89)$$

From (13.88), (13.89) a correlated equilibrium must satisfy the two conditions

$$\alpha_1 + 2\alpha_2 \leq 1 \text{ and } \alpha_2 \leq \alpha_1. \quad (13.90)$$

These conditions ensure that the players will follow their respective recommendations, conditional on other players complying with their recommendations.

Suppose that the choreographer, in the spirit of a social planner, wishes to choose (α_1, α_2) in order to maximize the joint payoffs of the two players ($\alpha_3 = 1 - \alpha_1 - \alpha_2$ is determined residually). Recall that the payoff pairs $(5, 1)$, $(4, 4)$, and $(1, 5)$ occur with respective probabilities α_1, α_2 , and $1 - \alpha_1 - \alpha_2$. The problem of the choreographer (similar to that of a Stackelberg leader) is to maximize

$$\alpha_1 (5, 1) + \alpha_2 (4, 4) + (1 - \alpha_1 - \alpha_2) (1, 5)$$

subject to the constraints in (13.90). The solution to this linear programming problem is $(\alpha_1^*, \alpha_2^*) = (\frac{1}{3}, \frac{1}{3})$ and the symmetric payoffs of the two players, $(\frac{10}{3}, \frac{10}{3})$, exceed the joint payoffs from both pure strategy Nash equilibria as well as from the mixed strategy Nash equilibria. *Thus, social norms, in their manifestations as correlation devices, maximize aggregate payoff of both players.*

This result requires a *normative predisposition* on the part of the players to follow the choreographer, unless the players can do better by defecting to alternative strategies. The following quote is instructive (Gintis, 2009, pp. 143–4): “Humans have evolved psychological predispositions that render social norms effective. Social evolution . . . has favored the emergence both of social norms and human predispositions to follow social norms, to embrace common priors, and to recognize

common knowledge of many events. Nor is this process limited to humans.” Further, individuals are often willing to enforce social norms at a cost to themselves (see Part 2 of the book).

Two important remarks follow from this discussion that generalize to games with a finite number of players and strategies. First, a correlated equilibrium is a Nash equilibrium of the overall game in which the choreographer moves first and announces (α_1, α_2) as well as the recommendations to the two players, followed by the choice of players to follow or reject the recommendations. Second, any Nash equilibrium of the original game (in this case (U, L) and (D, R)) is a correlated equilibrium. The reason is that if the choreographer simply asked each player to play their part of a Nash equilibrium, it would be optimal for each player to do so. However, the set of correlated equilibria is much larger than the set of Nash equilibria, e.g., the Pareto optimal strategy profile (D, L) , which is not a Nash equilibrium was played as part of the correlated equilibrium in the example, above. Thus, social norms may lie outside the set of Nash equilibria.

13.7 Other behavioral models of how people play games

We now describe a range of other models in behavioral game theory that have been successful in explaining a subclass of the available games and the experimental results.

13.7.1 Analogy-based equilibrium

Jehiel (2005) considers how players might simplify complex games in order to make them cognitively more tractable. His proposal is that players may partition a set of games into exogenously given *analogy classes*. Thus, analogy classes are simply a primitive of the game, much as one takes information sets as part of the description of extensive form games. Each player attempts to learn the average behavior of the opponents in each analogy class. Each player then plays a best response to beliefs about the average behavior of others in an analogy class and these beliefs turn out to be correct in equilibrium. Players need not know the analogy classes of their opponents or the payoffs of the opponents. The resulting equilibrium is known as an *analogy-based expectations equilibrium* (ABEE).

The consistency of beliefs is a strong requirement, hence, it is unlikely to apply to games that are played only a few times. Rather, consistency may be viewed as the outcome of a learning process. In the finest possible partition, an ABEE coincides with a Nash equilibrium, but for coarser partitions, the two typically differ. Indeed, from an empirical perspective, interest centers on games where these two concepts differ.

Let us motivate an ABEE for extensive form games, using the extensive form game in Figure 13.31. There are two players A and B who move sequentially. At each node, $i = 1, 2, 3$, each player can either take action T_i to end the game or pass the move to the other player (action P_i), except at node $i = 3$ where the action P_3 ends the game. This is a finite game of perfect information, and the solution is easily found by backward induction. In the subgame perfect solution, player A plays T_3 at the last node, player B plays T_2 at the penultimate node, and player A plays P_1 at the first node.

Suppose that we now wish to find the ABEE. Assume that, player B exogenously groups together nodes 1 and 3, at which player A moves, into a single analogy class.

Proposition 13.41 (i) *The following is an ABEE for the game in Figure 13.31 when player B groups together nodes 1 and 3 into a single analogy class. Player A plays P_1 at node 1 and*




Figure 13.31 A sequential move game.

Source: Reprinted from *Journal of Economic Theory* 123: 81–104. Jehiel, P., “Analogy-based expectation equilibrium.” Copyright © 2005, with permission from Elsevier.

T₃ at node 3. Player B plays P₂ at node 2. (ii) Given the analogy partition in (i), the subgame perfect solution is not an ABEE.

Proof: (i) Player B groups together nodes 1 and 3 into a single analogy class, so he considers the average behavior of player A at these nodes. Thus, player B believes that he faces a player A who will pass half the time and take half the time. This is equivalent to believing that player A plays a mixed strategy that randomizes uniformly over the two choices. Hence at node 2, when it is player B's turn to move, he gets 2 from taking (action T₂) and $\frac{1}{2}(1 + 4) = 2.5 > 2$ from passing (action P₂). Thus, he prefers to play P₂ at node 2, as claimed. Player A finds it optimal to play T₃ at the third node ($2 > 0$). Conditional on the subsequent optimal actions, player A's choice of P₁ at node 1 is optimal (because $-1 < 2$).

(ii) Conditional on the analogy classes, it is never optimal for player B to play T₂ at the second node. Hence, the subgame perfect outcome cannot be sustained as an ABEE. ■

The benefits from cognitive simplifications arising from analogy classes are obvious from Proposition 13.41. Players do not need to know (1) the payoffs of others, or (2) how their opponents group nodes into analogy classes. At the moment, we do not have a satisfactory theory of how players form analogy classes. The concept of an ABEE will be strengthened if the foundations for analogy classes can be developed.

Let us now consider the concept of an ABEE for two-player normal form games. Suppose that we have a set of normal form games $G = \{G_1, G_2, \dots, G_T\}$; to save on notation we shall label a particular element by G . The probability that game G is played is $p(G) \geq 0$ and $p(G) = 1$. Suppose that we index the two players by r and c (respectively, for the row and the column player). The respective pure strategy sets of the two players, S_r and S_c , are identical across all games. For any game, $G \in G$, the action profile is given by $(s_r, s_c) \in S_r \times S_c$ and the payoffs of the players are given by $u_j(s_r, s_c | G)$, $j = r, c$. The mixed strategy sets of players r and c are given, respectively, by Σ_r and Σ_c ; these are probability distributions over the sets S_r and S_c .

A strategy for any player specifies the mixed strategy that he intends to play for each game in the set G . Thus, a mixed strategy for player $j = r, c$ is denoted by $\sigma : G \rightarrow \Sigma_j$ so $\sigma_j(s_j | G)$ denotes the probability with which player j plays strategy s_j in game G .

Consider for player $j = r, c$, a partition A_j of the set G and call this an analogy partition. Let the cells of the partition A_j be denoted by $a_1, a_1, \dots, a_{n_j}$. If a game $G \in G$ appears in the cell $a_j \in A_j$, then $a_j(G)$ is called the analogy class of player j at G . Suppose that $\sigma_j(s_j | G)$ is the actual strategy of player $j = r, c$, and game G belongs to the analogy partition $a_i(G)$ of the other player, player $i = r, c$, $i \neq j$. Player i is then assumed to *know* the average strategy of player j , $\bar{\sigma}_j$, for each of the analogy classes but not the disaggregated behavior of player j across

	a	b	c	d	e		a	b	c	d	e
A	25, 10	0, 10	10, 20	0, 0	0, 0	A	15, 0	20, 10	10, 0	5, 5	0, 0
B	20, 15	15, 0	5, 0	10, 0	0, 0	B	0, 10	25, 10	0, 0	10, 20	0, 0
C	15, 0	10, 0	0, 0	5, 25	25, 0	C	10, 0	15, 0	5, 25	0, 0	25, 0
Game G1						Game G2					

Figure 13.32 Games G1 and G2 in Huck et al. (2011).

Source: Reprinted from *Games and Economic Behavior* 71: 351–65. Huck, S., Jehiel, P., and Rutter, T., “Feedback spillover and analogy-based expectations: a multi-game experiment.” Copyright © 2011, with permission from Elsevier.

the individual games in a_i .⁷⁷ Note that it is not assumed that player i can “compute” the average strategy from disaggregated information, but only that player i somehow “knows” the average strategy, perhaps as a steady state of some process of learning that occurs outside the model; not all readers will find this aspect of the model persuasive. Player i believes that the probability with which player j plays the strategy s_j for any game in the analogy partition $a_i(G)$, is computed as follows.

$$\bar{\sigma}_j(s_j | G) = \frac{\sum_{G' \in a_i(G)} p(G') \sigma_j(s_j | G)}{\sum_{G' \in a_i(G)} p(G')}. \quad (13.91)$$

In computing the average strategy of player j in (13.91), the actual strategy $\sigma_j(s_j | G)$ of player j for each game $G' \in a_i(G)$ is required but player j does not observe the disaggregated strategies. Given the beliefs in (13.91) about their opponents and the analogy partitions of both players, each player plays a best response. The mutual best responses computed in this manner, constitute an ABEE, as stated formally in the next proposition.

Definition 13.23 (*Analogy-based expectations equilibrium, ABEE*) : A strategy profile (σ_1, σ_2) is an analogy-based expectations equilibrium (ABEE), given the analogy partitions A_1, A_2 , if for all $G \in \mathbf{G}$ and all s_j^* in the support of $\sigma_j(G)$, $j = r, c$ we have

$$s_j^* \in \arg \max_{s_j \in S_j} \sum_{s_i \in S_i} \bar{\sigma}_i(s_i | G) u_j(s_j, s_i | G), \quad i \neq j \text{ and } i = r, c,$$

and $\bar{\sigma}_i$ is given in (13.91).

From Definition 13.23, when players use the finest possible analogy partitions, then $\bar{\sigma}_j(s_j | G)$ coincides with $\sigma_j(s_j | G)$, hence, an ABEE coincides with a Nash equilibrium. Further, none of the players is required to know the payoff of the other player—just the aggregate distribution of play of the opponent in each analogy class.

Huck et al. (2011) test for an ABEE in simple two-player games. Consider games G1 and G2 in Figure 13.32. Each of the two games has a unique pure strategy Nash equilibrium that can be found by iterated deletion of strictly dominated strategies- (A, c) in G1 and (B, d) in G2.

In order to consider an ABEE in these games one needs to specify the analogy classes and the probability that each of the two games occurs. Here we have the following: $S_r = \{A, B, C\}$,

⁷⁷ This behavior is similar to that of a fully cursed player in a cursed equilibrium; see Section 13.7.3.

$S_c = \{a, b, c, d, e\}$, $G = \{G1, G2\}$ and $p(G1) = p(G2) = 0.5$. The analogy partition used by the column player is always fine, i.e., $A_c = \{\{G1\}, \{G2\}\}$. However, we consider two possible cases of the analogy partitions of the row player. A fine partition in which $A_r = \{\{G1\}, \{G2\}\}$ and a coarse partition in which $A_r = \{\{G1, G2\}\}$.

When the row player uses a fine partition, as already noted, an ABEE coincides with a Nash equilibrium. Hence, there is a unique ABEE for games G1 and G2 in this case. The results when the row player uses a coarse partition are given in the next proposition that the reader is asked to prove as an exercise.

Proposition 13.42 *Suppose that the row player uses the coarse analogy partition $A_r = \{\{G1, G2\}\}$ and the column player uses the fine analogy partition $A_c = \{\{G1\}, \{G2\}\}$. Then in the unique ABEE, the row player plays B in game G1 and A in game G2. The column player plays a in game G1 and b in game G2.*

From Proposition 13.42, the ABEE is different from the Nash equilibrium of the game. It is, however, the case that strategy a in Game 1 and b in Game 2 give the column player the highest payoffs in each row, which could be one reason that subjects in experiments play these strategies; an even more stringent test would be to differentiate the predictions of Nash and ABEE when this is not the case.

Four different treatments are considered by Huck et al. (2011). These treatments are differentiated by the type of feedback that is given to a player about the play of the opponent in the last five rounds of play. Column players are given accurate feedback but the row player is given different feedback in the alternative treatments. For instance, in the COARSE treatment, the row player is simply informed about the actions of the column player in the last five rounds without identifying which of the two games, G1 or G2, those actions come from. In the FINE treatment, the difference from the COARSE treatment is that the games which gave rise to these actions are identified. In two other treatments, FEASIBLE and HARD, players are given the joint distribution of actions and games in which these actions were chosen in the last five rounds. With unlimited cognitive powers, subjects can, in principle, convert this data to the same form as in the treatment FINE, but with limited cognitive powers that would be difficult.

In each treatment, the behavior of subjects in experiments stabilizes by about round 15, but it converges to different equilibria. In treatments FINE and FEASIBLE, behavior converges to the Nash equilibrium but in treatments COARSE and HARD it converges to the unique ABEE of the game. The behavior of row players in the treatment COARSE fits well players who play a best response to the aggregate play of the column players over both games G1 and G2, which is consistent with an ABEE.

Experiments in which data is processed from multiple games at the same time are not very common. These results call for more empirical testing of games in this category. The ABEE can be extended to other classes of games. For an extension to static two-player games of incomplete information, see Jehiel and Koessler (2008).

13.7.2 *Subjective heterogeneous quantal response equilibrium*

The QRE and the level- k /CH models are two distinct approaches to relaxing the assumptions of a Nash equilibrium. The former relaxes the assumption of strict best replies and the latter relaxes the assumption of equilibrium in beliefs. Rogers et al. (2009) propose the *subjective heterogeneous quantal response equilibrium* (SQRE) that combines and clarifies the link between QRE and CH.

Consider the logit QRE framework for normal form games with n players that is described in Section 13.2.1 above; we carry the same notation forward in this section. The logit density is a one-parameter family with a single parameter λ . Assume that we have a distribution of types indexed by i such that $\lambda_i \in [0, \infty)$ is the type of player i . Let us introduce heterogeneity among the players about their perceived distribution of types of other players. The subjective belief of a player of type λ_k about the type of player $i \neq k$ is denoted by $F_i^k(\lambda_i | \lambda_k)$; we denote the corresponding density by $f_i^k(\lambda_i | \lambda_k)$.

The logit choice probabilities of players may now be computed in a manner similar to that in (13.7). The probability that a player of type λ_i chooses the action s_{ij} is given by

$$\phi_{ij}(\lambda_i, \bar{u}_i(\sigma)) = \frac{\exp(\lambda_i \bar{u}_{ij}(\sigma))}{\sum_{k=1}^{n_i} \exp(\lambda_i \bar{u}_{ik}(\sigma))}. \quad (13.92)$$

A strategy of player i , $\phi_i : [0, \infty) \rightarrow \Sigma_i$ gives the optimal choice probabilities with which each of the possible types plays each of the pure strategies in the set S_i .

Thus, the beliefs of a type λ_k of player k that player i will play the action s_{ij} is given by

$$\sigma_{ij}^k(\phi_i | \lambda_k) = \int_0^\infty \phi_{ij}(\lambda_i, \bar{u}_i(\sigma)) f_i^k(\lambda_i | \lambda_k) d\lambda_i. \quad (13.93)$$

One may analogously compute the beliefs of a player i of type λ_i about the probability with which all other types of all other players play their respective pure actions; denote this by $\sigma_{-i}^i(\phi_{-i} | \lambda_i)$. Thus, we can find the expected payoff of a type λ_i of player i by $\bar{u}_{ij}^{\lambda_i}(\sigma_{-i}^i) = \{\bar{u}_{i1}^{\lambda_i}, \bar{u}_{i2}^{\lambda_i}, \dots, \bar{u}_{in_i}^{\lambda_i}\}$ where $\bar{u}_{ij}^{\lambda_i}(\sigma_{-i}^i)$ is the payoff to player i from the strategy s_{ij} when others play σ_{-i}^i .

$$\bar{u}_{ij}^{\lambda_i}(\sigma_{-i}^i) = \sum_{s_{-i} \in S_{-i}} \sigma(s_{-i}) u_i(s_{ij}, s_{-i}). \quad (13.94)$$

In a SQRE, the three equations, (13.92), (13.93), (13.94) hold simultaneously. A SQRE reflects the fact that beliefs of players are subjective, so there is no assumption of common priors. The SQRE concept is too general and requires empirically plausible restrictions to make it operational.

A special case of SQRE arises when there is heterogeneity with respect to the types of the players but this heterogeneity is common knowledge. In this case, $F_i^k(\lambda_i | \lambda_k) = F_i(\lambda_i | \lambda_k)$ for all $i, k = 1, 2, \dots, n$. This formulation gives rise to *heterogenous quantal response equilibrium*, HQRE. Thus, $F_i(\lambda_i | \lambda_k)$ is common knowledge but the type of any player, λ_i , is private knowledge to player i . When the heterogeneity in types becomes degenerate, then HQRE reduces to the logit QRE that is described in Section 13.2.

It is possible that players might not be able to imagine types of other players above a certain bound but accurately perceive types lower than this bound—the authors' term for this is an *imagination bound*. Possible reasons include not being able to perceive higher types than one's own (overconfidence) or being cognitively limited in making the computations required to perceive higher types.

Suppose that type λ_i of player i believes that any type λ_{-i} of another player cannot be higher than the imagination bound, $b_i(\lambda_i)$. Higher types of players (higher values of λ_i) are interpreted as more skilled players in terms of payoff responsiveness. The function b_i is commonly known, increasing, and continuous. The positive slope of b_i implies that higher types imagine other players to be also higher or more skilled (although possibly lower than their own type). In this case,

the beliefs of a player k of type λ_k about any players i of type λ_i are given by

$$F_i^k(\lambda_i | \lambda_k) = \begin{cases} F_i(\lambda_i)/F_i(b_k(\lambda_k)) & \text{if } \lambda_i \leq b_k(\lambda_k) \\ 1 & \text{if } \lambda_i > b_k(\lambda_k) \end{cases}. \quad (13.95)$$

Player k of type λ_k cannot imagine other types to be higher than the upper bound $b_k(\lambda_k)$. The division by $F_i(b_k(\lambda_k))$ normalizes the cumulative density of all types λ_i of player i lower than type $b_k(\lambda_k)$. Such a model is known as a *truncated quantal response equilibrium model* (TQRE). In the special case where $b_i \rightarrow \infty$ for all $i = 1, 2, \dots, n$, the TQRE collapses to the standard HQRE model. At the other extreme when the imagination bound $b_i \approx 0$, then player i believes that all other types of all other players choose randomly (similar to the beliefs of level $L1$ players in level- k models). If $b_i(\lambda_i) \leq \lambda_i$, then player i of type λ_i is said to be *self-limited* because he cannot imagine that any of the other types is more skilled in the sense of payoff responsiveness.

The existence of a HQRE and TQRE in finite games is guaranteed; see Rogers et al. (2009).

In a TQRE, heterogeneity is indexed by payoff responsiveness, λ_i , while in level- k models and CH, it is indexed by k . In level- k models, for $k \geq 1$, although the source of the heterogeneity is different, we have $b_k(\lambda_k) \leq \lambda_k$. Rogers et al. (2009) give two restrictions on TQRE models such that the behavioral predictions of TQRE and the CH model can be made arbitrarily close and the approximating equilibrium in TQRE is unique. Suppose that the allowed values of λ are discrete: $0, \gamma, 2\gamma, \dots, k\gamma, \dots$. The parameter γ is known as the grid size. A player of type k or a level- k player is associated with $\lambda = k\gamma$. The first restriction is that for any value of λ we must have $b_i(\lambda) = \frac{k}{1+k}\lambda$, where k is the type associated with the chosen value of λ . Thus for the lowest possible value of $\lambda = 0$, we have $b_i(0) = 0$; these players choose randomly like type $L0$ players. Level-1 players correspond to $\lambda = \gamma$ and, in this TQRE model, they quantal best-respond, assuming that others are type $L0$ and so on.

In level- k and CH models, recall that players best-respond to their inaccurate beliefs about others. If we let $\gamma \rightarrow \infty$, then we have $\lambda \rightarrow \infty$. This corresponds to players playing best replies to their beliefs about others (as in a Nash equilibrium). In the main result of the paper, Rogers et al. (2009) show that as $\lambda \rightarrow \infty$, in almost all finite games, the behavioral predictions of TQRE and CH are essentially equivalent in the approximate sense.

Rogers et al. (2009) test the four different models, QRE, HQRE with uniform distribution of types, Poisson TQRE, and Poisson CH on 17 different full information games. Maximum likelihood methods are used for model estimation and model fits. Across the games, the parameter values for each of the models is similar but there are also a few games that are outliers. Results are quite different for pooled and game-specific data. For instance, CH gives the best fit in game-specific results but the worse fit for the pooled data. Heterogeneity in types appears not to be too important. However, once downward looking heterogeneity of the truncated types is assumed, as in the TQRE, then TQRE gives a better fit than CH in 10 out of 17 games and a better fit than the QRE in 15 out of 17 games. The authors conclude (p. 1456): “Overall, our results indicate that a combination of ‘strategic’ heterogeneity and stochastic choice are significant factors in explaining behavior in these games, and a downward looking belief structure helps explain behavior in some of the games.”

Further support for the HQRE has come from the bubbles game in Moinas and Pouget (2013). Consider sequential trading games in which trading occurs at discrete intervals and at any point in time, the owner of a single asset can sell or not sell a single unit of an object. Assume also that it is commonly known that the fundamental value of the object is zero. Under common knowledge of rationality, a price cap should not result in any bubbles. The argument proceeds by backward

induction (as in centipede games). Suppose that an investor finds himself saddled with an asset whose price equals the price cap. Then he cannot sell it and make a profit. Going back a period, the investor would not buy such an asset. Proceeding backwards in this manner, one can rule out a bubble ever forming on the asset. However, the higher the price cap, the greater is the number of steps of iterated reasoning that are required to make the argument (as in centipede games with a higher number of moves by players). In contrast to the theoretical predictions, Moinas and Pouget (2013) find that bubbles form in the presence as well as absence of price caps.⁷⁸ The model that best fits their data is HQRE and the analogy-based equilibrium model (ABEE).

13.7.3 Cursed equilibrium

Eyster and Rabin (2005) introduce the notion of a *cursed equilibrium*, a solution concept that can be applied to Bayesian games. Relative to fully rational players, cursed players do not fully take account of the correlation between the actions of other players and their types. Rather, each player believes that with some probability $\chi \in [0, 1]$, all types of the opponent play their average optimal action (and not the type-contingent optimal action). With the complementary probability, $1 - \chi$, the opponent is believed to play the optimal type-contingent actions. Such players are said to be χ -cursed.

Consider a simple, static, game of incomplete information, the Akerlof (1970) model of the market for lemons. Suppose that there are two types of used cars, lemons and peaches, that are distinguished by their quality. The proportion of each type among the used cars is commonly known to be $\frac{1}{2}$. A lemon is worth nothing to either the buyer or the seller. A peach is worth \$3000 to the buyer and \$2000 to the seller. Suppose that the market price is \$1000. Since this price is lower than the value of the car to the sellers of a peach, there is no supply of peaches in the market. Hence, at a price of \$1000, the only cars on the market are lemons. The optimal type-contingent actions of the sellers are: the peach seller sells with probability 0, and the lemon seller sells with probability 1. Thus, the average action over the two types is to sell with probability $\frac{1}{2}$.

Knowing this, a fully rational buyer will not buy cars in this market (market failure). However, a χ -cursed buyer thinks that with probability χ , each seller plays the average action, i.e., sell with probability $\frac{1}{2}$. Hence, the probability, in the mind of a χ -cursed buyer, that the car on the market is a “peach” is given by:

$$(1 - \chi) \times 0 + \chi \times \frac{1}{2} = \frac{\chi}{2}. \quad (13.96)$$

In (13.96), a fully rational buyer corresponds to $\chi = 0$, and correctly believes that a peach has a zero probability of being on the market at the given price. However, a fully cursed buyer sets $\chi = 1$ and mistakenly believes that the car is a peach with probability $\frac{1}{2}$. Hence, a χ -cursed buyer thinks that the expected value of the car is $\frac{\chi}{2} \times \$3000$. Since the market price is \$1000, it is therefore optimal to buy if $\chi > \frac{2}{3}$.

We now give a more formal account of a χ -cursed equilibrium. Recall the notation and definitions for finite Bayesian games in the Appendix on game theory, which we summarize first. In a finite Bayesian game, $G(N, S, u, T, \mu)$, $N = \{1, 2, \dots, n\}$ is the set of n players, nature is denoted

⁷⁸ The classic experimental paper by Smith et al. (1988) showed that even in the absence of asymmetric information, one gets price bubbles on assets followed by crashes. The difficulty of accounting for bubbles within a neoclassical framework is discussed in Shiller (2000).

by player $i = 0$; $S_i \subset \mathbb{R}$ is the finite pure strategy set of player i and $S = \times_i S_i$ is the set of all pure strategy profiles (nature's strategy set S_0 is empty); $S_{-i} = \times_{j \neq i} S_j$ is the set of pure strategy profiles of all players other than player i . T_i is the finite set of possible types of player i ; $T = \prod_i T_i$ is the set of all possible types. $\mu : T \rightarrow \mathbb{R}$, which is common to all players, gives the set of conditional probability distributions over T ; $u_i : S \times T \rightarrow \mathbb{R}$ is the payoff of player i and it depends on the chosen strategies and player types; the vector of payoffs of all players is denoted by $u = (u_1, u_2, \dots, u_n)$. ΔS_i is the set of all probability distributions over S_i ; the type-contingent mixed strategy of player i is denoted by $\sigma_i : T_i \rightarrow \Delta S_i$, where $\sigma_i(s_i | t_i)$ is the probability with which type t_i of player i plays strategy s_i under the mixed strategy $\sigma_i(t_i)$; $\sigma_{-i}(s_{-i} | t_{-i})$ is the probability with which other players of type $t_{-i} \in T_{-i}$ play some pure strategy profile $s_{-i} \in S_{-i}$ under the mixed strategy $\sigma_{-i}(t_{-i})$.

Definition 13.24 (*Bayesian Nash equilibrium*): Consider the static Bayesian game $G = \{N, S, u, T, \mu\}$. The strategy vector (σ_i, σ_{-i}) is a Bayesian Nash equilibrium if for all $i \in N$, all types $t_i \in T_i$ and all actions s_i^* such that $\sigma_i(s_i^* | t_i) > 0$ we have

$$s_i^* \in \arg \max_{s_i \in S_i} \sum_{t_{-i} \in T_{-i}} \mu(t_{-i}, t_i) \sum_{s_{-i} \in S_{-i}} \sigma_{-i}(s_{-i} | t_{-i}) u_i(s_i, s_{-i}; t_i, t_{-i}). \quad (13.97)$$

In a cursed equilibrium, players imagine that with some probability, all the types of other players play the average strategy. Consider the mixed strategy of all players other than player i , $\sigma_{-i}(t_{-i})$. The average strategy of all other players is then given by

$$\bar{\sigma}_{-i}(s_{-i} | t_i) = \sum_{t_{-i} \in T_{-i}} \mu(t_{-i}, t_i) \sigma_{-i}(s_{-i} | t_{-i}). \quad (13.98)$$

When other players actually play $\sigma_{-i}(s_{-i} | t_{-i})$, player i believes that with some probability $\chi \in [0, 1]$ they play the average strategy $\bar{\sigma}_{-i}(s_{-i} | t_i)$. The size of χ gives the degree of cursedness of player i . The types of players can be correlated, so, the computation of $\bar{\sigma}_{-i}(s_{-i} | t_i)$ may depend on the type t_i . Thus, different types of the same player may have different beliefs about the average actions of other players. However, if the types of the players are independent, i.e., $\mu(t_{-i} | t_i) = \mu(t_{-i} | t'_i)$ for all $t_i, t'_i \in T_i$, then all players have the same beliefs about the average actions of others.

The χ -cursed equilibrium is defined as follows.

Definition 13.25 (χ -cursed equilibrium): Consider the static Bayesian game $G = \{N, S, u, T, \mu\}$. The strategy vector (σ_i, σ_{-i}) is a χ -cursed equilibrium, if for all $i \in N$, all types $t_i \in T_i$ and all actions s_i^* such that $\sigma_i(s_i^* | t_i) > 0$ we have

$$s_i^* \in \arg \max_{s_i \in S_i} \sum_{t_{-i} \in T_{-i}} \mu(t_{-i}, t_i) \sum_{s_{-i} \in S_{-i}} [\chi \bar{\sigma}_{-i}(s_{-i} | t_i) + (1 - \chi) \sigma_{-i}(s_{-i} | t_{-i})] u_i(s_i, s_{-i}; t_i, t_{-i}). \quad (13.99)$$

When $\chi = 0$, from Definitions 13.24 and 13.25 we have that the cursed equilibrium and the Bayesian Nash equilibrium coincide. When $\chi = 1$, players are said to be *fully cursed*. In this case, they completely ignore the correlation between the types and actions of the other players and believe that each type of the other player plays the average action only. This case appears quite similar to the case of the analogy-based equilibrium of Jehiel (2005); see Jehiel and Koessler (2008). It is as if, for each opponent, all types have been put into one analogy class and the average action is assigned to all types. The theory allows different players to be cursed to different degrees.

Proposition 13.43 (Eyster and Rabin, 2005): For any finite Bayesian game $G = \{N, S, u, T, \mu\}$, and each $\chi \in [0, 1]$, a χ -cursed equilibrium exists.

One drawback of a cursed equilibrium is that the psychological underpinnings are not as persuasive as some of the other behavioral game theory models. Indeed, as the authors themselves point out (p. 1633): “All said, however, our primary motivation for defining cursed equilibrium is not based on learning or any other foundational justification, but rather on its pragmatic advantages as a powerful empirical tool to parsimoniously explain data in a variety of contexts.” However, the similarity with the analogy-based equilibrium and other justifications that the authors suggest (e.g., overconfidence) may lead to better foundations for such an equilibrium concept in the future. A second drawback is that the concept applies to Bayesian games only. For games of full information, the concept of a cursed equilibrium reduces to the standard Nash equilibrium and its refinements. This, we know, does not explain the experimental results from a variety of full information games. The case for a cursed equilibrium would be even stronger, if it could be shown that it is a heuristic that is somehow turned-on only when individuals encounter games of incomplete information.

One of the main applications of the χ -cursed equilibrium is to the explanation of bidding behavior in auctions that we consider in Section 13.8.

13.7.4 Evidential equilibrium

In static games, players are often uncertain of the actions of their opponents. This could be because, say, the game represents a relatively new strategic situation, or the players have not played much against each other before. In these situations, evidence, mainly from psychology, suggests that players assign *diagnostic significance* to their own actions in forming inferences about how others may behave. Such reasoning by players is known as *evidential reasoning* (ER), also sometimes known as *social projection*.

ER is most likely to apply when one views the opponents in a game as *like-minded*. Following the literature in psychology, we may refer to the reasoning used in classical game theory as *causal reasoning*. Causal reasoners assign no diagnostic value to their own actions. We argue that the evidence is best described by a mixture of evidential and causal reasoners in the population; indeed the theoretical framework, below, allows for such a mixture.

Robbins and Krueger (2005) describe ER as follows: “Using their own disposition or preferences as data, people can make quick predictions of what others are like or what they are likely to do.” In a survey article, Krueger (2007) writes “The concept of social projection is once again generating vigorous theory development and empirical research . . . social projection is among the simplest, oldest, and arguably most central concepts of the field.”

The rest of this section draws on al-Nowaihi and Dhami (2015) in giving a formal account of evidential reasoning for static games of full information.⁷⁹ The claim is not that ER helps us to explain the data for all static games but that it adds an important element to the toolkit of behavioral economists, enabling them to explain puzzles that are difficult to explain otherwise.

Consider some suggestive examples of evidential reasoning.

Example 13.18 (*False consensus effect*): Ross et al. (1977) asked subjects if they would walk around a university campus wearing a sandwich board with the word “REPENT” written on it. Those who agreed to do so, estimated that 63.5% of their peers would also agree to do so,

⁷⁹ For static games of incomplete information, see Dhami and al-Nowaihi (2013).

while those who refused, expected 76.7% of their peers to also refuse. Clearly these fractions add up to more than one and so cannot constitute consistent beliefs. This evidence is consistent with many subjects using evidential reasoning to impute diagnostic value to their own actions in forming beliefs about the likely actions of other like-minded people (the student population in the university in this case). This is an example of the “false consensus effect.”

Example 13.19 (Contributions among like-minded people in public goods games): In public goods games under causal reasoning (see Part 2), it is a Nash equilibrium for each player to free-ride (zero contributions). The experimental evidence shows that in the absence of punishment of non-cooperators, contributions decay over successive rounds. Gächter and Thöni (2005) investigate whether cooperation in public goods games is higher among like-minded people. They initially run a single-round of the public goods experiment. The subjects are then grouped by the amount of contributions they made in this round. For instance, the top three contributors are grouped into a separate group (the TOP group) as having the greatest inclination to contribute. The public goods game is then played separately for each group. Over the next ten rounds, contributions are much higher and free riding much lower in the TOP group, which achieves nearly the first best level of contributions in several rounds.⁸⁰ Evidential reasoning would predict that when high contributors have greater confidence that they are grouped with like-minded people, they take their own predisposition to contribute more as evidence that like-minded others are also likely to contribute more. In conjunction with reciprocity, this may achieve high levels of contributions.

Example 13.20 (Evidential reasoning about states of other humans): In an experiment conducted by Van Boven and Loewenstein (2005), participants were asked to imagine a hiker who is lost in the woods and asked whether the hiker was more likely to be thirsty or hungry. When the experimental participants were made to experience these states (thirst or hunger) they attributed similar states to the hypothetical hiker despite their own states having no correlation or cue value with the state of the lost hiker. In each case, they seem to assign diagnostic significance to their own state in inferring the state of the lost hiker.

Example 13.21 Table 13.4 reports survey evidence from successive US Presidential elections on the beliefs of voters, who voted for each political party, about the winning chances of their preferred political party. Voters who intend to vote Democrat typically assign higher probabilities to the Democrat candidate winning relative to those who intend to vote Republican. By contrast, voters who intend to vote Republican assign higher probabilities to a Republican win relative to voters who vote for Democrats. The assigned probabilities do not add up to one despite the fact that most people get their information from common sources such as TV, newspapers, and opinion polls. This is consistent with voters taking their own actions as diagnostic of what other like-minded voters might do. These findings show that, in some circumstances at least, people behave as if their actions were causal even when they are merely diagnostic or evidential. Quattrone and Tversky (1984) use these insights to explain why voters vote in elections. Voters who have a predisposition to vote (say, due to a sense of civic responsibility) believe that other like-minded responsible voters will also vote. On the other hand, if

⁸⁰ Even the endgame effect, i.e., the sharp drop in contributions in the last round of the experiment is most pronounced among the BOTTOM group.

Table 13.4 Percentage of voters from each political party expecting their own candidate to win in US presidential elections.

Year	Presidential candidates	% of Democrat voters expecting Democrat win	% of Republican voters expecting Republican win
1988	Dukakis vs. Bush	51.7	94.2
1984	Mondale vs. Reagan	28.8	99.9
1980	Carter vs. Reagan	87.0	80.4
1976	Carter vs. Ford	84.2	80.4
1972	McGovern vs. Nixon	24.7	99.6
1968	Humphrey vs. Nixon	62.5	95.4
1964	Johnson vs. Goldwater	98.6	30.5
1960	Kennedy vs. Nixon	78.4	84.2
1956	Stevenson vs. Eisenhower	54.6	97.6
1952	Stevenson vs. Eisenhower	81.4	85.9

Source: Forsythe et al. 1992.

they procrastinate and choose not to vote for their party, other like-minded people might also reason in a similar manner, and the other party could win the elections. So they vote.

Under each of the examples above, players who use evidential reasoning do not believe that by altering their own action they are able to alter the actions of the other players. They simply take their own predisposition to follow an action as diagnostic evidence for what other, like-minded, players might do.

IS EVIDENTIAL REASONING RATIONAL OR IS IT A USEFUL HEURISTIC?

One may be interested in the epistemic foundations of equilibrium concepts in economics (Aumann and Brandenburger, 1995; Polak 1999). Empirical violations of the Nash equilibrium (see Chapter 12) also violate the epistemic conditions for a Nash equilibrium. A related question is about the *rationality* of classical reasoning in games (causal reasoning) and non-classical reasoning such as ER (Rapoport, 1966, pp. 139–41; Bicchieri and Green, 1997; Bonanno, 2013). Economists accustomed to traditional notions of rationality may find evidential reasoning to be less than fully rational.

Our interest is neither in the epistemic foundations nor in the rationality or otherwise of ER, although the independent study of these issues is very worthwhile.⁸¹ The relevant issue is whether the evidence on human behavior supports ER. The results of experimental games appear to be best explained by a mixture of people; some appear to use evidential reasoning while others appear to use causal reasoning. At the same time, it appears that there is no single theory that is capable of explaining the evidence from all static games.

The motivation for evidential reasoning may be given with an analogy. Ever since the work of Allais (1953), the evidence from a large number of experiments turned out to be inconsistent with the predictions of models of formal decision theory that were based on the axioms of rationality (Kahneman and Tversky, 2000). In the early 1970s Daniel Kahneman and Amos Tversky proposed an alternative approach, the *heuristics and biases* approach (Kahneman et al., 1982; Kahneman,

⁸¹ The older literature on conjectural variations (or reaction function approach, see Friedman, 1977) and the newer literature on psychological game theory, that requires a weaker consistency of beliefs with actions relative to a Nash equilibrium, are not relevant to our arguments, since both use causal reasoning.

	C	D
C	2, 2	0, 3
D	3, 0	1, 1

Figure 13.33 The prisoner's dilemma game.

2003). In this approach, instead of optimizing in the classical sense, decision makers follow simple heuristics that are *fast* and *frugal* in the use of information. The heuristics often work well but sometimes lead to errors relative to a strictly optimizing behavior. This approach is considered in Part 7.

Indeed, many other alternatives to a Nash equilibrium, such as level- k models, may also be viewed in a similar manner. For instance, ex-post, players who use level- k models are often surprised by the outcome; their beliefs and the equilibrium actions need not be mutually consistent, unlike in a Nash equilibrium. However, under level- k models, one should not observe the play of dominated strategies by any level- k player, $k > 1$. This makes it difficult to explain why there is so much cooperation in prisoner's dilemma games or why people vote.

We propose here, a heuristics and biases approach to game theory, using ER as the appropriate vehicle for static games. A range of other evidence supports the interpretation of ER as a heuristic. People who use evidential reasoning are not aware of using it despite their behavior being obviously consistent with evidential reasoning. Evidential reasoning appears to arise as an *automatic* response, rather than as a *deliberate* response, i.e., it does not require awareness, effort, or intention. Evidence supporting this view comes from experiments which show that evidential reasoning was not hampered by cognitive load or time required to complete an action; see Krueger (2007). Furthermore, other evidence, also reported in Krueger (2007), suggests that considerable cognitive effort is required to suspend evidential reasoning. The evidence from Acevedo and Krueger (2005) indicates that evidential reasoning applies to human–human interaction but not to human–nonhuman interaction.

WHY IS THERE SO MUCH COOPERATION IN THE PRISONER'S DILEMMA GAME?

Consider the static prisoner's dilemma shown in Figure 13.33. Defection is a strictly dominant strategy. Hence, a player using causal reasoning should defect. However, experimental evidence indicates relatively high cooperation rates. Rapoport (1988) finds cooperation rates of 50% in the prisoner's dilemma game. In a study of the prisoner's dilemma based on high stakes outcomes from a British TV show called *Goldenballs*, Darai and Grätz (2010) find unilateral cooperation rates of 55% for stakes above £500 and cooperation rates of 74% for stakes below this level. Zhong et al. (2007) show that the cooperation rates in prisoner's dilemma studies go up to 60% when positive labels are used (such as a “cooperative game” rather than a “prisoner's dilemma game”). When purely generic labels are used (such as C and D), then the cooperation rates are about 50%. In a novel study, Khadjavi and Lange (2013) find that actual prison inmates cooperate far more than the student population. In particular, while the cooperation rates among students playing the static prisoner's dilemma game is 37%, the cooperation rate among prison inmates is 56%.

al-Nowaihi and Dhami (2015) show that one cannot explain cooperation in the prisoner's dilemma game simply by invoking social preferences. Since the action C is a dominated strategy it is not *rationalizable* and cannot be supported in a *correlated equilibrium* either. Cooperation

in a prisoner's dilemma game cannot be explained by using evolutionary games either. The reason is that the set of evolutionary stable equilibria is a subset of the Nash equilibria of the game, which, in this case, is the unique outcome (D,D); see Part 5 for evolutionary game theory. All models that require players to play their best responses, e.g., level- k models and CH models cannot also explain the cooperative outcomes in the prisoner's dilemma game because cooperation is a strictly dominated strategy. In QRE models, players play the strategies that give a higher expected payoff with a greater probability. Hence, they will always play the strategy defect with a greater probability. Thus, over time, we expect most of the outcomes to converge towards (D,D).

In the light of the evidence, the explanation of cooperative behavior in prisoner's dilemma games is an open problem in economics and social science. This cannot be dismissed as just a minor annoyance; the prisoner's dilemma game is the most widely used game in the social sciences (outside economics) as a metaphor for why cooperation may break down.

Lewis (1979) used ER to explain the unexpected levels of cooperation in the one-shot prisoner's dilemma game. Mutual cooperation is better than mutual defection. Furthermore, some players may be hard-wired to cooperate and/or they may be inequity averse. If players use ER, they may take their own preference for mutual cooperation as diagnostic evidence that their rival also has a preference for mutual cooperation, in which case both players are more likely to cooperate. These views are borne out by the evidence. Cooperators (those who play C) believe the probability that other players cooperate is between 0.6 and 0.7. Similarly, players who defect believe that other players will defect with probabilities between 0.6 to 0.7; see Krueger (2007). Darai and Grätz (2010) find that preplay promises of players to cooperate with each other and shake hands on it, increase the cooperation rates. In classical game theory this should have no effect, but evidential reasoning may provide one possible explanation. Such actions may increase the likelihood, in the minds of players, that they are facing like-minded players, hence, they cooperate more.

Like-mindedness is compatible with both outcomes, C and D, and we do observe both outcomes in the PD game. Those who play C (respectively, D) also believe that a disproportionately large share of the other players will play C (respectively, D). However, why then do we observe so much cooperation in the static prisoner's dilemma game? The reason is that most humans intrinsically like to cooperate, especially when the cooperative outcome is also Pareto optimal. Consider, for instance, the following quote from Gintis (2009, p. 145):

It is not reason but humanity that leads us to believe that the Pareto-superior equilibrium is obvious. We humans by virtue of our gene-culture coevolutionary history and our civilized culture, harbour a default frame that says "in coordination games, unless you have some special information that suggests otherwise, conjecture that the other player also considers the frame to be a default frame and reasons as you do, and choose the action that assumes your partner is trying to do well by you."

This also explains the label salience of outcomes in the prisoner's dilemma game: When framed as a cooperative game, the cooperation rate is higher (see above).

The explanation of cooperation based on evidential reasoning appears compelling. The fact that a sizeable fraction of the subjects in experiments also defect suggests that the results are best accounted for by a mixture in the population of people who use evidential reasoning and causal reasoning. A more formal account of the prisoner's dilemma game under evidential reasoning is given in Section 13.7.4, below.

EVIDENTIAL GAMES AND EVIDENTIAL EQUILIBRIUM

We now give an exposition of evidential games and provide a general framework of analysis for static games that is suitable for extensions to other heuristics. Consider a static game of complete information, $\{N, S, u\}$. The set of players is $N = \{1, 2, \dots, n\}$. $S_i \subseteq \mathbb{R}$ is the set of actions open to player i . $S = \times_{i=1}^n S_i$ gives all possible action profiles of the players. $S_{-i} \subseteq \mathbb{R}^{n-1}$ is the set of actions open to the other players. Σ_i is the set of probability distributions over S_i . Any $\sigma_i \in \Sigma_i$ is a mixed strategy of player i ; $\sigma_i(s_i)$ is the probability with which player i plays action $s_i \in S_i$. Since a pure strategy corresponds to $\sigma_i(s_i) = 1$, we will simply refer to σ_i as a strategy for player i .

A strategy profile of all players is denoted by $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n) \in \Sigma$, where $\Sigma = \times_{i=1}^n \Sigma_i$ is the set of all possible profiles of strategies. Players choose their strategies simultaneously. A particular profile of the strategies of the other players is denoted by $\sigma_{-i} = (\sigma_1, \dots, \sigma_{i-1}, \sigma_{i+1}, \dots, \sigma_n) \in \Sigma_{-i} = \times_{j \neq i} \Sigma_j$. The von-Neumann utility function of player i is $u_i : \Sigma \rightarrow \mathbb{R}$ and u is the vector of payoffs of the players. Given a strategy profile $\sigma = (\sigma_i, \sigma_{-i}) \in \Sigma$, the payoff of player i is $u_i(\sigma_i, \sigma_{-i}) \in \mathbb{R}$. The structure of the game, $\{N, S, u\}$, is common knowledge among the players.

A social projection function captures the beliefs of a player about the strategies of the other players, conditional on his own strategy.⁸² These beliefs are based on introspection that is founded on evidential reasoning. Evidential reasoning rules out the possibility that these beliefs cause the other players to play the expected strategies.

Definition 13.26 (Social projection functions, SPF): A social projection function for player i (SPF for short), is a mapping $P_i : \Sigma_i \rightarrow \Sigma_{-i}$, that assigns to each strategy, $\sigma_i \in \Sigma_i$, for player i , an $(n-1)$ vector of strategies for the other players. We write $P_{ij}(s_j|\sigma_i)$ for the “subjective belief” that player i assigns to player j playing $s_j \in S_j$, conditional on player i playing $\sigma_i \in \Sigma_i$. Hence, $P_{ij}(s_j|\sigma_i) \geq 0$ and $\sum_{s_j \in S_j} P_{ij}(s_j|\sigma_i) = 1$. We may write $P_i(\sigma_i) = \sigma_{-i}^e(\sigma_i) = (P_{i1}, \dots, P_{ii-1}, P_{ii+1}, \dots, P_{in})$ to indicate that $P_i(\sigma_i)$ is the $(n-1)$ vector of strategies that player i “expects” that the other players will follow if player i adopts the strategy σ_i .

We now define causal reasoning, the dominant type of reasoning used in classical game theory, and evidential reasoning.

Definition 13.27 (Causal reasoning): We say that player i uses causal reasoning if $P_{ij}(s_j|\sigma_i)$ is independent of σ_i for each $s_j \in S_j$ and each $j \neq i$, i.e., if $P_i(\sigma_i) = P_i(\sigma'_i)$ for all $\sigma_i, \sigma'_i \in \Sigma_i$.

Definition 13.28 (Evidential reasoning) We say that player i uses evidential reasoning if it is not necessarily the case that $P_i(\sigma_i) = P_i(\sigma'_i)$ for all $\sigma_i, \sigma'_i \in \Sigma_i$.

We now discuss Definitions 13.26–13.28. In a static game of full information, players may be uncertain of the actions taken by others. Under evidential reasoning, player i resolves this uncertainty by assigning diagnostic significance to his own choice of strategy, σ_i , in inferring the strategies of the other players, σ_{-i} , using his social projection function, P_i . For this reason, Definition 13.28 allows for $P_i(\cdot|\sigma_i)$ to change as σ_i changes.

It is quite important to realize that there is no causal connection between σ_i and σ_{-i} . The choice of σ_i by player i merely influences his “belief” about the strategies, σ_{-i} , of the other players. In particular, players who use evidential reasoning know that their own actions have no causal effects in altering the actions of others when they change their own actions. Further, a SPF (Definition 13.26)

⁸² We use the term social projection function because, on the one hand, it is obviously connected with social projection and evidential reasoning and, on the other hand, to distinguish it from the term projection function as commonly used in mathematics.

specifies the beliefs of a player for all possible actions of others, including out-of-equilibrium actions. The beliefs of a player need not be fulfilled in equilibrium, as in other disequilibrium-in-beliefs models such as the level- k model.

If player i uses causal reasoning as in classical game theory (see Definition 13.27) then he assigns no diagnostic significance to his own strategy, σ_i , in inferring the strategies, σ_{-i} , followed by the other players.⁸³ Thus, under causal reasoning, $P_i(\sigma_i)$ remains fixed as σ_i changes. From Definitions 13.27, 13.28, causal reasoning is a special case of evidential reasoning.

Example 13.22 *This framework is quite general, hence, there is merit in considering it in some detail. The social projection function simply captures beliefs that a player has about other players. We have introduced evidential reasoning in this section. However, there are other plausible methods of belief formation that could be modeled using the machinery introduced here. For instance, in level- k models applied to static games, denote the strategy of a type L_k player by σ_k and the pure strategy set by S_k that has n_k elements. Suppose also that type L_0 players choose uniformly randomly. Then the social projection function of a level L_1 player is given by $P_{10}(s_0|\sigma_1) = \frac{1}{n_0}$ for all $\sigma_1 \in \Sigma_1$ and for all $s_0 \in S_0$. Suppose that, conditional on these beliefs, the best response of type L_1 is, say, σ_1^* . Then the social projection function for any level L_2 player is $P_{21}(s_1|\sigma_2) = \sigma_1^*(s_1)$ for all $\sigma_2 \in \Sigma_2$ and for all $s_1 \in S_1$. The social projection functions of higher types can be constructed analogously.*

Players need not impute diagnostic significance to their actions when others are perceived not to be like-minded. Here we assume the special case that all players are like-minded. al-Nowaihi and Dhami (2015) consider the case where other players may be a mixture of like-minded and non like-minded players (ingroup and outgroup members).

Many different types of social projection functions (SPFs) are possible. There are possibly various degrees of like-mindedness. However, the examples above suggest a particularly salient SPF that imparts the main predictive content to theory. This SPF seems important when choices are low dimensional, such as the binary choices “cooperate” or “defect” in a prisoners’ dilemma game, or vote Democrat or Republican in US Presidential elections, or coordinate or fail to coordinate in simple coordination games. In each of these cases, players may think that other like-minded players will choose *identical actions* to their own. The corresponding SPF in these cases is the *identity social projection function*. In this case, conditional on playing any strategy σ_i , player i believes that like-minded players will play an identical strategy; for more general social projection functions, see al-Nowaihi and Dhami (2015).

Definition 13.29 (*Identity social projection function*): Suppose that all players in N have the same action set, i.e., $S_i = S_j = S$ for all $i, j \in N$. Let P_i be the social projection function for player $i \in N$. Recall that $P_{ij}(s|\sigma_i)$ is the probability that player i assigns to player j playing action s when the strategy of player i is given by σ_i . If $P_{ij}(s|\sigma_i) = \sigma_i(s)$ for all $s \in S$ and all $j \in N$, then we say that P_i is an identity social projection function.

Definition 13.30 (*Evidential game*): Consider the static game of complete information, $\{N, S, u\}$. Let $P = (P_1, P_2, \dots, P_n)$ be a profile of social projection functions, where P_i is the social projection function of player $i \in N$ (Definition 13.26). Then, the game augmented

⁸³ By contrast, in a dynamic game (under causal reasoning), if (say) player 1 moves first, choosing the strategy σ_1 , followed by player 2 who chooses strategy σ_2 , having observed a realization of σ_1 , then σ_2 may very well depend on σ_1 . When choosing σ_1 , player 1 will take into account the influence of his choice on the future behavior of player 2. This should not be confused with evidential reasoning.

with the vector of social projection functions, P , denoted by $\Gamma = \{N, S, u, P\}$, is called an evidential game. We say that players in such a game use evidential reasoning. In particular, if each $P_i(\sigma_i)$ is independent of σ_i , then we say that Γ is a causal game.

From Definition 13.30, a causal game is a special case of an evidential game. Suppose $P_i(\sigma_i)$ is independent of σ_i , for each player, i , so that $\Gamma = \{N, S, u, P\}$ is a causal game. Γ is still richer than the static game of complete information, $\{N, S, u\}$, because Γ incorporates players' beliefs about other players' actions, as given by P .

Definition 13.31 (Optimal strategies): An optimal strategy for player i , $\sigma_i^* \in \Sigma_i$, in the evidential game $\Gamma = \{N, S, u, P\}$ (Definition 13.30), maximizes the payoff function of player i , $u_i(\sigma_i, P_i(\sigma_i)) \equiv u_i(\sigma_i, \sigma_{-i}^e(\sigma_i))$.

Definition 13.32 (Evidential equilibria): The strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*, \dots, \sigma_n^*) \in \Sigma$ is an evidential equilibrium of the evidential game $\Gamma = \{N, S, u, P\}$ if σ_i^* is an optimal strategy for each $i \in N$ (Definition 13.31).

The optimal actions of each player in an evidential equilibrium are found by solving a decision theoretic problem in what is classically a strategic game. There are no higher order beliefs. Each player does not think about strategically exploiting the SPF of the other player. Indeed the game $\Gamma = \{N, S, u, P\}$ does not involve any assumptions about the mutual or common knowledge of P . Requiring u_i and P_i to be continuous functions on compact mixed strategy spaces guarantees that an equilibrium exists in the decision theoretic problem..

Note that Definition 13.32 only requires that a strategy for a player be optimal given *his* beliefs. But, of course, beliefs may not turn out to be correct, ex-post. Ultimately, the choice among models in all science is guided by the evidence. The evidence from many static games reviewed in this book, shows that beliefs about others can often turn out to be incorrect, at least when the game is new and played only a few times. Nevertheless, it is of interest to consider the special case where beliefs turn out to be correct, at least in equilibrium. This is the subject of the next two definitions.

Definition 13.33 (Mutually consistent strategies): A strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*, \dots, \sigma_n^*) \in \Sigma$ of the evidential game $\Gamma = \{N, S, u, P\}$ (Definition 13.30) is a mutually consistent vector of strategies if $P_i(\sigma_i^*) = \sigma_{-i}^*$ for all $i \in N$, i.e., if $P_{ij}(s_j | \sigma_i^*) = \sigma_j^*(s_j)$, for all $i, j \in N$, $i \neq j$, and all $s_j \in S_j$.

In other words, a strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*, \dots, \sigma_n^*)$ is a mutually consistent vector of strategies, if for all players $i, j \in N$, $i \neq j$ and all actions, s_j , open to player j , the probability $P_{ij}(s_j | \sigma_i^*)$ that player i “believes” that player j will play action s_j (given σ_i^*) is equal to the “actual” probability $\sigma_j^*(s_j)$ with which player j plays s_j . Indeed, imposing such a social projection is quite strong, as are the epistemic conditions for a Nash equilibrium.

Definition 13.34 (Consistent evidential equilibria): A consistent evidential equilibrium of the evidential game $\Gamma = \{N, S, u, P\}$ is an evidential equilibrium, $\sigma^* \in \Sigma$, which is also a mutually consistent vector of strategies (Definitions 13.32 and 13.33).

As one might expect, there is a natural correspondence between Nash equilibria and consistent evidential equilibria. This is formally stated and established by the following proposition, which is a special case of the famous result of Aumann and Brandenburger (1995) on the epistemic foundations of a Nash equilibrium. The proof, which is straightforward, is left as an exercise for the reader.

Proposition 13.44 (*al-Nowaihi and Dhami, 2015*):

- (a) Let $\sigma^* \in \Sigma$ be a Nash equilibrium in the static game of complete information, $\{N, S, u\}$. Consider the (constant) social projection functions: $P_i(\sigma_i) = \sigma_{-i}^*$, for all $\sigma_i \in \Sigma_i$ and for all $i \in N$. Then σ^* , is a consistent evidential equilibrium in the evidential game $\Gamma = \{N, S, u, P\}$. Furthermore, Γ is a causal game.
- (b) Let $\sigma^* \in \Sigma$ be an evidential equilibrium in the evidential game $\Gamma = \{N, S, u, P\}$, where P is the profile of constant social projection functions $P_i(\sigma_i) = \sigma_{-i}^*$, for all $\sigma_i \in \Sigma_i$ and for all $i \in N$ (hence, σ^* is a consistent evidential equilibrium and Γ is a causal game). Then σ^* is a Nash equilibrium in the static game of complete information $\{N, S, u\}$.

AN APPLICATION: THE PRISONER'S DILEMMA GAME

We now use the prisoner's dilemma game in Figure 13.33 to illustrate some of the implications of evidential reasoning. Call the row player "player 1" and the column player "player 2." Consider the following cases.

- Case 1. Both players use evidential reasoning.
- Case 2. One player uses evidential reasoning but the other uses causal reasoning.
- Case 3. Both players use causal reasoning but beliefs turn out to be wrong ex-post.
- Case 4. Both players use causal reasoning and beliefs turn out to be correct ex-post.

The last case illustrates Proposition 13.44a, namely, that there will always be a profile of social projection functions for which a given Nash equilibrium corresponds to a consistent evidential equilibrium. A comparison of Cases 3 and 4 shows that causal reasoning does not guarantee mutual consistency of beliefs.

For the prisoner's dilemma, $N = \{1, 2\}$, $S_1 = S_2 = \{C, D\}$, $S = \{C, D\} \times \{C, D\}$, and u is shown in the payoff matrix in Figure 13.33. Each player has a dominant action, D , thus, the unique Nash equilibrium of this game is (D, D) . By contrast, the empirical evidence reviewed above, shows that 50% or more of the outcomes involve the play (C, C) . We can set up this game as either an evidential game or a causal game. Since the Nash equilibrium for a prisoner's dilemma game is in pure strategies, we focus on pure strategies.

From Definition 13.28, in general, $P_{ij}(\cdot|\sigma_i)$ will vary with σ_i , but under causal reasoning $P_{ij}(\cdot|\sigma_i)$ will be independent of σ_i (Definition 13.27). For all cases below, we consider the identity social projection function (Definition 13.29).

Case 1: The social projection function for each player is given by

$$\begin{cases} P_{12}(C|C) = 1, P_{12}(D|C) = 0; P_{12}(D|D) = 1, P_{12}(C|D) = 0, \\ P_{21}(C|C) = 1, P_{21}(D|C) = 0; P_{21}(D|D) = 1, P_{21}(C|D) = 0. \end{cases} \quad (13.100)$$

From (13.100), both players use evidential reasoning, so this is an evidential game. Given the social projection functions, the unique optimal choice is C for each player. Hence, (C, C) is the unique evidential equilibrium. Notice that (C, C) is a mutually consistent vector of strategies (Definition 13.33). Hence, (C, C) , the Pareto optimal outcome, is a consistent evidential equilibrium (Definition 13.34).

Readers trained in classical game theory could wish to make the following argument. Since $P_{12}(C|C) = 1$ and players cannot change the actions of opponents by changing their action, why does not player 1 defect and increase his outcome to 3, hence, breaking the cooperative equilibrium? Rapoport (1966, pp. 139–41) argues along the following lines. It is also the case that

$P_{12}(D|D) = 1$. Player 1 knows that player 2, who is like-minded, will also benefit by thinking in an analogous manner and, hence, both will defect, in which case both get a payoff of 1. However, given (13.100), this does not maximize the individual payoffs, which is required for an evidential equilibrium (see Definitions 13.31, 13.32). Hence, player 1 finds it worthwhile to cooperate. For analogous reasons, player 2 also cooperates. Thus, mutual cooperation is sustained. Whether such a simple heuristic, independent of epistemic foundations, is reasonable or not, is entirely an empirical matter.

Remark 13.4 *We now make a set of observations that apply to each of the four cases that we consider in this section. The SPF for each player in (13.100) also specifies out-of-equilibrium beliefs. For instance, $P_{12}(D|C) = 0$ specifies the belief of player 1 that if he were to find it optimal to play C, then the probability that player 2 will play D is zero. It does not mean that by playing C, player 1 can induce player 2 to play D; indeed there is no such causal link. The optimal actions of each player, (C, C) in this case, are found from solving a decision-theoretic problem, using the heuristic of evidential reasoning, and none of the players attempt to strategically exploit the SPF of the other player. Indeed there is no requirement in an evidential game that there even be mutual knowledge of the SPF.*

Case 2. The social projection function for each player is given by

$$\begin{cases} P_{12}(C|C) = 1, P_{12}(D|C) = 0; P_{12}(D|D) = 1, P_{12}(C|D) = 0, \\ P_{21}(C|C) = 1, P_{21}(D|C) = 0; P_{21}(C|D) = 1, P_{21}(D|D) = 0. \end{cases}$$

Player 1 uses evidential reasoning and, in particular, his identity social projection function. Player 2 uses causal reasoning and, in particular, mistakenly assumes that player 1 will always cooperate. This is an evidential game. The unique evidential equilibrium (Definition 13.32) is (C, D). It is an evidential equilibrium because each player's chosen action is optimal, given his beliefs, which are captured by his social projection function. It is not a consistent evidential equilibrium because the belief of player 1 turns out to be mistaken in equilibrium ($P_{12}(C|C) = 1$ but player 2 plays D instead). By contrast, the belief of player 2 that player 1 always plays C turns out to be correct in equilibrium.

Case 3. The social projection function for each player is given by

$$\begin{cases} P_{12}(C|C) = 0, P_{12}(D|C) = 1; P_{12}(D|D) = 0, P_{12}(C|D) = 1, \\ P_{21}(C|C) = 0, P_{21}(D|C) = 1; P_{21}(C|D) = 1, P_{21}(D|D) = 0. \end{cases}$$

Both players use causal reasoning (Definition 13.27). Given these social projection functions, the unique payoff maximizing strategy for each player is to play D (Definition 13.31). Hence, (D, D), is the unique evidential equilibrium (Definition 13.32). However, (D, D), is not a mutually consistent vector of strategies (Definition 13.33) because each player expects his opponent to play C in response to D but the opponent's response is D. Hence (D, D) is not a consistent evidential equilibrium (Definition 13.34). The importance of this case arises because it is sometimes believed that a Nash equilibrium under causal reasoning requires mutual consistency of beliefs.

Case 4. The social projection function for each player is given by

$$\begin{cases} P_{12}(C|C) = 0, P_{12}(D|C) = 1; P_{12}(D|D) = 1, P_{12}(C|D) = 0, \\ P_{21}(C|C) = 0, P_{21}(D|C) = 1; P_{21}(C|D) = 0, P_{21}(D|D) = 1. \end{cases}$$

	S	R
S	10, 19	4, -7
R	(-14, -7, -1), 4	19, 10

Figure 13.34 A coordination game.

As in Case 3, both players use causal reasoning. Given the social projection function, playing D is the unique optimal strategy for player 1 and player 2 (Definition 13.31). Hence (D, D) is the unique evidential equilibrium (Definition 13.32). Furthermore, (D, D) is a mutually consistent vector of strategies (Definition 13.33) because each player expects his rival to play D and, in fact, his rival does play D . Hence, (D, D) is a consistent evidential equilibrium (Definition 13.34). The unique Nash equilibrium of this game is, of course, (D, D) . Hence this case illustrates Proposition 13.44a, namely, a Nash equilibrium of the game $\{N, S, u\}$ is also a consistent evidential equilibrium of the game $\{N, S, u, P\}$ with a suitable choice of social projection functions, P .

The machinery developed above applies to static games of full and incomplete information and as argued above (see Example 13.22) to potentially many other belief formation processes. For further applications of evidential reasoning to oligopolistic competition and Nash bargaining, see al-Nowaihi and Dhami (2015). However, the empirical evidence on evidential reasoning comes largely from static games. It would be of interest to see if the evidence also develops in the direction of dynamic games because the theoretical framework is amenable to such an analysis.

13.7.5 Noisy introspection

Goeree and Holt (2004) propose a model of noisy introspection (NI). The essential idea is to relax the assumption of rationality of beliefs in the QRE. Consider the coordination game in Figure 13.34 from Guyer and Rapoport (1972).

The row and column players have two pure strategies, S and R . Three different versions of this game can be generated by choosing, successively, the three alternative payoffs to the row player from the strategy profile (R, S) : -14 , -7 , and -1 . The two pure strategy Nash equilibria are (S, S) and (R, R) in the three different versions. The equilibrium (S, S) is also the maximin equilibrium. When the payoff from (R, S) is $(-14, 4)$, there is also a mixed strategy Nash equilibrium in which the column player plays S with probability 0.39 and the row player plays S with probability 0.19.⁸⁴ The strategies S and R are short for, respectively, “safe” and “risky.”

The strategy profile (R, R) is risky because if any of the players unilaterally deviates, then the opponent’s payoffs become negative. Thus, a player is likely to play R only if there is a high probability that the opponent will also play R . In the Harsanyi–Selten (1988) sense, the outcome (S, S) is risk dominant. These equilibria are identical in the three different versions of the game. Hence, the equilibrium that the players choose to coordinate on should not change. Experimental evidence, however, shows that as the payoff of the row player from the strategy profile (R, S) increases, successively, from -14 to -7 to -1 , the incidence of S falls from 89% to 88% to 75%. The likely reason is that the strategy R tends to be perceived as less risky.

⁸⁴ One can also compute the mixed strategy when the payoffs from (R, S) are, respectively, $(-7, 4)$ and $(-1, 4)$.

Denote the expected payoffs of player i (where i can be the row or the column player) from playing, respectively, the strategy S and R by u_{iS} and u_{iR} . Then the logit probabilities in a QRE, as in (13.7), with $\mu^{-1} = \lambda$, require that the probability of playing S by player i is

$$\pi_{iS} = \frac{\exp(\mu^{-1}u_{iS})}{\exp(\mu^{-1}u_{iS}) + \exp(\mu^{-1}u_{iR})}, \quad (13.101)$$

where μ is a noise parameter. As $\mu \rightarrow 0$, the strategy with the greatest payoff is chosen with probability 1 (check that when $\mu \rightarrow 0$ and $u_{iS} > u_{iR}$, the RHS of (13.101) equals 1).

In Figure 13.35, the darker step-functions show the logit best responses for a small level of noise, $\mu = 0.1$. This accounts for the slightly rounded edges of the step-functions (for $\mu = 0$ we get sharp edges as in a Nash equilibrium). The interior intersection of the Nash best response functions (for $\mu = 0$) give the mixed strategy equilibrium of the game, (0.39, 0.19). The Nash best responses also intersect at the southwest and the northeast edges, picking out the two pure strategy Nash equilibria (S, S) and (R, R).

For a higher value of $\mu = 6.6$, estimated from the Guyer and Rapoport (1972) data, the logit best responses are shown as the lighter curves. These intersect close to the northeast edge of Figure 13.35 and their point of intersection gives the QRE at (0.97, 0.98). In a QRE, the row and column players play S with respective probabilities 0.97 and 0.98.

By contrast, the experimental evidence indicates that beliefs are given by the spot marked with a cross (southwest of the QRE). Clearly, any explanation for such behavior must depart from equilibrium theories such as QRE and Nash. Goeree and Holt (2004) argue that mutual consistency of beliefs may be a reasonable assumption if players play repeatedly against each other and learn. However, many strategic interactions last for very few periods, or involve entirely new circumstances such as legal disputes, election campaigns, international conflicts, and 3G spectrum auctions. In such situations, mutual consistency of beliefs appears less reasonable.




Figure 13.35 Logit best responses for the game in Figure 13.34 when the payoff from (R, S) is $(-14, 4)$.

Source: Reprinted from *Games and Economic Behavior* 46: 365–82. Goeree, J. K. and Holt, C. A., “A model of noisy introspection.” Copyright © 2004, with permission from Elsevier.

Suppose that players do not have experience of playing against each other and each player believes that the other will play S (the safe strategy) with probability 0.7. The points marked with a star on the logit response functions in Figure 13.35, for each player, indicate the best response of each player on such a supposition. Allowing for one round of best responses to the response of the opponent leads us to the point marked with a cross in Figure 13.35 that is consistent with the experimental evidence. Of course, the beliefs in this case are not self-reproducing in equilibrium. However, moving away from the rational expectations assumption gives us a better fit with the evidence. Since the point marked with a cross is to the southwest of the QRE, and northeast of the Nash equilibrium point, this point is referred to as a *pull to the centre*. If, however, the logit response functions were downward sloping, then we would get a *push to the edge* phenomenon.

In order to deal with such behavior, Goeree and Holt (2004) relax the “mutual consistency of beliefs” assumption. They do so by using “iterated conjectures” in a noisy version (that introduces noise as in (13.101)) of the rationalizability notion introduced by Bernheim (1984) and Pearce (1984). Let

$$\pi_{iS}(q | \mu_0) = \frac{\exp(\mu_0^{-1} u_{iS})}{\exp(\mu_0^{-1} u_{iS}) + \exp(\mu_0^{-1} u_{iR})}$$

denote the logit probability that player i plays S, conditional on a belief that the opponent plays S with probability q and an error rate $\mu = \mu_0$. q represents a player’s *first order beliefs* about the other player. So we can recursively define higher order beliefs as follows.

1. ρ^0 represents a player’s choice probability ($\rho^0 \equiv \pi_{iS}(q | \mu_0)$).
2. ρ^1 represents a player’s first order beliefs about the strategy choice of the opponent ($\rho^1 \equiv q$).
3. ρ^2 represents a player’s second order beliefs. These are the beliefs of a player about the first order beliefs of the other player. And so on ...

The first order beliefs of player i respect the fact that player j is playing a logit best response, conditional on what player j thinks about the choice probability of player i . This implies that $q = \pi_{jS}(\rho^2 | \mu_1)$. This, the second level of introspection by player i involves his guess of what player j thinks about i ’s choice probability, so might involve a greater degree of error relative to first order beliefs, hence, we would reasonably expect that $\mu_1 > \mu_0$. In general, as we consider increasingly higher order beliefs, we get increasing error rates, $\mu_0 < \mu_1 < \mu_2 < \mu_3 < \dots$. One suggestion made by the authors is the geometrically increasing case, $\mu_k = t^k \mu_0$, $t > 1$.

Several other models can be recovered from the NI model. For instance, setting $\mu_0 = \mu_1 = \mu_2 = \mu_3 = \dots = \infty$, all pure strategies are played with an identical probability (check in this case that $\pi_{iS} = 1/2$ from (13.101)). This is like a level- k model with type $L0$. Now consider $\mu_0 = 0$ and $\mu_1 = \mu_2 = \mu_3 = \dots = \infty$. This corresponds to a type $L1$ in a level- k model. In this manner, all the types in a level- k model can be recovered.

Rationalizability in the sense of Bernheim (1984) and Pearce (1984) requires that players do not play those strategies that are never a best response to a plausible set of beliefs about the play of the opponents. The final step in the model of NI is to extend this notion to modify best response to a noisy best response, as in (13.101), when beliefs are constructed iteratively, as explained above.

13.7.6 A brief note on some other models

Colman et al. (2014) introduce a set of twelve, 3×3 and 4×4 , games that distinguish between three different solution concepts—level- k models, team reasoning, and *strong Stackelberg reasoning*. The last of these concepts, due to Colman and Bacharach (1997), may not be familiar to many economists. Stackelberg reasoning implies that even in static normal form games, players believe that their opponents will play their best replies to their strategies. Taking into account this one-step behavior of others, players then choose their optimal strategies. The best replies of others may not be unique which limits the domain of such reasoning. Colman et al. (2014) strengthen the concept to strong Stackelberg reasoning that restricts the concept to games where the best replies of others are unique.

Figure 13.36 shows one of the two player, 3×3 experimental games in Colman et al. (2014). Consider the predictions made by the three alternative theories for the row player. Start with level- k models. Suppose that a type $L0$ player randomizes uniformly over the column strategies. Thus, the optimal strategy for a type $L1$ row player is B. Type $L2$ row players believe that the column players are of type $L1$. A type $L1$ column player plays C on the supposition that the row player is of type $L0$. Thus, the best response by a type $L2$ row player is B. Types are truncated at $L2$.

Team reasoning (see Bacharach, 1999, 2006) is a relaxation of the assumption of methodological individualism. Its central insight is that in certain contexts, individuals switch off their preference for selfish own-payoff maximization and choose instead to maximize the payoff of the team. For instance, individual players in a football team may pursue the team objective of winning the game. Of, course players may wish to maximize the team outcome only if they believe that other players will do likewise. Suppose that players maximize the average payoff, as a proxy for the team payoff. In Figure 13.36, the unique outcome under team reasoning is (C,C).

A row player who uses Stackelberg reasoning will choose A because the best response by the column player of the strategies A, B, and C by the row player is, respectively, A, B, and C. Likewise under such reasoning, the column player will also choose A. Thus, the unique outcome under Stackelberg reasoning is (A, A).

The main experimental findings over two sets of experiments (the second set was used to check the robustness of results) are as follows. The choices of players are consistent with type $L1$ players in 37–46% of the cases in 3×3 games and in 46–50% of the cases in 4×4 games.⁸⁵ The percentage of choices consistent with team reasoning increased from 13–20% in 3×3 games to 26–27%

	A	B	C
A	3, 3	1, 0	0, 2
B	1, 1	2, 2	3, 0
C	2, 0	1, 2	2, 5

Figure 13.36 An experimental game from Colman et al. (2014).

Source: Colman, A., Pulford, B., and Lawrence, C. "Explaining strategic coordination: Cognitive hierarchy theory, strong Stackelberg reasoning, and team reasoning." *Decision* 1(1): 35–58. Copyright © 2014, American Psychological Association, reprinted with permission.

⁸⁵ The percentage of choices across all three types of theories can sum up to greater than 100 because the competing theories make identical predictions in some games.

in 4×4 games. Based on raw frequencies, the percentages of choices that are consistent with Stackelberg reasoning in 3×3 games is 38 (the corresponding figures for $L1$ types and team reasoning are 43 and 29). However, there is a significant drop in choices that are consistent with Stackelberg reasoning in 4×4 games to 18% (corresponding figures for type $L1$ players and team reasoning are 46 and 27). The authors conjecture that the higher cognitive effort required in 4×4 games relative to 3×3 games hampers the use of Stackelberg reasoning. Furthermore, it seems that when games get cognitively more challenging, Stackelberg reasoning is abandoned in favour of level- k and team reasoning. Exploring the foundations of such heuristic switches is an exciting and unexplored research area.

We now briefly note a few other equilibrium concepts. Yildiz (2007) considers the possibility of *wishful thinking* in which players partition the set of strategy profiles, then pick their own strategy that corresponds to the highest possible payoff. This is an extreme form of optimism. Halpern and Pass (2012) introduce the concept of *iterated regret minimization*. As the name indicates, this involves the iterated deletion of strategies that do not minimize regret. The authors apply their concept to a range of games where the standard Nash equilibrium concept has proved problematic in explaining the empirical results, such as the traveller's dilemma, centipede games, and Bertrand competition. Cubitt and Sugden (2011) offer another iterative procedure, called the *reasoning-based expected utility procedure*. This procedure typically deletes more strategies than the standard iterated deletion of strategies.

13.8 Behavioral economics of auctions

This section considers the evidence from auctions of a single unit of an object and theoretical attempts to explain the evidence. Our main focus is on common value auctions.

13.8.1 Some basic results in auction theory

We begin by establishing some terminology, definitions, and some basic results in auctions for a single unit of an item. Multi-unit auctions have become more empirically important since the 1990s and have been an important area of theoretical and empirical research. However, for the experimental results on auctions and the behavioral models that we report, the single unit auctions will mostly suffice. The purpose of this section is to build some intuition about the basics of the economics of auctions at a heuristic level.⁸⁶

Suppose that we have n bidders, $i = 1, 2, \dots, n$, for a single indivisible object in an auction. Bidder i observes a private signal, X_i , and some public signal S that is informative about the value of the object (X_i and S can also be vectors). Let the vector of private signals be denoted by $X = (X_1, X_2, \dots, X_n)$. In general, the value of the object for bidder i , v_i , depends on all private signals and the public signal, i.e.,

$$v_i = v_i(X, S).$$

The two main cases considered in the experimental literature are as follows.

1. Independent Private Values (IPV): In this case, bidder i derives utility solely from his own private signal. Furthermore, the private signals are independent, so

⁸⁶ Readers interested in more formal treatments can consult several possible sources, e.g., Krishna (2010) and Milgrom (2004).

$$v_i = v_i(X_i).$$

All bidders know the probability distribution that generates the private signals of all competitors.

2. Common Values (CV): In this case, bidders do not use (or receive) private signals and their value for the object depends solely on the common public signal received, so, for any bidder i ,

$$v_i = v_i(S).$$

AUCTIONS WITH INDEPENDENT PRIVATE VALUES (IPV)

In the *first price sealed-bid auction*, all bidders simultaneously submit bids, $b_1, b_2, \dots, b_n$ and the highest bidder wins. Some tie breaking rule is invoked if there is more than one winning bid, for instance, one among the winning bids may be randomly picked and given the object. The winner pays a price, p^* , equal to the highest bid, so $p^* = b^* = \max \{b_1, b_2, \dots, b_n\}$. Hence, the utility of the winner is given by

$$v_i(X_i) - p^*.$$

In the Vickery or *second price sealed-bid auction*, the winner is the bidder who makes the highest bid. However, the price paid, p^s , equals the second highest bid, b^s . Hence, the utility of the winner is given by

$$v_i(X_i) - p^s.$$

Two important auction formats in actual practice are the English (or ascending) price auction and the Dutch (or descending) price auction. As the names indicate, the English price auction begins with some reserve price, r . The auction price is then gradually raised in discrete steps until only one bidder, who makes the winning bid, remains active. In the Dutch auction, on the other hand, one starts from a high price. The price is then lowered in discrete steps until one of the bidder cries “Mine” who is then declared the winner.

In order to motivate the experimental results, we now give some heuristic insights into some of the properties of the four kinds of auctions that we identify above (first price, second price, English, and Dutch); these results are summarized in Proposition 13.45, below.

Consider, an English auction with the IPV assumption. Starting at some reserve price, an auctioneer elicits increasingly higher bids. Suppose that the value of the object to a bidder is v_i . Then the bidder continues to increase his bids until the time when the bid b_i equals his value for the object, v_i . Any further increase in bids forces him to drop out. Hence, bidding one’s valuation is a dominant strategy.

Now consider a second price sealed-bid auction with the IPV assumption. We claim that it is a dominant strategy for bidder i to bid $b_i^* = v_i$. Suppose that all other bidders bid their valuation, i.e., $b_j^* = v_j, j \neq i$. Suppose $b_i < v_i$, then there is some probability that another individual might bid an amount $b_j^* : b_i < b_j^* \leq v_i$, and win. Since, bidder i can increase profits from bidding his true valuation, hence, $b_i < v_i$ is dominated by $b_i^* = v_i$. Similarly $b_i > v_i$ is dominated by $b_i = v_i$, because there is then some probability that bidder i will win and incur a loss. Thus, in the English and the second price auctions, each bidder bids his/her true valuation, the identity of

the winner is identical and the seller derives the same revenue. This is stated in Remark 13.5b(i), below.

The following example from Milgrom (1989) conveys some of the essential insights on first price sealed-bid auctions under the IPV assumption.

Example 13.23 Suppose that there are two risk-neutral bidders, bidding for one indivisible unit of an object. The first bidder has a valuation of $v_1 = 101$, which is public knowledge. The other bidder's valuation, v_2 , is private information but it is known to be distributed as follows: v_2 equals 50 with probability $\frac{4}{5}$ and 75 with probability $\frac{1}{5}$. Assuming common knowledge of rationality and one dollar as the lowest unit of currency, if bidder 1 bids 76 then he wins for sure, making a profit of $101 - 76 = 25$. If he bids $b_1 = 51$, then his expected payoff is $\frac{4}{5}(v_1 - b_1) = 40$. Thus, in his optimal bid, player 1 wins with probability $\frac{4}{5}$, and bidder 2 wins with probability $\frac{1}{5}$. The two important insights are the following.

- (1) The outcome in this asymmetric information first price sealed-bid auction, unlike the English auction and the second price sealed-bid auction need not be efficient (because bidder 2 wins one out of every five times).
- (2) Bidders bid less than their valuation.

While the outcome in Example 13.23 is not efficient, under certain conditions, first price sealed-bid auctions can produce the efficient outcome. One of the exercises at the end of the chapter shows this to be the case.

Suppose that we are now interested in the question of which auction design, among the four, generates the highest amount of revenue for the seller. In the second price auction, bidders bid their valuation but pay a price to the seller that equals the second highest bid. In first price and Dutch auctions, bidders shade their bids below their valuation, however, they pay a price to the seller that equals their bids (see Example 13.23). Hence, a priori it would seem difficult to decide which auction generates more revenues for the seller. However, a remarkable result, the revenue equivalence theorem, states that under IPV, risk-neutral bidders, and a further set of auxiliary assumptions, all four auction formats yield the same expected revenues for the sellers and same expected payoffs to the buyers.⁸⁷

We now loosely summarize these observations in Remark 13.5. For the precise statement of the result and proofs, see Fudenberg and Tirole (1991), Krishna (2010), and Milgrom (2004).

Remark 13.5

- (a) [Revenue equivalence theorem] Under IPV, when bidders are risk neutral, the four auction formats—first price, second price, English and Dutch—yield the same social surplus and bidders' utility. Additionally, if the valuations are iid, the highest bidder wins in equilibrium, and the lowest bidder gets zero profits, then, the expected payoffs of the seller are identical.
- (b) [Equivalence among auction formats] Under IPV and risk neutrality of bidders, the following two results hold.
 - (b1) The second price auction and the English auction induce identical bidding behavior and revenue.
 - (b2) The first price auction and the Dutch auctions induce identical bidding behavior and revenue.

⁸⁷ See, for instance, Vickrey (1961), Myerson (1981), and Riley and Samuelson (1981).

AUCTIONS WITH COMMON VALUES (CV)

Consider a common value (CV) auction of a single indivisible unit of an object. The most important result that has attracted a large body of empirical work and behavioral explanations is the *winner's curse*, which was first discovered in oil lease auctions. In these auctions, bidders bid for drilling rights for oil on certain tracts of land. The value of the oil is the same for all companies, so in this sense this is a common value auction. However, companies might differ in their estimate of the actual cost, c , of drilling for oil. An estimate of the actual cost by company $i = 1, \dots, n$ is given by $C_i = c + \varepsilon_i$, where the term ε_i is independently and identically distributed across companies and has zero mean.

It follows that the cost estimate of each bidder is unbiased, i.e., $E[C_i] = c$. Since bidders are alike in all other respects in a CV auction, the winner is the most optimistic bidder who has the lowest realization of ε_i . Thus, an unpleasant surprise might await the winner. One of two things might happen when actual costs turn out to be higher than the estimated costs. (1) The bidder might go bankrupt, or (2) the profits might turn out to be lower than expected. These two cases are sometimes known as versions 1 and 2 of the winner's curse.

The source of the problem that leads to the winner's curse is that bidders do not correctly calculate the value of an object or the cost, *conditional on their bid being successful*. Since auctions involve complicated calculations of beliefs about rivals, their rationality, and considerations of strategic uncertainty, it might be useful to illustrate the issues involved in a simpler game. Consider the "Acquire a company" game introduced by Samuelson and Bazerman (1985).

Example 13.24 [*Acquire a company or Takeover game; Samuelson and Bazerman, 1985*]
An acquiring firm (A) bids an amount b to acquire a target firm (T). Firm A announces a bid, which firm T can either accept or reject. Thus, the action set of the target firm is $S_T = \{0, 1\}$, where 0 means reject offer and 1 means accept offer. The value of the firm, v , is private information to T. Firm A knows, however, that v is uniformly distributed in the interval $[0, 100]$. Furthermore, firm A also knows that if acquired, the value of firm T will increase to γv , $\gamma > 1$ (say, because firm A has a commonly known reputation to add value to its acquisitions). On its part, firm T agrees to the bid by firm A if $b \geq v$. How much should A bid?

When $\gamma = 1.5$, experimental results suggest that about 90% of the subjects in their role as firm A, bid in the range 50–75.⁸⁸ The reasoning revealed by subjects in experiments in their role as firm A, is typically as follows. "The average value of the company is 50. By acquiring it, the value increases by 50%, so that post-acquisition, the expected value is 75." The problem with this reasoning is that it ignores the adverse selection problem by making unconditional calculations; it does not calculate the correct value of the firm, conditional on the bid being accepted.

Here is how the optimal bid should be calculated. The expected profit of firm A, for any bid b , is

$$\pi(b) = \Pr(v \leq b) \int_0^b [\gamma v - b] dv, \quad (13.102)$$

where $\Pr(v \leq b)$ is the probability that the bid is no lower than the value of the firm, otherwise firm A's bid is not accepted. If a bid b is accepted, then it reveals that the value of firm T lies uniformly in the interval $[0, b]$, which explains the upper and lower limits on the integral.

⁸⁸ See, for instance, Samuelson and Bazerman (1985), Ball et al. (1991), and Tor and Bazerman (2003).

The optimal, or the rational bid is zero. Formally, since v is uniformly distributed in the interval $[0, 100]$, so $\Pr(v \leq b) = \frac{b}{100}$. Substituting this in (13.102), we get $\pi(b) = \frac{b}{100} \int_0^b [\gamma v - b] dv$. Evaluating this integral we get

$$\pi(b) = \frac{b^3}{100} \left[\frac{\gamma}{2} - 1 \right].$$

For future reference, we write the rational bid as

$$b^r = \begin{cases} 0 & \text{if } 1 < \gamma \leq 2, \\ 100 & \text{if } \gamma > 2. \end{cases} \quad (13.103)$$

The intuition is as follows. If a bid b is accepted, then firm T is worth no more than b if run by the management of firm T. Since the value of firm T is uniformly distributed, each value of firm T below b is equally likely, so, on average, firm T is worth no more than $b/2$. Under the management of firm A, however, the company is worth $\gamma \frac{b}{2}$. For any $\gamma \leq 2$ the expected value of the company is below the bid value, so it is optimal to bid zero.

In contrast to the theoretical prediction of a zero bid, when $\gamma \leq 2$, most subjects in experiments bid between 50 and 75. In other words, they face a *winner's curse*.

13.8.2 Some empirical evidence on the predictions of the rational auctions model

We briefly highlight some of the experimental findings with a view to providing the reader with a flavor of the findings.⁸⁹ Our main focus is on the winner's curse. However, we will first review some findings from the IPV auctions. We will abbreviate the *risk-neutral Nash equilibrium*, for private or common value auctions, by "RNNE."

The predictions of the RNNE in Remark 13.5(b1), (b2) are violated by the experimental findings. Kagel et al. (1987) find that in second price auctions, the bids average 11% more than the dominant strategy, which is to bid one's valuation. Kagel and Levin (1993) find that 62% of the bids are above the dominant strategy and 8% are below it. Furthermore, while the bids converge rapidly to the dominant strategy equilibrium in the English auctions, this is not the case in second price auctions.

There is also evidence of overbidding in first price sealed-bid auctions. Coppinger et al. (1980) give evidence that contradicts Remark 13.5(b2)—the equivalence between first price and Dutch auctions. Goeree et al. (2002) find that for the case of two bidders with six discrete possible values of an object, bidders overbid relative to the predictions of the RNNE in all treatments. Neugebauer and Selten (2006) look at various treatments that are differentiated by the extent of information, e.g., bids of players and the identity of the winner. They also find evidence of overbidding in a first price sealed-bid auction. In one treatment, 75% of the bidders bid above the RNNE, while in another, 41%–48% of the bidders overbid relative to the RNNE.

The tendency to overbid relative to the RNNE does not necessarily vanish with experienced bidders. Garratt et al. (2012) conduct experiments with individuals who are experienced participants in eBay auctions, which have a similarity with second price auctions (although these two auctions are not identical). In experiments that expose experienced bidders on eBay to a standard

⁸⁹ Readers interested in the details can refer to the survey article by Kagel and Levin (2010) that we draw on in this section.

second price sealed-bid auction, 21.2% bid their value, 41.3% under-bid, and 37.5% over-bid. Subjects who had experience with selling on eBay, bid lower as compared to subjects with buying experience on eBay.

The RNNE performs better if one is primarily interested in the signs of the comparative static effects. For instance, an increase in the number of bidders is predicted to increase equilibrium bids. This finds support in the experimental results (Kagel and Levin, 1993). There is also evidence that bidders bid closer to their valuations in second price auctions relative to first price auctions (Kagel and Levin, 2010).

We now turn to the evidence from common value auctions. The first field account of the winner's curse came from three engineers Capen, Clapp, and Campbell (1971). They examined the outcomes of the US Federal government's outer continental shelf oil lease auctions, which have the characteristic of CV auctions. They note that winners of auctions were consistently disappointed with the rates of return. There was wide variability in the bids. Thaler (1988) highlights the following findings from this study and the follow-up study by Mead et al. (1983) that suggests a winner's curse. (1) Restricting oneself to serious bids, the ratio of the highest to the lowest bids is commonly between 5 to 10 and as high as 100. (2) The sum of winning bids in the 1969 Alaska North Slope sale was \$900 million, while the sum of the second highest bids was \$370 million. Furthermore, the winning bid was at least four times the next best bid in 26% of the auctioned tracts and at least two times the next best bid in 77% of the cases. (3) Calculating before-tax returns on 1223 leases over the period 1954–69 in the Gulf of Mexico, only 22% of the leases turned out to be profitable, but the after-tax earnings were not very high; 62% of the leases were dry and another 16% were unprofitable.

Bazerman and Samuelson (1983) provide the first experimental evidence for the winner's curse, using as their subjects, MBA students at Boston University. They conducted 48 first price sealed-bid auctions in 12 classes. In each auction, the number of bidders varied between 4–26. The object for sale was a jar full of coins or nickels or small/large paper clips with a value of \$8. Although all bidders observed the same object, they did not know the exact value of the items in the jar and had to guess the value. Hence, the setting was designed, as far as possible, to be a CV auction. The bid that was closest to the actual value of the item won a \$2 prize. The mean bid across all auctions was \$5.13, which is \$2.87 below the value of the object. This probably reflects risk aversion. However, in contrast to the average overall bid, the average winning bid was \$10.01, which is about \$2 above the actual value. Furthermore, the average winning bid led to losses in half of the auctions. The authors then followed up this game with the Acquire a company game in Samuelson and Bazerman (1985). In this game, the rational bid is zero, however, 90% of the subjects bid between 50–75. These results are suggestive of a winner's curse.

Similar findings of a winner's curse are documented for above-average ability students by Lind and Plott (1991) and also in variations of the game where one bidder knows the value of the object with certainty (Dyer et al., 1989). The finding that inexperienced subjects suffer from a winner's curse and fail to adjust adequately to avoid the winner's curse has been replicated in many other studies. Dyer et al. (1989) ran first price sealed-bid auction experiments with inexperienced subjects and find that only 17% of the winners earn any profits. It also seems that one cannot ascribe these findings to bad luck alone. Examining average and high bids as separate groups, they find that 59% of all bids and 82% of the high bids were greater than the expected value of the object, conditional on winning it.

Explanations of the winner's curse have also been given in terms of “limited liability” or the “joy of winning.” The former explanation, given by Hansen and Lott (1991), argues that the subjects in experiments do not lose more than the cash balances that they start off with, hence, limited liability that is inherent in experimental setups, promotes more aggressive bidding. The

joy of winning explanation simply supposes that subjects get a thrill (or extra utility) from winning the auction, which predisposes them to bid more aggressively. Both these explanations have been countered by experimental results.⁹⁰

The predictions of the standard auction model have been falsified in two other respects. First, theory predicts that in CV auctions, as the number of bidders increases, optimal bids should fall. The intuition is that it is even worse to win in a larger group because the winner's signal of the cost C is then such that the measurement error, ε_i , is more extreme. In contrast, experimental results show that bids do not fall in response to an increase in the number of bidders (Kagel et al., 1995). Second when public information is provided to subjects in experiments about the value of the object, it reduces revenue when a winner's curse is operative. The reason is that bidders subject to a winner's curse bid too aggressively. However, when provided with the information about the true value of the object, they revise their bids downwards. However, theory predicts the opposite. This is because in first price auctions, individuals bid below their valuation and when the true value is publicly announced they should revise their bids upwards. The experimental findings are contrary to the predictions of theory (Kagel and Levin, 1986).

Experienced auction players in lab experiments also exhibit the winner's curse (Dyer et al., 1989). For moderately experienced bidders, Kagel and Levin (1986) find that there is adjustment that takes account of the winner's curse through a "hot stove" learning process. Such bidders reduce their bids with experience, and are able to avoid the winner's curse. However, they do not seem to correctly calculate the expected value of an object, conditional on winning it. When bidders are extremely experienced, greater public information increases revenues in sealed-bid auctions, in line with the theory's predictions. Nevertheless, the profits of these highly experienced bidders still fall below those that would be consistent with the RNNE. The reason is that they bid far too aggressively relative to the predictions of RNNE, although they avoid the winner's curse (Kagel and Richard, 2001). However, Kagel and Levin (2001, pp. 32, 33) caution us in interpreting these results by arguing that the winner's curse is avoided in experimental settings only with strong information feedback that is unlikely to be present in field settings. There are additional reasons, they point out, relating to the incentive and organizational structure, to suspect that other factors in field settings might work in favor of producing the winner's curse rather than removing it.

Harrison and List (2008) conduct two experiments with sports cards dealers; see also Harrison and List (2004). In the first experiment, they use a symmetric and an asymmetric information procedure in which bidders bid for an object. Under symmetric information, all dealers know that the value of the object is drawn randomly from a uniform distribution with a known midpoint. Under asymmetric information, one of the dealers, the insider, knows the objectively true value of the item with certainty, while others believe that it is randomly drawn from a known distribution. Dealers do not suffer from the winner's curse under symmetric information, but under asymmetric information they suffer from the winner's curse in 24–30% of the cases. The authors interpret this result as evidence that in familiar environments (symmetric information case) dealers protect themselves from the winner's curse. In a second experiment, dealers bid on a familiar item, a package of ten sports cards with a known retail value of \$9–\$10. Dealers do not exhibit the winner's curse in this case. In both experiments, more non-dealers than dealers exhibit the winner's curse.

⁹⁰ Kagel and Levin (1991) re-examined the data in Kagel and Levin (1986) and showed that the limited liability argument is not supported and the joy of winning explanation has been rejected by the experimental findings of Holt and Sherman (1994) and Avery and Kagel (1997).

Kagel (2015) argues persuasively that there are alternative and plausible explanations for the results from both experiments reported above. In the first experiment, dealers could simply be using their common heuristic of buying at a low price to sell at a higher price, which protects them from the winner's curse. In the second experiment, there is no scope for a winner's curse because the package of ten sports cards has a publicly known retail value.

The winner's curse has been documented in a variety of settings and not just auctions (see Kagel and Levin, 2010 for the references). Experiments that simultaneously include IPV and CV factors are relatively rare. In one such set of experiments, conducted by Goeree and Offerman (2002), actual bidding behavior is intermediate between the RNNE bids and the naive benchmark in which subjects do not take account of the expected value of winning, conditional on winning the object. The outcome is still relatively efficient because most bidders are subject to the winner's curse, so relatively, the bidder with the highest valuation wins.

13.8.3 *Cursed equilibrium as an explanation of the winner's curse*

We now give an explanation for the winner's curse using the χ -cursed equilibrium of Eyster and Rabin (2005) in Section 13.7.3. Consider the Takeover game in Example 13.24 above, for the case of a fully cursed player, $\chi = 1$. Recall that players in a cursed equilibrium know the distribution of types of the opponent, but ignore the correlation between types and actions. In the context of the Takeover game, types of the target firm, T, are distributed uniformly over $[0, 100]$.

If the acquiring firm A is not cursed, then it knows that the target firm will accept any bid $b \geq v$ and reject any bid $b < v$. Thus, it would know that the probability that a bid b is successful is $\frac{b}{100}$. However, a fully cursed firm A assigns the average action to all types of the target firm. Recalling from Example 13.24 that firm T can either accept or reject the bid ($S_T = \{0, 1\}$), the average strategy of firm T, using (13.98), is $\bar{\sigma}_T = \int_0^b \frac{1}{100} dv + \int_b^{100} \frac{0}{100} dv = \frac{b}{100}$. Thus, a fully cursed firm A believes that any bid b will be accepted with probability $\frac{b}{100}$; if the bid is accepted, the utility of firm A is $\gamma v - b$. These beliefs are not correct because any type of firm T will reject the bid if $b < v$. Thus, a fully cursed firm A maximizes

$$\pi(b; \chi = 1) = \frac{b}{100} \int_0^{100} [\gamma v - b] dv. \quad (13.104)$$

Evaluating the integral, we get $\pi(b; \chi = 1) = 50\gamma b - b^2$. Therefore, the optimal fully cursed bid is $b^c = 25\gamma$. Let us compare it to the optimal bid, b^r , in the absence of cursedness, given in (13.103). When $1 < \gamma \leq 2$, $b^r = 0$ but $b^c \in [25, 50]$, thus, in this case, the fully cursed bidder bids too high. This is an example of the winner's curse. The fully cursed bidder imagines that even when $b < v$ some types of firm T will continue to sell to her and this prevents firm A from lowering the bid to the optimal bid of zero. However, when $2 < \gamma < 4$, $b^r = 100$ but $b^c \in (50, 100)$. In this case, the fully cursed bidder bids too low and fails to realize that by raising the bid, all types of sellers will sell with probability 1.

What if $0 < \chi < 1$? In this case, the profit of a χ -cursed firm A is given by

$$\begin{aligned} \pi(b; \chi) &= \int_0^b \left[(1 - \chi) \times 1 + \chi \frac{b}{100} \right] [\gamma v - b] \frac{1}{100} dv \\ &\quad + \int_b^{100} \left[(1 - \chi) \times 0 + \chi \frac{b}{100} \right] [\gamma v - b] \frac{1}{100} dv. \end{aligned}$$

All types of firm T sell with probability 1 if $b \geq v$. Hence, in the first term, a χ -cursed firm A believes that with probability $(1 - \chi)$ firm T will accept and with probability χ it plays the average strategy $\bar{\sigma}_T = \frac{b}{100}$. But when $b < v$ it knows that with probability $(1 - \chi)$, firm T will not buy (i.e., choose the action $s_T = 0$) and with probability χ it plays the average strategy $\bar{\sigma}_T = \frac{b}{100}$. In each case, the payoff to firm A is $\gamma v - b$ if the bid is accepted.

Example 13.25 [100VT game and shifted 100VT game; Charness and Levin, 2009]: Consider the Takeover game in Example 13.24. Modify it to eliminate firm T in a robot design as follows. A single decision maker, firm A, chooses a bid, b , and then chooses one of 100 cards, numbered 1–100, that are placed face down on a computer screen. If the bid is at least as great as the number, v , on the card, then firm A gets $1.5v - b$ (i.e., $\gamma = 1.5$). Otherwise, the bidder gets zero. The rational bid can be found as in (13.103) and, since $\gamma \leq 2$, it equals $b^r = 0$. The equilibrium bid in a cursed equilibrium maximizes (13.104), so $b^c = 37.5$. Averaging across all 100VT experiments in Charness and Levin (2009), the prediction of the cursed equilibrium is close to the actual values in experiments. Charness and Levin (2009) also consider the shifted 100VT game in which the value of firm T, lies uniformly in the interval $[20, 120]$. In this case, the cursed equilibrium is given by the solution to

$$b^c \in \arg \max \int_{20}^{120} [1.5v - b] \frac{b}{100} dv.$$

A simple calculation shows that the solution is $b^c = 52.5$. This is consistent with the experimental results of Charness and Levin (2009). One might wonder if the robot design in which other players are absent (and so beliefs about other players are absent), is a valid application of a cursed equilibrium? Charness and Levin (2009) interpret Eyster and Rabin's (2005, p. 1632) cursed equilibrium to imply that "in a less literal sense, the notion of cursed equilibrium is not that each player is certain that every other player's action does not depend on her type but rather that they don't look for the correlations this mechanism implies." Crawford et al. (2013, p. 33) disagree with this view. They write: "However, we doubt that such mechanical generalizations would be empirically very useful, because we suspect people use different cognitive routines in games against real other players than in decision problems based on theoretical other players."

Example 13.26 [2VT game; Charness and Levin, 2009] This problem reduces the complexity of the Bayesian updating process in the Takeover game to minimize the scope of other explanations that rely on bounded rationality. Suppose that the value of firm T is equally likely to be 0 or 99 and $\gamma = 1.5$. A rational bidder knows that a bid $b \geq 99$ will be accepted with probability 1 and a bid $0 \leq b < 99$ will be accepted with probability $1/2$. Thus, his expected payoff from bidding $b = 99$ is $\frac{1}{2}(1.5 \times 99 - 99) + \frac{1}{2}(1.5 \times 0 - 99) = -24.75$, while from bidding $b = 0$ it is $(1.5 \times 0 - 0) = 0$. Hence, the rational bid $b^r = 0$.

Suppose that the bidder is fully cursed, so $\chi = 1$. Such a bidder imagines that any bid $b < 99$ is accepted with probability 0.5, the average action of the two likely types of firm T. There is no point in bidding above 99, so by bidding $b = 99$ his payoff is $-24.75 < 0$. On the other hand when he decides to bid below 99, then there is no point in bidding anything more than zero, so his payoff in this case is $\frac{1}{2}(1.5 \times 99 - 0) + \frac{1}{2}(1.5 \times 0 - 0) = 74.25 > 0$. Hence, a fully cursed bidder bids zero, so $b^c = 0$. In this case, the cursed bid coincides with the rational bid. Averaged across the experimental sessions over 60 periods, the experimental results of Charness and Levin (2009) show that between 30–60% of the bids are zero, reducing the extent of the winner's curse relative to the 100VT treatment. The average bid was, however, relatively




Figure 13.37 Actual and fitted bids as a function of the signals.

Source: Eyster and Rabin (2005) with permission of The Econometric Society.

higher in the 2VT design implying that subjects lose money when the rationality requirement is relaxed.

Avery and Kagel (1997) consider two-bidder, second price, common value auctions. The value of the object is not known but each bidder privately receives a signal z_i drawn uniformly from the interval $[1, 4]$. The value of the object to each bidder is $v = z_1 + z_2$. It can be shown that a χ -cursed bidder bids $b_i(z_i) = 2.5\chi + (2 - \chi)z_i$. When $\chi = 0$, $b_i(z_i) = 2z_i$ and when $\chi = 1$, $b_i(z_i) = 2.5 + z_i$, where 2.5 is the average signal of the other player. Figure 13.37 shows a plot of b_i versus z_i for experienced subjects⁹¹ that is much closer to the theory's prediction relative to the plot for inexperienced subjects. Figure 13.37 also shows the various predictions of the model for $\chi = 0, 1, 2$ as well as an OLS fitted curve. Recall that the case $\chi = 0$ gives the Bayesian Nash equilibrium; this is shown as the bold curve in the figure.

In general there is no presumption that cursedness alone explains the bidding behavior of subjects in experiments. The behavior of subjects is much richer than can be explained by cursedness alone. For instance, some subjects play weakly dominated strategies, there appears to be heterogeneity in cursedness across subjects, and subjects play with significant noise around the predicted cursed equilibrium strategies. Relative to inexperienced subjects, experienced subjects are less likely to play weakly dominated strategies and they exhibit lower heterogeneity. However, in both cases, focusing on the Bayesian Nash equilibrium ($\chi = 0$), subjects bid above their best response for low signals and below their best response for high signals. Such behavior is better explained when bidders are cursed, $\chi > 0$. The best fitting value of χ , assuming that subjects play normally distributed noisy best responses is 0.63. In sum, while cursedness alone cannot explain the experimental data on auctions, it fits better than a Bayesian Nash equilibrium.

⁹¹ These are subjects who have experience of 7 practice rounds and 18 actual rounds of the experimental game.

13.8.4 A level- k explanation of the winner's curse

Crawford et al. (2009) and Crawford and Iriberri (2007b) give a level- k analysis of the auction of a single unit of an object with two risk-neutral bidders.⁹² The valuation of each bidder, v_i , $i = 1, 2$, is private information. Each v_i is independently and identically distributed with distribution function $F(v)$ and density $f(v)$. It is useful to distinguish between the two cases of increasing and decreasing densities, subscripted by the alphabets I and D respectively. The increasing density case is given by

$$F_I(v) = v^\gamma; f_I(v) = \gamma v^{\gamma-1}; v \in [0, 1]; \gamma > 1. \quad (13.105)$$

The assumption $\gamma > 1$ ensures that the density is increasing in v . The decreasing density case is given by

$$F_D(v) = \frac{\beta}{\beta - \alpha} \left(1 - \frac{\alpha}{v}\right); f_D(v) = \frac{\alpha\beta}{(\beta - \alpha)v^2}; v \in [\alpha, \beta]; 0 < \alpha < \beta. \quad (13.106)$$

It is obvious that $f'_D(v) < 0$. To derive the benchmark results for the case of fully rational bidders, we first check for Myerson's (1981) *regularity condition*, which requires that $v - \frac{1-F(v)}{f(v)}$ is increasing in v (Fudenberg and Tirole, 1991). For the two cases in (13.105), (13.106), we have

$$v - \frac{1 - F_I(v)}{f_I(v)} = \frac{(1 + \gamma)v}{\gamma} - \frac{1}{\gamma v^{\gamma-1}} \text{ and } v - \frac{1 - F_D(v)}{f_D(v)} = \frac{v^2}{\beta}.$$

Noting that $\gamma > 1$, the distributions satisfy the regularity conditions in both cases. Myerson (1981) showed that if the regularity conditions hold, then the first price auction belongs to the class of optimal mechanisms. For this reason, we will now focus on first price auctions only. We begin by stating the results for the case of fully rational bidders; the proof is omitted because it is standard and readily available (Myerson, 1981; Krishna, 2010). Denote the bids in the increasing and the decreasing density cases by B and b respectively; a superscript star denotes the optimal bid.

Proposition 13.45 (Fully rational bidders benchmark; Crawford et al., 2009):

- (a) *Increasing density case: Consider a first price auction with a reserve price $\rho \in [0, 1]$. For any realization of value $v \geq \rho$, the optimal bid is given by*

$$B^*(v) = \frac{\rho^{1+\gamma}}{(1 + \gamma)v^\gamma} + \left(\frac{\gamma}{1 + \gamma}\right)v. \quad (13.107)$$

And the optimal reserve price is $\rho^* = (1 + \gamma)^{\frac{-1}{\gamma}}$.

- (b) *Decreasing density case: Consider a first price auction with a reserve price $\rho \in [\alpha, \beta]$. For any realization of value $v \geq \max\{\alpha, \rho\}$, the optimal bid is given by*

$$b^*(v) = \frac{v}{v - \alpha} [\rho - \alpha + \alpha (\ln v - \ln \rho)]. \quad (13.108)$$

And the optimal reserve price is $\rho^* = \alpha$.

⁹² We omit here an analysis of *zero-sum betting*. In such games, classical theory predicts that we should not observe any trade. Yet, experimentally, one observes a great deal of trade (Brocas et al., 2014). See Section 5 in Crawford et al. (2013) for more details on a level- k explanation of this phenomena based on a random specification of type L0 players that is independent of the realization of private information of other players.

With the benchmark results in place we can now undertake a level- k analysis of the first price auction model in the two cases, increasing and decreasing densities. It is not clear, a priori, how one might proceed to model the actions of the type L_0 players. Crawford et al. (2009) specify two possibilities.

1. Random L_0 : In this case, type L_0 players simply choose their bids randomly over the relevant support, and any L_k type, $k \geq 1$, is called a “random L_k type.”
2. Truthful L_0 : In this case, type L_0 players truthfully bid their valuations, and any higher L_k type, $k \geq 1$, is called a “truthful L_k type.”

Denote the bid by a type L_k player in the increasing and the decreasing density cases, respectively, by B_k and b_k , with a superscript star denoting an optimal value. For exposition purposes we shall assume that there are only two types of level- k players L_1 and L_2 . Thus, players of type L_0 exist only in the minds of type L_1 players as a model of the type of other players.

THE INCREASING DENSITY CASE WITH LEVEL- k PLAYERS

(a) Random L_k players: Random L_0 types choose their bid, B_0 , uniformly over the interval $[\rho, 1]$, $\rho \geq 0$; bidding over $[\rho, 1]$ appears consistent with the data rather than bidding over $[0, 1]$, hence, this is part of the specification of type L_0 players in the model. If the valuation, v , is lower than the reserve price, ρ , for any of the higher types, they will not bid a positive amount.

Type L_1 players believe that other bidders are of type L_0 . Thus, the perceived probability that the bid B_1 made by a type L_1 player wins is $\frac{B_1 - \rho}{1 - \rho}$. Given that the payoff of type L_1 is his private valuation, v , minus his bid, B_1 , the optimal bid of type L_1 solves the following expected utility problem.

$$B_1^* \in \arg \max \left(\frac{B_1 - \rho}{1 - \rho} \right) (v - B_1).$$

Thus, the optimal bid for a type L_1 player is

$$B_1^* = \frac{\rho + v}{2} \geq B^*(v). \quad (13.109)$$

Since we assume that $v \geq \rho$, so $B_1^* \geq \rho$. A type L_2 bidder believes that the other bidder is of type L_1 who bids according to (13.109). Hence, the probability that the bid of type L_2 , denoted by B_2 , wins, is $\Pr(B_2 \geq B_1^*)$. Using the fact that v is distributed over $[0, 1]$ with distribution function $F_1(v) = v^\gamma$, gives $\Pr(B_2 \geq B_1^*) = \Pr(B_2 \geq \frac{\rho + v}{2}) = \Pr(2B_2 - \rho \geq v) = (2B_2 - \rho)^\gamma$. Hence, the optimal bid of type L_2 , B_2^* , solves:

$$B_2^* \in \arg \max (2B_2 - \rho)^\gamma (v - B_2).$$

Solving the first order condition and noting that bidders will not bid if $v < \rho$, we get

$$B_2^* = \max \left\{ \rho, \frac{\gamma}{1 + \gamma} + \left(\frac{\gamma - 1}{1 + \gamma} \right) v \right\} \leq B^*(v).$$

Thus, type L_2 players underbid relative to the fully rational prediction.

(b) Truthful L_k players: We omit the results here as an exercise for the reader. The bids are typically greater than the fully rational bids, which is consistent with the winner's curse.

THE DECREASING DENSITY CASE WITH LEVEL- k PLAYERS

In this case, the distributional assumptions are given in (13.106). When the valuation exceeds the reservation price, $v \geq \rho$, the equilibrium bids of players are computed as follows.

(a) Random L_k players: In order to facilitate comparison with the fully rational case, assume, as in Proposition 13.45, that $\rho \in [\alpha, \beta]$. Random L_0 types choose uniformly over the interval $[\rho, \beta]$, $\rho \geq \alpha$. Type L_1 players believe that the other bidder is of type L_0 , thus, for any bid they make, B_1 , their probability of winning is $\frac{B_1 - \rho}{\beta - \rho}$. The optimal bid of a type L_1 player, $\rho \leq B_1^* \leq \beta$ can be solved from

$$B_1^* \in \arg \max \left(\frac{B_1 - \rho}{\beta - \rho} \right) (v - B_1).$$

The optimal solution is

$$B_1^* = \frac{v + \rho}{2} \geq b^*(v), \quad (13.110)$$

where the optimal rational bid $b^*(v)$ is defined in (13.108).

Type L_2 players bid an amount B_2 and they believe that the other bidder is of type L_1 who bids according to (13.110). $\Pr(B_2 \geq B_1^*) = \Pr(2B_2 - \rho \geq v)$. Hence, using (13.106), the expected profits of a type L_2 are given by

$$\frac{\beta}{\beta - \alpha} \left(1 - \frac{\alpha}{(2B_2 - \rho)} \right) (v - B_2). \quad (13.111)$$

When $v \geq \rho$, the bid which maximizes (13.111) such that $B_2 \in [\rho, \beta]$ is given by

$$B_2^* = \frac{\rho}{2} + \frac{\alpha}{2} \sqrt{\frac{2v - \rho}{\alpha}} \geq b^*(v), \quad (13.112)$$

where $b^*(v)$ is defined in (13.108). From (13.110), (13.112), the bids by both types of players exceed the fully rational bids. This is consistent with the winner's curse.

(b) Truthful level- k players: We leave this case as an exercise. The main results are similar to the increasing density case.

In the decreasing density case, level- k players overbid relative to the fully rational players because their beliefs are anchored on the play of type L_0 players. Type L_0 players bid uniformly, thus, higher types overestimate the probability of high valuations relative to the actual density (which is decreasing). The intuition for underbidding in the increasing density case is simply the converse of this.

Charness and Levin (2009) test the level- k model and the cursed equilibrium model; these are jointly lumped by them as *belief-based models*. Consider, for instance, the 100VT and the 2VT experiments described in Examples 13.25, 13.26 above. In these games, a bidder faces a computer. Suppose that the bidder imagines the computer to be a type L_0 player and that the bidder is a type L_1 player in a 100VT or 2VT game. Then he believes that the opponent will choose each of her two actions (accept or reject bid) with equal probability, irrespective of the size of the bid. Hence, it is optimal to bid zero. We also saw in Example 13.26 that the cursed bid was also zero in this case. Thus, the level- k models and the cursed equilibrium models cannot explain the winner's curse in these experiments. Charness and Levin (2009) suggest that the winner's curse is caused

by cognitive processes that are not a feature of these two belief-based models. Thus, for instance, the winner's curse is lessened in their experiments when one moves from the 100VT treatment to the less cognitively challenging 2VT treatment. Other such cognitive factors that impact on the degree of the winner's curse in their experiments include the amount of information and instructions provided to subjects in experiments.

Ivanov et al. (2010) provide experimental evidence that casts doubt on belief-based models (they consider cursed equilibrium, level- k , and QRE models). They introduce the maximal game, which is a second price common value auction with n players. Each player in this game receives an iid signal, $x_i \in [0, 10]$, about the common value, V , of the object. The objectively correct common value is the highest value of the signal received, i.e., $V = \max_{i=1}^n x_i$. Players submit simultaneous bids, b_i , and the winner pays the second highest bid. Players can either underbid, $b < x$, overbid $b > x$, or bid the value of their own signal, $b = x$.

Starting from $b < x$ suppose that a bidder raises his bid to $b = x$. Suppose all other players bid equal to their signals, then (1) if at $b < x$ the bidder won the object, then he would still win it and pay the same price (second highest bid), and (2) if at $b < x$ the bidder does not win the object, then he receives a zero payoff. By raising the bid to $b = x$, there is some probability that he will win the object and receive a positive payoff. Hence, underbidding is weakly dominated by $b = x$ for any player. Analogously, the reader can check that $x < b \leq 10$ is also weakly dominated by $b = x$. Thus, bidding one's received signal, i.e., $b_i = x_i$ for all i is a Nash equilibrium in the maximal game; in fact it is the unique equilibrium.

The experimental results of Ivanov et al. (2010) show that the largest proportion of bids satisfies $x < b \leq 10$. Hence, there is a winner's curse. To check for the importance of beliefs, the authors run the experiments in two parts. In part I, bidders participate in two-player bidder auctions (where they suffer from the winner's curse). In part II, they play against a computer that plays according to the bid function of the same individual in part I. So, if the individuals are responding to beliefs, they should realize that they are playing against someone who overbids relative to the signal. Winning against such an individual conveys even worse news. So bids should either stay the same or come down. The experimental results show that individuals do not revise their bids downwards, continuing to bid in a manner that is similar to their bids in part I.

The authors conclude that their results cast doubts on belief-based models of auctions. However, this conclusion has been robustly challenged. Ivanov et al. (2010) select only a subset of subjects whose data is analyzed, based on a selection criterion. Costa-Gomes and Shimoji (2011) raise concerns over the subset of subjects whose data is analyzed. When they analyze the data for all subjects, they find that the overbidding in the part II result is relevant for only a third of the subjects; for half of these subjects, the data is in conformity with belief-based models. A similar method—non-exclusion of subjects—is also employed by Camerer et al. (2012) to re-analyze the data in Ivanov et al. (2010). They find strong support for belief-based models that include cursed equilibrium and cognitive hierarchy models. Furthermore, Crawford et al. (2013, p. 34) raise the objection that the Ivanov et al. data is very noisy with very low signal to noise ratio that can be used to reject any theory.

Crawford and Iriberri (2007b) run a three-horse race between level- k models, cursed equilibrium, and QRE using auction data from three datasets. These datasets are the first and second price common value auctions in Kagel and Levin (1986) and Avery and Kagel (1997), and the independent private values auctions in Goeree et al. (2002) (respectively KL, AK, and GHP). Their interest was in examining data from the early rounds of play.⁹³ Players are assumed to play with logistic noise; both subject-specific and type-specific noise is allowed.

⁹³ For instance, they used the data from AK for inexperienced subjects, while the Eyster and Rabin (2005) results that we examined above, based on the same dataset, were for experienced subjects.

The type distribution of level- k players is fairly stable across the three datasets, particularly when subject-specific noise is considered. This is true to a lesser degree under type-specific noise. With a few exceptions,⁹⁴ the frequency of level- k types is comparable to earlier research (see Section 13.4.5). The frequency of type $L0$ players was quite low—4% in KL first price auctions and 0% in AK second price and GHP first price auctions. For the GHP data with subject-specific noise, level- k models fit the data much better than QRE and can account for some of the deviations from equilibrium that GHP explain by invoking other factors such as “joy of winning” and “risk aversion.”

Interesting, for KL first price and AK second price auctions, when a cursed equilibrium model is used, the Eyster and Rabin (2005) results for experienced subjects are quite similar to the Crawford and Iriberri (2007b) results for inexperienced subjects. Different types of cursed subjects are estimated from the data assuming that $\chi \in [0, 1]$ and types go up in increments of 0.1. There are spikes in the data for estimated cursed types at $\chi = 0$ (equilibrium play) and $\chi = 1$ (fully cursed types). There is a roughly even distribution of types in between (except for spikes at $\chi = 0.2$ in KL second price and $\chi = 0.7$ in AK second price). Except for KL first price auctions, the level- k models fit the data better than cursed equilibrium models.

⁹⁴ In particular, in the KL second price auctions, the estimated frequency of type $L1$ players is too low (21%) and that of type $L2$ is too high (32%), relative to earlier estimates. The authors provide likely explanations for why this might be the case.

Exercise for Part 4

- Recall our discussion of the findings of the two-player game in Beard and Beil (1994) in Section 12.2.1. Goeree and Holt (2001) replicate their findings for an extensive form game shown in Figure 12.1, but with different payoffs as shown in Table 4.E1, which also gives a summary of the results from their replication (q is the actual probability with which a player 2 responds to player 1's choice of R).
 - Recall that p is the conditional probability (in the mind of player 1) of the play of r by player 2, following the play of R by player 1. Compute the threshold probability p_c such that for probabilities greater than p_c , it is optimal for player 1 to play R.
 - Interpret the results in Table 4.E1 and compare these findings to those of Beard and Beil (1994).

Table 4.E1 Goeree and Holt's replication of the Beard–Beil (1994) findings.

Treatment	Numbers of pairs	Payoffs from various strategy combinations			% of Player 1s choosing strategy L	q
		L, l and L, r	R, l	R, r		
1	25	70, 60	60, 10	90, 50	12	1.00
2	25	70, 60	60, 48	90, 50	32	0.53
3	15	80, 50	20, 10	90, 70	13	1.00
4	25	80, 50	20, 68	90, 70	52	0.75
5	25	400, 25	100, 348	450, 350	80	0.80

- (Sefton and Yavas, 1996). Consider the following symmetric two-player strategic form game:

	R	B
R	960, 960	960, 0
B	0, 960	1200, 1200

As in Section 12.2.4, consider an implementation of the strategy profile (B, B) through the Abreu–Matsushima mechanism. Consider three values of T : $T = 4$, $T = 8$, and $T = 12$. The fine is set at $F = 160$. Which of the three cases is implementable with this level of fine? The results in these three cases are shown in Figure 4.E1 with the percentage of all-B play plotted along the vertical axis. Interpret the results.

- Consider the following game from Goeree and Holt (2001).

	Left (26%)	Middle (8%)	Non-Nash (68%)	Right (0%)
Top (68%)	200, 50	0, 45	10, 30	20, -250
Bottom (32%)	0, -250	10, -100	30, 30	50, 40




Figure 4.E1 Percentage of (B, B) choices in various rounds.

Source: Reprinted from *Games and Economic Behavior* 16: 280–302. Sefton, M. and Yavas, A. "Abreu-Matsushima mechanisms: experimental evidence." Copyright © 1996, with permission from Elsevier.

Payoffs in the table are in cents and the percentage of row and column players who play each of the strategies is shown in brackets, next to each strategy.

- (a) Find all the Nash equilibria in the game. Why is the majority choice of strategy for the column players, Non-Nash, problematic for a Nash equilibrium?
 - (b) How can the strategy Non-Nash be explained as a maximin choice?
 - (c) How can we rule out loss aversion as a possible explanation of the observed play?
4. (Goeree and Holt, 2001) Consider the sender–receiver game in Table 4.E2. The sender has two possible types, A, B, each with a prior probability of 1/2. The receiver observes the sender's message, either Left, L, or Right, R, and then chooses one of three actions C, D, or E.
- (a) The experimental results indicate the presence of the following outcome. Type A senders send R and type B senders send L. The receivers always play C. Show how this separating equilibrium can be explained by the senders using the *principle of insufficient reason*, namely, that they think each response of the receiver is equally likely.
 - (b) A second possible equilibrium of this game is a pooling equilibrium. In this equilibrium, each sender sends the message L. On receiving a message L, the receiver plays C and on receiving the out-of-equilibrium message R, the receiver plays D. However, the experimental evidence in (a) indicates that the separating equilibrium was played. Does

Table 4.E2 A signaling game from Goeree and Holt (2001).

		Response to Left signal					Response to Right signal		
		C	D	E			C	D	E
type A	sends	300, 300	0, 0	500, 300	type A	sends	450, 900	150, 150	1000, 300
	Left					Right			
type B	sends	500, 500	300, 450	300, 0	type B	sends	450, 0	0, 300	0, 150
	Left					Right			

the pooling equilibrium in (b) survive the Cho and Kreps (1987) intuitive criterion? What can you conclude about the plausibility of the pooling criterion?

(c) Now consider the following change to the payoff matrix. The outcome 500,500 in the southeast corner of the payoff matrix is replaced by 300,300. The experimental results are that 10 of the 13 type A senders choose R and 9 of the 11 type B senders choose L. Hence, we seem to have a separating equilibrium. Is such a separating equilibrium also a perfect Bayesian Nash equilibrium (PBE) in classical game theory? If not then construct a PBE for this game. Are the experimental results consistent with your constructed equilibrium?

5. Consider the Van Huyck et al. (1997) continental divide game in Section 12.4.9. The payoffs are given in (12.7) and the payoff matrix is described in Table 12.17.

(a) Show that the results given in Figure 12.19 can be explained by either of the following two models. (i) A model of myopic best response dynamics in which players play a best response to the observed median in the previous period. (ii) A model of fictitious play (see Part 5 for details) in which players best respond, given the observed frequency of the medians so far. So if the j^{th} median value, M_j , has occurred f_j times, then player i maximizes the following payoff function

$$\sum_j f_j \pi(s_i, M_j),$$

where the payoff function $\pi(s_i, M_j)$ is defined in (12.7).

(b) Show that the myopic best reply dynamics converge within 5 time periods, while the fictitious best reply dynamics converge within 22 periods.

6. Prove Proposition 12.16.
 7. (Jung et al., 1994) Prove Proposition 12.18.
 8. Verify the equilibria claimed in experiments H and J in Table 12.29.
 9. Consider the three two-player games from Goeree and Holt (2001) shown in Figure 4.E2. Five cohorts of ten subjects were randomly matched to play these games. The percentage of players choosing a particular action in each game is shown in brackets, next to the strategy.
 (a) What is the mixed strategy equilibrium for the symmetric matching pennies game? Do you think that the subjects' choices are consistent with this equilibrium?

Symmetric matching pennies game		
	L (48%)	R (52%)
T (48%)	80, 40	40, 80
B (52%)	40, 80	80, 40
Asymmetric matching pennies game		
	L (16%)	R (84%)
T (96%)	320, 40	40, 80
B (4%)	40, 80	80, 40
Reversed asymmetry		
	L (16%)	R (84%)
T (96%)	320, 40	40, 80
B (4%)	40, 80	80, 40

Figure 4.E2 Three simple strategic form games and the choices of players.

Source: Reprinted from *American Economic Review* 91(5): 1402–22. Goeree, J. K. and Holt, C. A. "Ten little treasures of game theory and ten intuitive contradictions." Copyright © 2004, with permission from Elsevier.

(b) Now consider the next two games. The only difference in these games from the symmetric matching pennies game is that the row player's payoff from the strategy profile (T, L) changes (to 320 and 44 respectively). Notice that the column player's payoffs have not altered. What is the new equilibrium prediction? Is the data consistent or inconsistent with classical game theory?

10. Prove Proposition 12.23.
11. Prove Proposition 12.24.
12. (Goeree et al. 2005) Consider the asymmetric matching pennies game in Figure 4.E3 in which $x \geq -1$.

	L	R
U	$x, -1$	$-1, 1$
D	$-1, 1$	$1, -1$

Figure 4.E3 An asymmetric matching pennies game in Goeree et al. (2005).

Label the row player and the column players, respectively, as player 1 and player 2. Denote the probability with which the row player plays U by π_{1U} and the probability with which the column player plays L by π_{2L} .

- (a) Suppose that $x \geq 1$. Denote a regular QRE by (π_{1U}^*, π_{2L}^*) . Derive the conditions on (π_{1U}^*, π_{2L}^*) that are required for a regular QRE.
- (b) Suppose that $-1 < x < 1$. Denote a regular QRE by (π_{1U}^*, π_{2L}^*) . Derive the conditions on (π_{1U}^*, π_{2L}^*) that are required for a regular QRE.
- (c) In (π_{1U}, π_{2L}) space, draw the relevant regions that satisfy the restrictions of a QRE for the two cases: (i) $x = 9$, and (ii) $x = 1$. In the same diagram also indicate the points corresponding to a Nash equilibrium and to random play.
- (d) Show that in a regular QRE, the probability with which the row player plays U is increasing in x and the probability with which the column player plays D is decreasing in x .
13. Consider the properties of responsiveness and monotonicity required for a *regular quantal response function* (Definition 13.8). Consider a decision problem in which there is only one player, player 1, and there are only two actions, s_{11} and s_{12} . Using (13.5), (13.6), the probability that player 1 plays strategy s_{11} is $\pi_{11} = \Pr[\bar{u}_{11} + \varepsilon_{11} \geq \bar{u}_{12} - \varepsilon_{12}] = G(\bar{u}_{11} - \bar{u}_{12})$, where G is the distribution function of $\varepsilon_{12} - \varepsilon_{11}$ and $\bar{u}_{1j}, \varepsilon_{1j}$ are, respectively, the mean payoff and the errors arising from the play of strategy $s_{1j}, j = 1, 2$.
 - (a) Suppose that $G'(x) > 0$ for all x and $G(0) = 3/4$. What can we conclude about the responsiveness and monotonicity of the regular quantal response function?
 - (b) Suppose that $G(x) < 0.5$ for $x < 0$ and $G(x) > 0.5$ for $x > 0$ and there exists some interval $[x_1, x_2] \subset [0, 1]$ over which $G'(x) < 0$. What can we conclude about the responsiveness and monotonicity of the regular quantal response function?
14. Prove part (ii) of Proposition 13.26.
15. (Crawford et al., 2013) Consider the following pursuer–evader game known as the Far Pavilions game. The column player is pursuing the row player (the evader). Each player must simultaneously choose the action, head North (N) or South (S). The payoff matrix is shown in Figure 4.E4. The South route is the easier of the two. If both players head in the same direction, then the row player is caught and has a negative payoff. The payoff is more

	S	N
S	-1, 3	1, 0
N	0, 1	-2, 2

Figure 4.E4 The Far Pavilions Game.

Source: Crawford et al. (2013) with permission from the American Economic Association.

negative when both players head North because of the more difficult route. If both players head in different directions, then the row player gets a non-negative payoff. We wish to consider the wisdom of the following recommendation made by a “mentor” to the row player: “ride hard for the north, since they will be sure you will go southward, where the climate is kinder.”

- (a) Show that the game has a unique mixed strategy Nash equilibrium in which the probability of the mentor’s recommended outcome (N,S) being played out is $9/16$.
 - (b) Show that the recommendation of the mentor to the row player is equivalent to the optimal response of a $L3$ player when the anchoring type $L0$ player randomizes equally between N and S.
 - (c) Suppose that the frequency of types $L0$, $L1$, $L2$, and $L3$ players is respectively 0, 50, 30, and 20. Show that the probability of the outcome (N,S) being played is $2/5$ when the anchoring type $L0$ chooses uniformly randomly between N and S. How does this compare with a situation where (i) all players choose uniformly randomly between their strategies, and (ii) all players play the mixed strategy Nash equilibrium in part (b)?
16. (Crawford and Iriberri, 2007b) Consider the class of zero-sum games in which one player wins by matching actions with the other player, while the other player wins by mismatching actions. Figure 4.E5 shows a zero-sum game of hide-and-seek from Rubinstein (1999).

A hider hides some prize in one of the four boxes marked either A or B, which are ordered as shown in Figure 4.E5. The seeker does not know which box the prize is hidden in. Which box should the hider choose to hide the prize in? Which box should the seeker choose to open? If the seeker guesses the correct box, he gets 1 and the hider gets zero. Otherwise the hider gets 1 and the seeker gets zero.

- (a) Draw the normal form of the game and show that the game has a unique mixed strategy equilibrium. Show that seekers are at a disadvantage in this equilibrium because they can find the prize only 24% of the time.
- (b) Notice that the boxes marked A at both ends are salient/focal. The central box marked B is also salient because it is the only box not marked A. Hence, central A is the least salient location. From experimental results in Rubinstein et al. (1996), central A is a strongly modal choice. Hiders choose central A about 37% of the time and seekers choose to look in it about 46% of the time. Thus, in the experimental results, seekers are able to find the prize much more often than the equilibrium prediction in (a). What should a classical game theorist conclude?

**Figure 4.E5** A zero-sum game of hide-and-seek.

(c) Consider now the explanation of the results in Crawford and Iriberri (2007b) in terms of the level- k model. Each player (hider and seeker) has one of five possible types, $L0, L1, L2, L3, L4$. It does not make sense to assume that type $L0$ (whether hider or seeker) chooses each box with the same probability. The reason is that the labels on the box are an example of non-neutral framing. Suppose, therefore, that type $L0$ chooses the boxes with the following probability. The two end boxes are each chosen with a probability $\frac{p}{2}$, $p > \frac{1}{4}$. The central B box is chosen with a probability $q > \frac{1}{4}$, while the central A is chosen with a probability of $1 - p - q$. Thus, type $L0$ favours the salient locations: end As and the central B. Derive the optimal response of each type ($L0, L1, L2, L3, L4$) of each player (hider and seeker). Is this model able to explain the modal choice of central A for hiders and seekers when the percentages of the types $L0, L1, L2, L3, L4$ are respectively, 0, 19, 32, 24, 25?

(d) Can a quantal response equilibrium explain the modal choice?

17. (Crawford, 2007) Consider the BOS game shown in Figure 13.12 and its analysis in Section 13.4.2 with one round of communication. Make the following amendment to the preplay communication round. Instead of one round of communication, assume now that there are as many rounds of communication as players wish to have (abundant communication). Players in the preplay round simultaneously exchange messages. If they both decide to continue for another round of message exchange they can. However, if one or both players decide to terminate this round, then they choose their respective actions in the BOS game. Assume that players do not draw any inferences about the type of their opponent from the history of message exchange and there is a slight cost of continuing for an extra round of communication. Assume that $a < 2$. Show that with these assumptions, the outcome matrix is as in Proposition 13.31. What do you think happens when $a > 2$?
18. (Ellingsen and Östling, 2010) Derive the analogue of Proposition 13.32 when there is two-way communication in the preplay communication stage and both senders and receivers simultaneously exchange messages. Show that in this case, there is mis-coordination in half the cases when two type $L1$ players meet.
19. (Camerer et al., 2004a) Consider the following stag-hunt game that examines the effect of changes in group size. There are n players who simultaneously choose one of two strategies, L or H. The group choice is H if all players choose H. Otherwise, i.e., if at least one player chooses L, then the group choice is L. The payoffs of any player (row player) against the group average (column player) are shown in Figure 4.E6 and $0 < x < 1$. Thus, for instance, a player receives a payoff of x when he plays L and the group plays L. The Pareto optimal choice is for all players to play H.
 - (a) Consider a two-player stag-hunt game ($n = 2$). Show that if $x < 0.5$, type $L1$ players choose H, otherwise they choose L. Show that any type Lk player, $k > 1$, also behaves exactly like a type $L1$ player.
 - (b) Consider a three-player stag-hunt game ($n = 3$). Show now that type $L1$ players choose H if and only if $x \leq 0.25$ and L otherwise.

	L	H
L	x	0
H	0	1

Figure 4.E6 Payoffs in the n -player stag-hunt game. The column player represents the group average and the row player is any of the n players.

- (c) The experimental results of Van Huyck et al. (1990) show that an increase in group size reduces coordination. Is this consistent with the predictions of the level- k model for the n player stag-hunt game?
20. (Camerer et al., 2004a) Consider game 4 from Binmore et al. (2001) shown in Figure 4.E7. Each of the row and the column players has three strategies.
- (a) Show that the game has a mixed strategy equilibrium in which the row player plays B with probability 0.67 and the column players play R with probability 0.67.
- (b) Show that the predictions of a CH model (see Section 13.3.3) with, say, $\tau = 1.5$, are consistent with the experimental finding of the play of R by 83% of the players.
- (c) The experimental finding is that only 28% of the row players play their part of the Nash strategy profile, B. Why is B not very plausible? In a Nash equilibrium, B is chosen because the column players choose R and strategies and beliefs have to be mutually consistent. Show that the CH model predicts a 25% choice of each of T and M; a prediction that is closer to the experimental findings, relative to the Nash equilibrium.

	L	C	R	data	Nash	CH fit
T	0, 0	2, -2	-1, 1	0.33	0.17	0.25
M	2, -2	0, 0	-1, 1	0.29	0.17	0.25
B	-1, 1	-1, 1	0, 0	0.28	0.67	0.50
data	0.13	0.04	0.83			
Nash	0.17	0.17	0.67			
CH fit	0.08	0.08	0.84			

Figure 4.E7 Game 4 from Binmore et al. (2001).

Source: "Does minimax work? An experimental study." Binmore, K. G., Swierzbinski, J., and Proulx, C. (2001). *Economic Journal* 111(473): 445–64. © 2001 John Wiley and Sons.

21. (Geanakoplos et al. 1989) Consider the confidence game in Example 13.13. Derive the psychological Nash equilibria in the game.
22. Consider the two-player game Hawk–Dove shown in Figure 4.E8 that was used in Rabin (1993) to illustrate the reciprocity equilibria in a single stage, simultaneous move game. The game is symmetric and each player has two strategies, Hawk (h) and Dove (d). Player 1 is the row player and player 2 is the column player, and $X > 0$ is a scaling factor. The game has two Nash equilibria, (h, d) and (d, h) . Modify the multistage model in Dufwenberg and Kirchsteiger (2004) (see Section 13.5.3) to a single stage in order to compute the psychological equilibria for this game.
23. Prove Proposition 13.36(iii), (iv).

	Hawk	Dove
Hawk	$-2X, -2X$	$2X, 0$
Dove	$0, 2X$	X, X

Figure 4.E8 A hawk–dove game.




Figure 4.E9 A Trust game in extensive form.

Source: Reprinted from *Journal of Economic Theory* 144(1): 1–35. Battigalli, P. and Dufwenberg, M. "Dynamic psychological games." Copyright © 2009, with permission from Elsevier.

24. (Battigalli and Dufwenberg, 2009) Consider the extensive form game shown in Figure 4.E9 (the payoff of the first player is listed first). In the unique backward induction outcome, player 1 plays Don't Trust and ends the game at the first node with the inefficient payoffs (1, 1).
- Now introduce the possibility of guilt. Let player 2 play Share with probability p . Let the first order beliefs of player 1 about p be denoted by b_1^1 and let the beliefs of player 2 about b_1^1 be denoted by b_2^2 . Suppose that guilt takes a particularly simple form. When player 2 plays Grab he faces a reduction in his material payoff by $5b_2^2$. Use a *psychological forward induction* argument to show how the efficient outcome with payoffs (2, 2) can arise [Hint: If player 1 plays Trust, what does this tell player 2 about the lower bound on b_1^1 ?].
25. Consider the results of Charness and Dufwenberg (2006) in Section 13.5.5. Can these results be explained by the following three classes of models?
- Models of inequity aversion, considered in Part 2 of the book, such as those by Fehr and Schmidt (1999), Bolton and Ockenfels (2000), and Charness and Rabin (2002).
 - Models of intentions-based reciprocity, considered in Chapter 13, such as those by Rabin (1993) and Dufwenberg and Kirchsteiger (2004).
 - Models in which there is a fixed cost of lying, such as the model by Gneezy (2005) considered in Part 2 of the book.
26. Prove Proposition 13.42.
27. (Jehiel, 2005) Consider the analogy-based expectations equilibrium (ABEE) of Jehiel (2005). Suppose that we have the two normal form games G1 and G2, shown in Figure 4.E10

	L	M	R	R'		L	M	R	R'
U	5, 2	0, 2	2, 4	0, 0	U	3, 0	4, 2	2, 0	1, 1
D	4, 3	3, 0	1, 0	2, 0	D	0, 2	5, 2	0, 0	2, 4
Game G1					Game G2				

Figure 4.E10 Two normal form games from Jehiel (2005).

Source: Reprinted from *Journal of Economic Theory* 123: 81–104. Jehiel, P. "Analogy-based expectation equilibrium." Copyright © 2005, with permission from Elsevier.

Here we have the following data for the row player (r) and the column player (c): $S_r = \{U, D\}$, $S_c = \{L, M, R, R'\}$, $G = \{G1, G2\}$ and $p(G1) = p(G2) = 0.5$. The analogy partition used by the column player is always fine, i.e., $\Pi_c = \{\{G1\}, \{G2\}\}$. However, the row player has a coarse partition, so $\Pi_r = \{\{G1, G2\}\}$.

Each of the games has a unique Nash equilibrium: (U, R) in game G1 and (D, R') in game G2. Show that in the ABEE of the game, the row player plays D in game G1 and U in game G2 and the respective strategies of the column player are L in game G1 and M in game G2. Does the row player benefit by having a finer partition $\Pi_r = \{\{G1\}, \{G2\}\}$?

28. Prove Proposition 13.44.
29. (al-Nowaihi and Dhami, 2015) This question relates to evidential games in Section 13.7.4. Consider the Nash demand game (Nash, 1950a, 1953). Two players share a cake of size one. Player 1 demands $x \in [0, 1]$ and player 2 simultaneously demands $y \in [0, 1]$. If the demands are feasible, i.e., if $x + y \leq 1$, then each player receives what she demanded. However, if the demands are not feasible, i.e., if $x + y > 1$, then each player gets 0. Any pair, (x, y) , such that $x + y = 1$ is a Nash non-cooperative equilibrium. Thus, the Nash non-cooperative equilibrium concept does not pin down a unique solution in the Nash demand game.
 - (a) Show that if we allow arbitrary social projection functions, then any outcome, (a, b) , such that $a + b = 1$, $a > 0$ and $b > 0$, is an outcome of a consistent evidential equilibrium for the Nash demand game for suitably chosen social projection functions.
 - (b) Show that $(\frac{1}{2}, \frac{1}{2})$ is the unique consistent evidential equilibrium for the Nash demand game under the identity social projection functions: $P_{12}(y|x) = 1 \Leftrightarrow y = x$ and $P_{21}(x|y) = 1 \Leftrightarrow x = y$.
30. Consider two risk-neutral bidders, i, j , participating in a first price sealed-bid auction for one indivisible unit of an object. Both bidders use expected utility to evaluate risky payoffs. Bidders have valuations v_i, v_j that are privately known. Each bidder knows that the valuation of the other is drawn from a uniform distribution $[0, 1]$. The distributions for v_i, v_j are independent and identical. If a bidder wins the object and pays a price p , then the utility of the bidder is $v_i - p$. Suppose that bidders simultaneously submit bids b_i, b_j and the one who makes a higher bid wins. In the event of a tie, both bidders get the object with a probability of $1/2$. The payoff of the loser is zero. We are interested in the Bayesian Nash equilibrium (BNE) of this static game of incomplete information. Consider equilibrium bidding strategies that are linear in one's valuation,

$$b_i(v_i) = \alpha_i + \beta_i v_i; \quad b_j(v_j) = \alpha_j + \beta_j v_j. \quad (4.E1)$$

Notice that in the BNE, we are required to determine the optimal bid for each type of the player; (4.E1) satisfies this requirement with each type of a player associated with a distinct valuation.

- (a) Show that in the unique BNE of the game,

$$b_i = \frac{v_i}{2}; \quad b_j = \frac{v_j}{2}.$$

Notice that unlike the English auctions and the second price auctions, one's optimal bids are below one's valuation.

- (b) Show that when the valuations are uniformly distributed, an equilibrium in linear bidding strategies as in (4.E1) exists and is unique.
- (c) Is the outcome of the first price auction efficient? ■
31. (Crawford et al., 2009)
 - (a) Consider the increasing density case with level- k players in the single price simultaneous bid auction with two risk-neutral players, as in Section 13.8.4. Thus, while the valuation

of each bidder is private information, bidders know the distribution of each other's possible valuations, which is given by (13.105). Show that for truthful level- k players, the bids of types $L1$ and $L2$ are identical and given by

$$B_2^* = \max \left\{ \rho, \left(\frac{\gamma}{1+\gamma} \right) v \right\} \leq B^*(v)$$

where $B^*(v)$ is the optimal rational bid as defined in (13.107) and B_2^* is the optimal bid of a type $L2$ player.

(b) Show that for the decreasing density case, given in (13.106), the optimal bids of types $L1$ and $L2$ are consistent with the winner's curse.

REFERENCES FOR PART 4

- Abreu, D. and Matsushima, H. (1992). Virtual implementation in iteratively undominated strategies: complete information. *Econometrica* 60(5): 993–1008.
- Acevedo, M. and Krueger, J. I. (2005). Evidential reasoning in the prisoners' dilemma game. *American Journal of Psychology* 118(3): 431–57.
- Agranov, M., Caplin, A., and Tergiman, C. (2012). Naive play and the process of choice in guessing games. Unpublished.
- Akerlof, G. A. (1970). The market for “lemons”: quality uncertainty and the market mechanism. *Quarterly Journal of Economics* 84(3): 488–500.
- Akerlof, G. A. and Yellen, J. L. (1985). A near-rational model of the business cycle, with wage and price inertia. *Quarterly Journal of Economics* 100: 823–38.
- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'Ecole Americaine. *Econometrica* 21(4): 503–46.
- al-Nowaihi, A. and Dhami, S. (2013). A theory of reference time. University of Leicester, Discussion Papers in Economics 13/26.
- al-Nowaihi, A. and Dhami, S. (2015). Evidential equilibria: Heuristics and biases in static games of complete information. Forthcoming in *Games*.
- Amaldoss, W. and Jain, S. (2002). David vs. Goliath: an analysis of asymmetric mixed-strategy games and experimental evidence. *Management Science* 48(8): 972–91.
- Andreoni, J. (1988). Why free ride? Strategies and learning in public goods experiments. *Journal of Public Economics* 37(3): 291–304.
- Aoyagi, M. and Fréchette, G. (2009). Collusion as public monitoring becomes noisy: experimental evidence. *Journal of Economic Theory* 144(3): 1135–65.
- Arad, A. and Rubinstein, A. (2012). The 11-20 money request game: a level-k reasoning study. *American Economic Review* 102(7): 3561–73.
- Attanasi, G., Battigalli, P., and Manzoni, E. (2015). Incomplete information models of guilt aversion in the trust game. Forthcoming in *Management Science*.
- Attanasi, G. and Nagel, R. (2008). A survey on psychological games: theoretical findings and experimental evidence. In A. Innocenti and P. Sbriglia (eds.), *Games, Rationality and Behaviour: Essays on Behavioural Game Theory and Experiments*. Houndmills: Palgrave Macmillan.
- Aumann, R. J. (1974). Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1(1): 67–96.
- Aumann, R. J. (1985). On the non-transferable utility value: a comment on the Roth-Shafer examples. *Econometrica* 53(3): 667–78.
- Aumann, R. J. (1987). Correlated equilibrium as an expression of Bayesian rationality. *Econometrica* 55(1): 1–18.
- Aumann, R. J. (1992). Irrationality in game theory. In P. Dasgupta, D. Gale, O. Har and E. Maskin (eds.), *Economic Analysis of Markets and Games, Essays in Honor of Frank Hahn*. Cambridge, MA: MIT Press, pp. 214–27.
- Aumann, R. J. (1995). Backward induction and common knowledge of rationality. *Games and Economic Behavior* 8(1): 6–19.
- Aumann, R. J. and Brandenburger, A. (1995). Epistemic conditions for Nash equilibrium. *Econometrica* 63(5): 1161–80.
- Avery, C. and Kagel, J. H. (1997). Second-price auctions with asymmetric payoffs: an experimental investigation. *Journal of Economics and Management Strategy* 6(3): 573–603.
- Babcock, L. and Loewenstein, G. (1997). Explaining bargaining impasse: the role of self-serving biases. *Journal of Economic Perspectives* 11(1): 109–26.
- Babcock, L., Loewenstein, G., Issacharoff, S., and Camerer, C. F. (1995). Biased judgments of fairness in bargaining. *American Economic Review* 85(5): 1337–43.
- Babcock, L., Wang, X., and Loewenstein, G. (1996). Choosing the wrong pond: social comparisons that reflect a self-serving bias. *Quarterly Journal of Economics* 111(1): 1–19.
- Bacharach, M. (1999). Interactive team reasoning: a contribution to the theory of cooperation. *Research in Economics* 53(2): 117–47.
- Bacharach, M. (2006). *Beyond Individual Choice: Teams and Frames in Game Theory*. N. Gold and R. Sugden (eds). Princeton, NJ: Princeton University Press.

- Bacharach, M., Guerra, G., and Zizzo, D. J. (2007). The self-fulfilling property of trust: an experimental study. *Theory and Decision* 63(4): 349–88.
- Ball, S. B., Bazerman, M. H., and Carroll, J. S. (1991). An evaluation of learning in the bilateral winner's curse. *Organizational Behavior and Human Decision Processes* 48(1): 1–22.
- Bandyopadhyay, S. and Chatterjee, K. (2012). Models of coalition formation in multilateral negotiations. In: G. E. Bolton and R. T. A. Croson (eds.), *The Oxford Handbook of Economic Conflict Resolution*. New York: Oxford University Press, pp. 76–88.
- Banerjee, A. V. (1992). A simple model of herd behavior. *Quarterly Journal of Economics* 107(3): 797–817.
- Banks, J., Camerer, C. F., and Porter, D. (1994). An experimental analysis of Nash refinements in signaling games. *Games and Economic Behavior* 6(1): 1–31.
- Banks, J. S. and Sobel, J. (1987). Equilibrium selection in signaling games. *Econometrica* 55(3): 647–61.
- Bardsley, N., Mehta, J., Starmer, C., and Sugden, R. (2010). Explaining focal points: cognitive hierarchy theory versus team reasoning. *Economic Journal* 120(543): 40–79.
- Bar-Eli, M., Azar, O. H., Ritov, I., Keidar-Levin, Y., and Schein, G. (2007). Action bias among elite soccer goalkeepers: the case of penalty kicks. *Journal of Economic Psychology* 28(5): 606–21.
- Basu, K. (1994). The traveler's dilemma: paradoxes of rationality in game theory. *American Economic Review* 84(2): 391–5.
- Battalio, R., Samuelson, L., and Van Huyck, J. (2001). Optimization incentives and coordination failure in laboratory stag hunt games. *Econometrica* 69(3): 749–64.
- Battigalli, P. and Dufwenberg, M. (2007). Guilt in games. *American Economic Review* 97(2): 170–6.
- Battigalli, P. and Dufwenberg, M. (2009). Dynamic psychological games. *Journal of Economic Theory* 144(1): 1–35.
- Battigalli, P. and Siniscalchi, M. (1999). Hierarchies of conditional beliefs and interactive epistemology in dynamic games. *Journal of Economic Theory* 88: 188–230.
- Bazerman, M. H. and Samuelson, W. F. (1983). I won the auction but don't want the prize. *Journal of Conflict Resolution* 27(4): 618–34.
- Beard, T. R. and Beil, R. O., Jr. (1994). Do people rely on the self-interested maximization of others? An experimental test. *Management Science* 40(2): 252–62.
- Bereby-Meyer, Y. and Roth, A. E. (2006). The speed of learning in noisy games: partial reinforcement and the sustainability of cooperation. *American Economic Review* 96(4): 1029–42.
- Bernheim, B. D. (1984). Rationalizable strategic behavior. *Econometrica* 52(4): 1007–28.
- Bicchieri, C. and Green, M. (1997). Symmetry arguments for cooperation in the prisoner's dilemma. In G. Holmstrom-Hintikka and R. Tuomela (eds.), *Contemporary Action Theory: The Philosophy and Logic of Social Action*. Alphen aan den Rijn: Kluwer.
- Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as information cascades. *Journal of Political Economy* 100(5): 992–1026.
- Binmore, K. (1991). *Fun and Games: A Text on Game Theory*. Lexington, MA: D. C. Heath.
- Binmore, K. G., Shaked, A., and Sutton, J. (1985). Testing noncooperative bargaining theory: a preliminary study. *American Economic Review* 75(5): 1178–80.
- Binmore, K. G., Shaked, A., and Sutton, J. (1989). An outside option experiment. *Quarterly Journal of Economics* 104(4): 753–70.
- Binmore, K. G., Swierzbinski, J., and Proulx, C. (2001). Does minimax work? An experimental study. *Economic Journal* 111(473): 445–64.
- Blanchard, O. J. and Fischer, S. (1989). *Lectures on Macroeconomics*. Cambridge, MA: MIT Press.
- Blinder, A. S., Canetti, E. R. D., Lebow, D. F., and Rudd, J. B. (1998). *Asking About Prices. A New Approach to Understanding Price Stickiness*. New York: Russell Sage Foundation.
- Bloomfield, R. (1994). Learning a mixed strategy equilibrium in the laboratory. *Journal of Economic Behavior and Organization* 25(3): 411–36.
- Blume, A. and Ortmann, A. (2007). The effects of costless pre-play communication: experimental evidence from games with Pareto-ranked equilibria. *Journal of Economic Theory* 132(1): 274–90.
- Boğaçhan, Ç. and Kariv, S. (2004). Distinguishing informational cascades from herd behavior in the laboratory. *American Economic Review* 94(3): 484–98.

- Bolton, G. E. (1991). A comparative model of bargaining: theory and evidence. *American Economic Review* 81(5): 1096–36.
- Bolton, G. E. and Croson, R. T. A. (2012). *The Oxford Handbook of Economic Conflict Resolution*. New York: Oxford University Press.
- Bolton, G. E., Katok, E., and Zwick, R. (1998). Dictator same giving: rules of fairness versus acts of kindness. *International Journal of Game Theory* 27: 269–99.
- Bolton, G. E. and Ockenfels, A. (2000). ERC: a theory of equity, reciprocity and competition. *American Economic Review* 90(1): 166–93.
- Bolton, G. E. and Zwick, R. (1995). Anonymity versus punishment in ultimatum bargaining. *Games and Economic Behavior* 10(1): 95–121.
- Bonanno, G. (2013). Counterfactuals and the prisoner's dilemma. University of California, Department of Economics. Working Paper 137.
- Bosch-Domènech, A., Montalvo, J. G., Nagel, R., and Satorra, A. (2002). One, two, (three), infinity, . . . : Newspaper and lab beauty-contest experiments. *American Economic Review* 92(5): 1687–701.
- Bosch-Domènech, A. and Vriend, N. J. (2013). On the role of non-equilibrium focal points as coordination devices. *Journal of Economic Behavior and Organization* 94: 52–67.
- Brandts, J. and Cooper, D. J. (2006). A change would do you good. An experimental study on how to overcome coordination failure in organizations. *American Economic Review* 96(3): 669–93.
- Brandts, J. and Cooper, D. J. (2007). It's what you say, not what you pay: an experimental study of manager–employee relationships in overcoming coordination failure. *Journal of the European Economic Association* 5(6): 1223–68.
- Brandts, J. and Holt, C. A. (1992). An experimental test of equilibrium dominance in signaling games. *American Economic Review* 82(5): 1350–65.
- Brandts, J. and Holt, C. A. (1993). Adjustment patterns and equilibrium selection in experimental signaling games. *International Journal of Game Theory* 22(3): 279–302.
- Brandts, J. and McLeod, W. B. (1995). Equilibrium selection in experimental games with recommended play. *Games and Economic Behavior* 11(1): 36–63.
- Breitmöser, Y. (2012). Strategic reasoning in p-beauty contests. *Games and Economic Behavior* 75(2): 555–69.
- Brewer, P. J., Huang, M., Nelson, B., and Plott, C. R. (2002). On the behavioral foundations of the law of supply and demand: human convergence and robot randomness. *Experimental Economics* 5(3): 179–208.
- Brocas, I., Carrillo, J. D., Wang, S. W., and Camerer, C. F. (2014). Imperfect choice or imperfect attention? Understanding strategic thinking in private information games. Forthcoming in *Review of Economic Studies*.
- Brock, W. A. and Hommes, C. (1997). A rational route to randomness. *Econometrica* 65(5): 1059–95.
- Brown, J. N. and Rosenthal, R. W. (1990). Testing the minimax hypothesis: a re-examination of O'Neill's game experiment. *Econometrica* 58(5): 1065–81.
- Bryant, J. (1983). A simple rational expectations Keynes-type model. *Quarterly Journal of Economics* 98(3): 525–8.
- Budescu, D. V. and Rapoport, A. (1994). Subjective randomization in one- and two-person games. *Journal of Behavioral Decision Making* 7(4): 261–78.
- Bulow, J. I., Geanakoplos, J. D., and Klemperer, P. (1985). Multimarket oligopoly: strategic substitutes and strategic complements. *Journal of Political Economy* 93(3): 488–511.
- Burchardi, K. B. and Penczynski, S. P. (2014). Out of your mind: eliciting individual reasoning in one shot games. *Games and Economic Behavior* 84: 39–57.
- Cachon, G. P. and Camerer, C. F. (1996). Loss-avoidance and forward induction in experimental coordination games. *Quarterly Journal of Economics* 111(1): 165–94.
- Cadsby, C. B., Frank, M., and Maksimovic, V. (1990). Pooling, separating, and semiseparating equilibria in financial markets: some experimental evidence. *Review of Financial Studies* 3(3): 315–42.
- Cadsby, C. B., Frank, M., and Maksimovic, V. (1998). Equilibrium dominance in experimental financial markets. *Review of Financial Studies* 11(1): 189–232.
- Camerer, C. F. (2003a). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton University Press.

- Camerer, C. F. (2003b). Behavioural studies of strategic thinking in games. *Trends in Cognitive Sciences* 7(5): 225–31.
- Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2004a). A cognitive hierarchy model of games. *Quarterly Journal of Economics* 119(3): 861–98.
- Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2004b). Behavioral game theory: thinking, learning, and teaching. In S. Huck (ed.), *Advances in Understanding Strategic Behavior, Essays in Honor of Werner Güth*. London: Palgrave Macmillan.
- Camerer, C. F., Johnson, E., Rymon, T., and Sen, S. (1994). Cognition and framing in sequential bargaining for gains and losses. In K. Binmore, A. Kirman, and P. Tani (eds.), *Frontiers of Game Theory*. Cambridge, MA: MIT Press, pp. 27–47.
- Camerer, C. F. and Knez, M. (2000). Increasing cooperation in social dilemmas by establishing a precedent of efficiency in coordination games. *Organizational Behavior and Human Decision Processes* 82(2): 194–216.
- Camerer, C. F. and Weigelt, K. (1988). Experimental tests of a sequential equilibrium reputation model. *Econometrica* 56(1): 1–36.
- Camerer, C. F., Nunnari, S., and Palfrey, T. R. (2012). Quantal response and nonequilibrium beliefs explain overbidding in maximum-value auctions. Unpublished.
- Capen, E. C., Clapp, R. V., and Campbell, W. M. (1971). Competitive bidding in high-risk situations. *Journal of Petroleum Technology* 23(6): 641–53.
- Caplin, A., Dean, M., and Martin, D. (2011). Search and satisficing. *American Economic Review* 101(7): 2899–922.
- Capra, C. M., Goeree, J. K., Gomez, R., and Holt, C. A. (1999). Anomalous behavior in a travelers dilemma? *American Economic Review* 89(3): 678–90.
- Carrillo, J. D. and Palfrey, T. (2009). The compromise game: two-sided adverse selection in the laboratory. *American Economic Journal: Microeconomics* 1(1): 151–81.
- Carrillo, J. D. and Palfrey, T. (2011). No trade. *Games and Economic Behavior* 71(1): 66–87.
- Cason, T. N., Savikhin, A. C., and Sheremeta, R. M. (2012). Behavioral spillovers in coordination games. *European Economic Review* 56(2): 233–45.
- Cason, T. N. and Sharma, T. (2007). Recommended play and correlated equilibria: an experimental study. *Economic Theory* 33(1): 11–27.
- Chamberlin, E. H. (1948). An experimental imperfect market. *Journal of Political Economy* 56(2): 95–108.
- Charness, G. (2000). Self-serving cheap talk: a test of Aumann's conjecture. *Games and Economic Behavior* 33(2): 177–94.
- Charness, G. and Dufwenberg, M. (2006). Promises and partnership. *Econometrica* 74(6): 1579–601.
- Charness, G. and Levin, D. (2009). The origin of the winner's curse: a laboratory study. *American Economic Journal: Microeconomics* 1(1): 207–36.
- Charness, G. and Rabin, M. (2002). Understanding social preferences with simple tests. *Quarterly Journal of Economics* 117(3): 817–69.
- Chatterjee, K. and Samuelson, W. (1983). Bargaining under incomplete information. *Operations Research* 31(5): 835–51.
- Chaudhuri, A., Schotter, A., and Sopher, B. (2009). Talking ourselves to efficiency: coordination in inter-generational minimum effort games with private, almost common and common knowledge of advice. *Economic Journal* 119(534): 91–122.
- Chiappori, P. A., Levitt, S. D., and Groseclose, T. (2002). Testing mixed strategy equilibrium when players are heterogeneous: the case of penalty kicks. *American Economic Review* 92(4): 1138–51.
- Cho, I.-K. and Kreps, D. M. (1987). Signaling games and stable equilibria. *Quarterly Journal of Economics* 102(2): 179–221.
- Chong, J.-K., Camerer, C. F., and Ho, T.-H. (2005). Cognitive hierarchy: a limited thinking theory in games. In: R. Zwick and A. Rapoport (eds.), *Experimental Business Research: Marketing, Accounting and Cognitive Perspectives, Volume 3*. Dordrecht and New York: Springer, pp. 203–28.
- Cohen, R. B., Polk, C., and Vuolteenaho, T. (2005). Money illusion in the stock market: the Modigliani-Cohn hypothesis. *Quarterly Journal of Economics* 120(2): 639–68.
- Collins, R. and Sherstyuk, K. (2000). Spatial competition with three firms: an experimental study. *Economic Inquiry* 38(1): 73–94.

- Colman, A. M. and Bacharach, M. (1997). Payoff dominance and the Stackelberg heuristic. *Theory and Decision* 43(1): 1–19.
- Colman, A. M., Pulford, B. D., and Lawrence, C. L. (2014). Explaining strategic coordination: cognitive hierarchy theory, strong Stackelberg reasoning, and team reasoning. *Decision* 1(1): 35–58.
- Colman, A. M., Pulford, B. D., and Rose, J. R. (2008). Team reasoning and collective rationality: piercing the veil of obviousness. *Acta Psychologica* 128(2): 409–12.
- Colman, A. M. and Stirk, J. A. (1998). Stackelberg reasoning in mixed-motive games: an experimental investigation. *Journal of Economic Psychology* 19(2): 279–93.
- Cooper, D. J. (2006). Are experienced managers experts at overcoming coordination failure? The field experiments, Artefactual Field Experiments, Working Paper 00037.
- Cooper, D. J., Garvin, S., and Kagel, J. H. (1997a). Adaptive learning vs. equilibrium refinements in an entry limit pricing game. *Economic Journal* 107(442): 553–75.
- Cooper, D. J., Garvin, S., and Kagel, J. H. (1997b). Signaling and adaptive learning in an entry limit pricing game. *RAND Journal of Economics* 28(4): 662–83.
- Cooper D. J. and Van Huyck, J. (2003). Evidence on the equivalence of the strategic and extensive form representation of games. *Journal of Economic Theory* 110(2): 290–308.
- Cooper, R. W., DeJong, D. V., Forsythe, R., and Ross, T. W. (1989). Communication in the battle of the sexes game: some experimental results. *RAND Journal of Economics* 20(4): 568–87.
- Cooper, R. W., DeJong, D. V., Forsythe, R., and Ross, T. W. (1990). Selection criteria in coordination games: some experimental results. *American Economic Review* 80(1): 218–33.
- Cooper, R. W., DeJong, D. V., Forsythe, R., and Ross, T. W. (1992). Communication in coordination games. *Quarterly Journal of Economics* 107(2): 739–71.
- Cooper, R. W., DeJong, D. V., Forsythe, R. and Ross, T. W. (1993). Forward induction in the battle-of-the-sexes games. *American Economic Review* 83(5): 1303–16.
- Cooper, R. W., DeJong, D. V., Forsythe, R. and Ross, T. W. (1994). Alternative institutions for resolving coordination problems: experimental evidence on forward induction and preplay communication. In J. W. Friedman (ed.), *Problems of Coordination in Economic Activity*. Norwell, MA: Kluwer Academic Publishers.
- Cooper, R. W. and John, A. (1988). Coordinating coordination failures in Keynesian models. *Quarterly Journal of Economics* 103(3): 441–63.
- Coppinger, V. M., Smith, V. L., and Titus, J. A. (1980). Incentives and behavior in English, Dutch and sealed-bid auctions. *Economic Inquiry* 18(1): 1–22.
- Costa-Gomes, M. A. and Crawford, V. P. (2006). Cognition and behavior in two-person guessing games: an experimental study. *American Economic Review* 96(5): 1737–68.
- Costa-Gomes, M., Crawford, V. P., and Broseta, B. (2001). Cognition and behavior in normal-form games: an experimental study. *Econometrica* 69(5): 1193–235.
- Costa-Gomes, M. A. and Shimoji, M. (2011). A comment on “Can relaxation of beliefs rationalize the winner’s curse? An experimental study.” *Econometrica* 83: 375–83.
- Costa-Gomes, M. A. and Weizsäcker, G. (2008). Stated beliefs and play in normal-form games. *Review of Economic Studies* 75(3): 729–62.
- Crawford, V. P. (1995). Adaptive dynamics in coordination games. *Econometrica* 63(1): 103–43.
- Crawford, V. P. (2003). Lying for strategic advantage: rational and boundedly rational misrepresentation of intentions. *American Economic Review* 93(1): 133–49.
- Crawford, V. P. (2007). Let’s talk it over: coordination via preplay communication with level-k thinking. Mimeo, UCLA Department of Economics.
- Crawford, V. P. and Broseta, B. (1998). What price coordination? The efficiency enhancing effect of auctioning the right to play. *American Economic Review* 88(1): 198–225.
- Crawford, V. P., Costa-Gomes, M. A., and Iriberri, N. (2013). Structural models of non-equilibrium strategic thinking: theory, evidence, and applications. *Journal of Economic Literature* 51(1): 5–62.
- Crawford, V. P., Gneezy, U., and Rottenstreich, Y. (2008). The power of focal points is limited: even minute payoff asymmetry may yield large coordination failures. *American Economic Review* 98(4): 1443–58.
- Crawford, V. P. and Iriberri, N. (2007a). Fatal attraction: salience, naïveté, and sophistication

- in experimental hide-and-seek games. *American Economic Review* 97(5): 1731–50.
- Crawford, V. P. and Iriberry, N. (2007b). Level-k auctions: can a nonequilibrium model of strategic thinking explain the winner's curse and overbidding in private-value auctions? *Econometrica* 75(6): 1721–70.
- Crawford, V. P., Kugler, T., Neeman, Z., and Pauzner, A. (2009). Behaviorally optimal auction design: an example and some observations. *Journal of the European Economic Association* 7(2–3): 377–87.
- Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica* 50(6): 1431–51.
- Croson, R. T. A. (2000). Feedback in voluntary contribution mechanisms: an experiment in team production. *Research in Experimental Economics* 8: 85–97.
- Cubitt, R. P. and Sugden, R. (2011). The reasoning-based expected utility procedure. *Games and Economic Behavior* 71(2): 328–38.
- Dal Bó, P. (2005). Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. *American Economic Review* 95(5): 1591–604.
- Dal Bó, P. and Fréchette, G. R. (2011). The evolution of cooperation in infinitely repeated games: experimental evidence. *American Economic Review* 101(1): 411–29.
- Daniel, T. E., Scale, D. A., and Rapoport, A. (1998). Strategic play and adaptive learning in sealed bid bargaining mechanism. *Journal of Mathematical Psychology* 42(2–3): 133–66.
- Danz, D. N., Fehr, D., and Kübler, D. (2012). Information and beliefs in a repeated normal-form game. *Experimental Economics* 15(4): 622–40.
- Darai, D. and Grätz, S. (2010). Golden balls: a prisoner's dilemma experiment. University of Zurich. Working Paper 1006.
- Davis, D. D. and Holt, C. A. (1993). *Experimental Economics*. Princeton, NJ: Princeton University Press.
- De Bruyn, A. and Bolton, G. E. (2008). Estimating the influence of fairness on bargaining behavior. *Management Science* 54(1): 1774–91.
- de Sousa, J., Hollard, G., and Terracol, A. (2013). Non-strategic players are the rule rather than the exception. Unpublished.
- Dechenaux, E., Kovenock, D., and Lugovskyy, V. (2006). Caps on bidding in all-pay auctions: comments on the experiments of A. Rapoport and W. Amaldoss. *Journal of Economic Behavior and Organization* 61(2): 276–83.
- Dekel, E. and Siniscalchi, M. (2015). Epistemic game theory. In H. P. Young, and S. Zamir (ed.) *Handbook of Game Theory with Economic Applications, Volume 4*. North Holland. pp. 619–702.
- Dhami, S. and al-Nowaihi, A. (2013). Evidential equilibria: heuristics and biases in static games. University of Leicester, Department of Economics. Working Paper 13/25.
- Drouvelis, M., Müller, W., and Possajennikov, A. (2012). Signaling without common prior: results on experimental equilibrium selection. *Games and Economic Behavior* 74(1): 102–19.
- Duffy, J. (2006). Agent-based models and human subject experiments. In: L. Tesfatsion and K. L. Judd (eds.), *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*. Amsterdam: North-Holland, pp. 949–1011.
- Duffy, J. and Feltovich, N. (2002). Do actions speak louder than words? An experimental comparison of observation and cheap talk. *Games and Economic Behavior* 39(1): 1–27.
- Duffy, J. and Feltovich, N. (2006). Words, deeds, and lies: strategic behaviour in games with multiple signals. *Review of Economic Studies* 73(3): 669–88.
- Duffy, J. and Feltovich, N. (2010). Correlated equilibria, good and bad: an experimental study. *International Economic Review* 51(3): 701–21.
- Duffy J. and Ochs, J. (2009). Cooperative behavior and the frequency of social interaction. *Games and Economic Behavior* 66(2): 785–812.
- Dufwenberg, M. (2008). Psychological games. In S. N. Durlauf and L. E. Blume (eds.), *The New Palgrave Dictionary of Economics*. 2nd edition. London: Palgrave Macmillan.
- Dufwenberg, M., Gächter, S., and Hennig-Schmidt, H. (2011). The framing of games and the psychology of play. *Games and Economic Behavior* 73(2): 459–78.
- Dufwenberg, M. and Gneezy, U. (2000). Measuring beliefs in an experimental lost wallet game. *Games and Economic Behavior* 30(2): 163–82.
- Dufwenberg, M. and Kirchsteiger, G. (2004). A theory of sequential reciprocity. *Games and Economic Behavior* 47(2): 268–98.
- Dufwenberg, M., Sundaram, R., and Butler, D. (2010). Epiphany in the Game of 21. *Journal*

- of *Economic Behavior and Organization* 75(2): 132–43.
- Dyer, D., Kagel, J. H., and Levin, D. (1989). A comparison of naive and experienced bidders in common value offer auctions: a laboratory analysis. *Economic Journal* 99(394): 108–15.
- Eckel, C. C. (2007). Economic games for social scientists. In M. Webster and J. Sell (eds.), *Laboratory Experiments in Social Sciences*. Oxford: Elsevier.
- Ellingsen, T., Johannesson, M., Tjøtta, S., and Torsvik, G. (2010). Testing guilt aversion. *Games and Economic Behavior* 68(1): 95–107.
- Ellingsen, T. and Östling, R. (2010). When does communication improve coordination? *American Economic Review* (100)4: 1695–1724.
- Elster, J. (1998). Emotions and economic theory. *Journal of Economic Literature* 36(1): 47–74.
- Evans, G. W. and Honkapojha, S. (2001). *Learning and Expectations in Macroeconomics*. Princeton, NJ: Princeton University Press.
- Evans, G. W. and Ramey, G. (1992). Expectation calculation and macroeconomics dynamics. *American Economic Review* 82(1): 207–24.
- Evans, G. W. and Ramey, G. (1998). Calculation, adaptation and rational expectations. *Macroeconomic Dynamics* 2(2): 156–82.
- Eyster, E. and Rabin, M. (2005). Cursed equilibrium. *Econometrica* 73(5): 1623–72.
- Falk, A., Fehr, E., and Fischbacher, U. (2008). Testing theories of fairness and reciprocity—intentions matter. *Games and Economic Behavior* 62(1): 287–303.
- Falk, A. and Fischbacher, U. (2006). A theory of reciprocity. *Games and Economic Behavior* 54(2): 293–315.
- Farrell, J. (1987). Cheap talk, coordination, and entry. *RAND Journal of Economics* 18(1): 34–9.
- Farrell, J. (1988). Communication, coordination and Nash equilibrium. *Economic Letters* 27(3): 209–14.
- Farrell, J. and Rabin, M. (1996). Cheap talk. *Journal of Economic Perspectives* 10(3): 103–18.
- Fehr, E., Kirchsteiger, G., and Riedl, A. (1998). Gift exchange and reciprocity in competitive experimental markets. *European Economic Review* 42(1): 1–34.
- Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics* 114(3): 817–68.
- Fehr, E. and Tyran, J.-R. (2001). Does money illusion matter? *American Economic Review* 91(5): 1239–62.
- Fehr, E. and Tyran, J.-R. (2002). Limited rationality and strategic interaction: the impact of the strategic environment on nominal inertia. University of St. Gallen, Department of Economics. Working Paper 2002-25.
- Fehr, E. and Tyran, J.-R. (2005). Individual irrationality and aggregate outcomes. *Journal of Economic Perspectives* 19(4): 43–66.
- Fehr, E. and Tyran, J.-R. (2007). Money illusion and coordination failure. *Games and Economic Behavior* 58(2): 246–68.
- Fehr, E. and Tyran, J.-R. (2008). Limited rationality and strategic interaction: the impact of the strategic environment on nominal inertia. *Econometrica* 76(2): 353–94.
- Feltovich, N., Iwasaki, A., and Oda, S. H. (2012). Payoff levels, loss avoidance and equilibrium selection in games with multiple equilibria: an experimental study. *Economic Inquiry* 50(4): 932–52.
- Fey, M., McKelvey, R. D., and Palfrey, T. R. (1996). An experimental study of constant-sum centipede games. *International Journal of Game Theory* 25(3): 269–87.
- Fischer, S. (1977). Long-term contracts, rational expectations, and the optimal money supply rule. *Journal of Political Economy* 85(1): 191–205.
- Forsythe, R., Kennan, J., and Sopher, B. (1991). An experimental analysis of strikes in bargaining games with one sided private information. *American Economic Review* 81(1): 253–78.
- Forsythe, R., Nelson, F., Neumann, G. R., and Wright, J. (1992). Anatomy of an experimental political stock market. *American Economic Review* 82(5): 1142–61.
- Fouraker, L. E. and Siegel, S. (1963). *Bargaining Behavior*. New York: McGraw-Hill.
- Friedman, J. W. (1977). *Oligopoly and the Theory of Games*. Amsterdam: North-Holland.
- Fudenberg, D., Levine, D., and Maskin, E. (1994). The Folk Theorem with imperfect public information. *Econometrica* 62: 997–1039.
- Fudenberg, D., Levine, D., and Tirole, J. (1985). Infinite-horizon models of bargaining with one-sided incomplete information. In A. E. Roth (ed.), *Game-Theoretic Models of Bargaining*. Cambridge: Cambridge University Press, pp. 73–98.

- Fudenberg, D. and Maskin, E. (1986). The Folk Theorem for repeated games with discounting or with incomplete information. *Econometrica* 54: 533–54.
- Fudenberg, D. and Tirole, J. (1991). *Game Theory*. Cambridge, MA: MIT Press.
- Gächter, S. and Thöni, C. (2005). Social learning and voluntary cooperation among like-minded people. *Journal of the European Economic Association* 3(2–3): 303–14.
- Garratt, R. J., Walker, M., and Wooders, J. (2012). Behavior in second-price auctions by highly experienced eBay buyers and sellers. *Experimental Economics* 15(1): 44–57.
- Garvin, S. (1993). Experiments in entry deterrence. University of Pittsburgh, Ph.D. Dissertation.
- Geanakoplos, J., Pearce, D., and Stacchetti, E. (1989). Psychological games and sequential rationality. *Games and Economic Behavior* 1(1): 60–79.
- Georganas, S., Healy, P. J., and Weber, R. A. (2012). On the persistence of strategic sophistication. CESifo Group Munich. Working Paper 4653.
- Gibbons, J. D. and Chakraborti, S. (2003). *Nonparametric Statistical Inference*. 4th edition. New York: Dekker.
- Gilboa, I. and Schmeidler, D. (1988). Information dependent games: can common sense be common knowledge? *Economic Letters* 27(3): 215–21.
- Gintis, H. (2009). *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences*. Princeton, NJ: Princeton University Press.
- Glazer, J. and Perry, M. (1996). Virtual implementation in backwards induction. *Games and Economic Behavior* 15(1): 27–32.
- Glazer, J. and Rosenthal, R. W. (1992). A note on Abreu-Matsushima mechanisms. *Econometrica* 60(6): 1435–8.
- Gneezy, U. (2005). Deception: the role of consequences. *American Economic Review* 95(1): 384–94.
- Gneezy, U., Rustichini, A., and Vostroknutov, A. (2010). Experience and insight in the Race game. *Journal of Economic Behavior and Organization* 75(2): 144–55.
- Gode, D. K. and Sunder, S. (1993). Allocative efficiency of markets with zero intelligence traders: market as a partial substitute for individual rationality. *Journal of Political Economy* 101(1): 119–37.
- Goeree, J. K. and Holt, C. A. (2000). Asymmetric inequality aversion and noisy behavior in alternating-offer bargaining games. *European Economic Review* 44(4–6): 1079–89.
- Goeree, J. K. and Holt, C. A. (2001). Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review* 91(5): 1402–22.
- Goeree, J. K. and Holt, C. A. (2004). A model of noisy introspection. *Games and Economic Behavior* 46(2): 365–82.
- Goeree, J. K., Holt, C. A., and Palfrey, T. R. (2002). Quantal response equilibrium and overbidding in private-value auctions. *Journal of Economic Theory* 104(1): 247–72.
- Goeree, J. K., Holt, C. A., and Palfrey, T. R. (2003). Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior* 45(1): 97–113.
- Goeree, J. K., Holt, C. A., and Palfrey, T. R. (2005). Regular quantal response equilibrium. *Experimental Economics* 8(4): 347–67.
- Goeree, J. K. and Offerman, T. (2002). Efficiency in auctions with private and common values: an experimental study. *American Economic Review* 92(3): 625–43.
- Goeree, J. K., Palfrey, T. R., Rogers, B. W., and McKelvey, R. D. (2007). Self-correcting information cascades. *Review of Economic Studies* 74(3): 733–62.
- Goldfarb, A. and Xiao, M. (2011). Who thinks about the competition? Managerial ability and strategic entry in US local telephone markets. *American Economic Review* 101(7): 3130–61.
- Goldfarb, A. and Yang, B. (2009). Are all managers created equal? *Journal of Marketing Research* 46(5): 612–22.
- Grosskopf, B. and Nagel, R. (2008). The two-person beauty contest. *Games and Economic Behavior* 62(1): 93–9.
- Grossman, S. J. and Perry, M. (1986). Sequential bargaining under asymmetric information. *Journal of Economic Theory* 39(1): 120–54.
- Guerra, G. and Zizzo, D. J. (2004). Trust responsiveness and beliefs. *Journal of Economic Behavior and Organization* 55(1): 25–30.
- Gul, F. and Pesendorfer, W. (2007). The case for mindless economics. Mimeo, Princeton University.
- Gul, F., Sonnenschein, H., and Wilson, R. (1986). Foundations of dynamic monopoly and the

- Coase conjecture. *Journal of Economic Theory* 39(1): 155–90.
- Güth, W. (2012). Bargaining and negotiations: what should experimentalists explore more thoroughly? In G. E. Bolton and R. T. A. Croson (eds.), *The Oxford Handbook of Economic Conflict Resolution*. New York: Oxford University Press, pp. 76–88.
- Güth, W., Schmittberger, R., and Schwarze, B. (1982). An experimental analysis of ultimatum bargaining. *Journal of Economic Behavior and Organization* 3(4): 367–88.
- Güth, W. and Tietz, R. (1988). Ultimatum bargaining for a shrinking cake—an experimental analysis. In R. Tietz, W. Albers, and R. Selten (eds.), *Bounded Rational Behavior in Experimental Games and Markets*. Berlin: Springer, pp. 111–28.
- Güth, W. and Tietz, R. (1990). Ultimatum bargaining behavior: a survey and comparison of experimental results. *Journal of Economic Psychology* 11(3): 417–49.
- Guyer, M. J. and Rapoport, A. (1972). Two by two games played once. *Journal of Conflict Resolution* 16(3): 409–31.
- Haile, P. A., Hortaçsu, A., and Kosenok, G. (2008). On the empirical content of quantal response equilibrium. *American Economic Review* 98(1): 180–200.
- Halpern, J. Y. and Pass, R. (2012). Iterated regret minimization: a new solution concept. *Games and Economic Behavior* 74(1): 184–207.
- Haltiwanger, J. and Waldman, M. (1989). Limited rationality and strategic complements: the implications for macroeconomics. *Quarterly Journal of Economics* 104(3): 463–83.
- Hansen, R. G. and Lott, J. R. Jr. (1991). The winner's curse and public information in common value auctions: comment. *American Economic Review* 81(81): 347–61.
- Hargreaves Heap, S., Rojo-Arjona, D., and Sugden, R. (2014). How portable is level-0 behavior? A test of level-k theory in games with non-neutral frames. *Econometrica* 82(3): 1133–51.
- Harrison, G. W. and List, J. A., (2004). Field experiments. *Journal of Economic Literature* 42: 1013–59.
- Harrison, G. W. and List, J. A. (2008). Naturally occurring markets and exogenous laboratory experiments: a case study of the winner's curse. *Economic Journal* 118: 822–43.
- Harsanyi, J. C. (1973). Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points. *International Journal of Game Theory* 2(1): 1–23.
- Harsanyi, J. C. and Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*. Cambridge, MA: MIT Press.
- Harstad, R. M. and Selten, R. (2013). Bounded-rationality models: tasks to become intellectually competitive. *Journal of Economic Literature* 51(2): 496–511.
- Heinemann, F., Nagel, R., and Ockenfels, P. (2004). The theory of global games on test: experimental analysis of coordination games with public and private information. *Econometrica* 72(5): 1583–99.
- Hermalin, B. E. (1998). Toward an economic theory of leadership: leading by example. *American Economic Review* 88(5): 1188–1206.
- Ho, T., Camerer, C. F., and Weigelt, K. (1998). Iterated dominance and iterated best response in experimental “p-beauty contests.” *American Economic Review* 88(4): 947–69.
- Hoffman, E., McCabe, K., Shachat, K., and Smith, V. (1994). Preferences, property rights, and anonymity in bargaining games. *Games and Economic Behavior* 7(3): 346–80.
- Holt, C. A., Jr. and Sherman, R. (1994). The laser's curse and bidder's bias. *American Economic Review* 84: 642–52.
- Hotelling, H. (1929). Stability in competition, *Economic Journal* 39(153): 41–57.
- Huck, S., Jehiel, P., and Rutter, T. (2011). Feedback spillover and analogy-based expectations: a multi-game experiment. *Games and Economic Behavior* 71(2): 351–65.
- Hurwicz, L., Radner, R., and Reiter, S. (1975). A stochastic decentralized resource allocation process: part I. *Econometrica* 43(2): 187–221.
- Isoni, A., Poulsen, A., Sugden, R., and Tsutsui, K. (2013). Focal points in tacit bargaining problems: experimental evidence. *European Economic Review* 59: 167–88.
- Ivanov, A., Levin, D., and Niederle, M. (2010). Can relaxation of beliefs rationalize the winner's curse? An experimental study. *Econometrica* 78(4): 1435–52.
- Jehiel, P. (2005). Analogy-based expectation equilibrium. *Journal of Economic Theory* 123(2): 81–104.

- Jehiel, P. and Koessler, F. (2008). Revisiting games of incomplete information with analogy-based expectations. *Games and Economic Behavior* 62(2): 533–57.
- Johnson, E. J., Camerer, C. F., Sen, S., and Rymon, T. (2002). Detecting failures of backward induction: monitoring information search in sequential bargaining. *Journal of Economic Theory* 104(1): 16–47.
- Jung, Y. J., Kagel, J. H., and Levin, D. (1994). On the existence of predatory pricing: an experimental study of reputation and entry deterrence in the chain-store game. *RAND Journal of Economics* 25(1): 72–93.
- Kagel, J. H. (2015). Laboratory experiments: The lab in relationship to field experiments, field data, and economic theory. In G. R. Fréchette and A. Schotter (eds.), *Handbook of Experimental Economic Methodology*. Oxford: Oxford University Press, pp. 339–59.
- Kagel, J. H., Harstad, R. M., and Levin, D. (1987). Information impact and allocation rules in auctions with affiliated private values: a laboratory study. *Econometrica* 55(6): 1275–304.
- Kagel, J. H. and Levin, D. (1986). The winner's curse and public information in common value auctions. *American Economic Review* 76(5): 894–920.
- Kagel, J. H. and Levin, D. (1991). The winner's curse and public information in common value auctions: reply. *American Economic Review* 81(1): 362–9.
- Kagel, J. H. and Levin, D. (1993). Independent private value auctions: bidder behavior in first-, second- and third-price auctions with varying numbers of bidders. *Economic Journal* 103(419): 868–79.
- Kagel, J. H. and Levin, D. (2001). Behavior in multi-unit demand auctions: experiments with uniform price and dynamic Vickrey auctions. *Econometrica* 69(2): 413–54.
- Kagel, J. H. and Levin, D. (2010). Auctions: a survey of experimental research. In J. H. Kagel and A. E. Roth (eds.), *The Handbook of Experimental Economics, Volume II*. Princeton, NJ: Princeton University Press, pp. 501–86.
- Kagel, J. H., Levin, D., and Harstad, R. M. (1995). Comparative static effects of number of bidders and public information in second-price common value auctions. *International Journal of Game Theory* 24(3): 293–319.
- Kagel, J. H. and Richard, J.-F. (2001). Super-experienced bidders in first-price common value auctions: rules of thumb, Nash equilibrium bidding and the winner's curse. *Review of Economics and Statistics* 83(3): 408–19.
- Kagel, J. H. and Roth, A. E. (1995). *The Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press.
- Kahneman, D. (1988). Experimental economics: a psychological perspective. In R. Tietz, W. Albers, and R. Selten (eds.), *Bounded Rational Behavior in Experimental Games and Markets: Proceedings of the Fourth Conference on Experimental Economics, Bielefeld, West Germany, September 21–25, 1986*. New York, Berlin; London and Tokyo: Springer, pp. 11–18.
- Kahneman, D. (2003). Maps of bounded rationality: psychology for behavioral economics. *American Economic Review* 93(5): 1449–75.
- Kahneman, D., Slovic, P., and Tversky, A. (eds.) (1982). *Judgement Under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.
- Kahneman, D. and Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica* 47(2): 263–92.
- Kahneman, D. and Tversky, A. (2000). *Choices, Values and Frames*. Cambridge: Cambridge University Press.
- Kalai, E. and Smorodinsky, M. (1975). Other solutions to Nash's bargaining problem. *Econometrica* 43: 513–18.
- Katok, E., Sefton, M., and Yavas, A. (2002). Implementation by iterative dominance and backward induction: an experimental comparison. *Journal of Economic Theory* 104(1): 89–103.
- Kawagoe, T. and Takizawa, H. (2012). Level-k analysis of experimental centipede games. *Journal of Economic Behavior and Organization*. 82(2–3): 548–66.
- Kennan, J. and Wilson, R. (1990). Theories of bargaining delays. *Science* 249(4973): 1124–8.
- Keynes, J. M. (1936). *The General Theory of Employment, Interest, and Money*. New York: Harcourt Brace and Co.
- Khadjavi, M. and Lange, A. (2013). Prisoners and their dilemma. *Journal of Economic Behavior and Organization* 92: 163–75.
- Khalmetski, K., Ockenfels, A., and Werner, P. (2015). Surprising gifts. *Journal of Economic Theory* 159: 163–208.

- Kohlberg, E. and Mertens, J.-F. (1986). On the strategic stability of equilibria. *Econometrica* 54(5): 1003–37.
- Kooreman, P., Faber, R. P., and Hofmans, H. M. (2004). Charity donations and the euro introduction: some quasi-experimental evidence on money illusion. *Journal of Money, Credit, and Banking* 36(6): 1121–224.
- Kőszegi, B. and Rabin, M. (2006). A model of reference-dependent preferences. *Quarterly Journal of Economics* 121(4): 1133–65.
- Kreps, D. M. and Wilson, R. (1982). Sequential equilibria. *Econometrica* 50(4): 863–94.
- Krishna, V. (2010). *Auction Theory*. 2nd edition. San Diego: Academic Press.
- Krueger, J. I. (2007). From social projection to social behaviour. *European Review of Social Psychology* 18(1): 1–35.
- Kübler, D. and Weizsäcker, G. (2004). Limited depth of reasoning and failure of cascade formation in the laboratory. *Review of Economic Studies* 71(2): 425–41.
- Kuhn, H. W. (1953). Extensive games and the problem of information. In H. W. Kuhn and W. W. Tucker (eds.), *Contributions to the Theory of Games, Volume II*. Princeton, NJ: Princeton University Press, pp. 193–216.
- Ledyard, J. (1995). Public goods: a survey of experimental research. In J. H. Kagel and A. E. Roth (eds.), *The Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press, pp. 111–194.
- Leininger, W., Linhart, P. B., and Radner, R. (1989). Equilibria of the sealed-bid mechanism for bargaining with incomplete information. *Journal of Economic Theory* 48(1): 63–106.
- Levitt, S. D., List, J. A., and Reiley, D. H. (2010). What happens in the field stays in the field: professionals do not play minimax in laboratory experiments. *Econometrica* 78(4): 1413–34.
- Levitt, S. D., List, J. A., and Sadoff, S. E. (2011). Checkmate: exploring backward induction among chess players. *American Economic Review* 101(2): 975–990.
- Lewis, D. K. (1979). Prisoners' dilemma is a Newcomb problem. *Philosophy and Public Affairs* 8(3): 235–40.
- Lind, B. and Plott, C. R. (1991). The winner's curse: experiments with buyers and with sellers. *American Economic Review* 81(1): 335–46.
- List, J. A. (2004). Testing neoclassical competitive theory in multilateral decentralized markets. *Journal of Political Economy* 112(5): 1131–56.
- List, J. A. and Lucking-Reiley, D. (2002). The effects of seed money and refunds on charitable giving: experimental evidence from a university capital campaign. *Journal of Political Economy* 110(1): 215–33.
- Loewenstein, G. F., Issacharoff, S., Camerer, C. F., and Babcock, L. (1993). Self-serving assessments of fairness and pretrial bargaining. *Journal of Legal Studies* 22(1): 135–59.
- Loewenstein, G. F., Thompson, L., and Bazerman, M. H. (1989). Social utility and decision making in interpersonal contexts. *Journal of Personality and Social Psychology* 57(3): 426–41.
- Lucas, R. E., Jr. (1972). Expectations and the neutrality of money. *Journal of Economic Theory* 4(2): 103–24.
- McCabe, K. A., Mukherji, A., and Runkle, D. E. (2000). An experimental study of information and mixed-strategy play in the three-person matching-pennies game. *Economic Theory* 15(2): 421–62.
- McKelvey, R. D. and Palfrey, T. R. (1992). An experimental study of the centipede game. *Econometrica* 60(4): 803–836.
- McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior* 10(1): 6–38.
- McKelvey, R. D. and Palfrey, T. R. (1998). Quantal response equilibrium in extensive form games. *Experimental Economics* 1(1): 9–41.
- McKelvey, R. D., Palfrey, T. R., and Weber, R. A. (2000). The effects of payoff magnitude and heterogeneity on behavior in 2×2 games with unique mixed strategy equilibria. *Journal of Economic Behavior and Organization* 42(4): 523–48.
- Mailath, G. J., Okuno-Fujiwara, M., and Postlewaite, A. (1993). Belief-based refinements in signalling games. *Journal of Economic Theory* 60(2): 241–76.
- Mankiw, N. G. and Reis, R. (2002). Sticky information versus sticky prices: a proposal to replace the new Keynesian Phillips curve. *Quarterly Journal of Economics* 117(4): 1295–328.
- Maschler, M., Solan, E., and Zamir, S. (2013). *Game Theory*. Cambridge: Cambridge University Press.

- Mas-Colell, A., Whinston, M. D., and Green, J. (1995). *Microeconomic Theory*. New York: Oxford University Press.
- Mead, W. J., Moseidjord, A., and Sorensen, P. E. (1983). The rate of return earned by leases under cash bonus bidding in OCS oil and gas leases. *Energy Journal* 4(4): 37–52.
- Mehta, J., Starmer, C., and Sugden, R. (1994). The nature of salience: an experimental investigation of pure coordination games. *American Economic Review* 84(3): 658–73.
- Milgrom, P. (1989). Auctions and bidding: a primer. *Journal of Economic Perspectives* 3(3): 3–22.
- Milgrom, P. (2004). *Putting Auction Theory to Work*. Cambridge: Cambridge University Press.
- Milgrom, P. and Roberts, J. (1982). Limit pricing and entry under incomplete information: an equilibrium analysis. *Econometrica* 50(2): 443–459.
- Milgrom, P. and Roberts, J. (1991). Adaptive and sophisticated learning in normal form games. *Games and Economic Behavior* 3(1): 82–100.
- Miller, R. M. (2008). Don't let your robots grow up to be traders: artificial intelligence, human intelligence, and asset-market bubbles. *Journal of Economic Behavior and Organization* 68(1): 153–66.
- Miller, R. M., Plott, C. R., and Smith, V. L., (1977). Intertemporal competitive equilibrium: an empirical study of speculation. *Quarterly Journal of Economics* 91(4): 599–624.
- Moinas, S. and Pouget, S. (2013). The bubble game: an experimental study of speculation. *Econometrica* 81(4): 1507–39.
- Monderer, D. and Samet, D. (1989). Approximating common knowledge with common beliefs. *Games and Economic Behavior* 1(2): 170–90.
- Mookherjee, D. and Sopher, B. (1994). Learning behavior in an experimental matching pennies game. *Games and Economic Behavior* 7(1): 62–91.
- Mookherjee, D. and Sopher, B. (1997). Learning and decision costs in experimental constant sum games. *Games and Economic Behavior* 19(1): 97–132.
- Moreno, D. and Wooders, J. (1998). An experimental study of communication and coordination in non-cooperative games. *Games and Economic Behavior* 24(1–2): 47–76.
- Morris, S. and Shin, H. S. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review* 88(3): 587–97.
- Murnighan, J. K., Roth, A. E., and Schoumaker, F. (1988). Risk aversion in bargaining: an experimental study. *Journal of Risk and Uncertainty* 1(1): 101–24.
- Myers, S. C. and Majluf, N. S. (1984). Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics* 13(2): 187–221.
- Myerson, R. B. (1979). Incentive compatibility and the bargaining problem. *Econometrica* 47(1): 61–73.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research* 6(1): 58–73.
- Myerson, R. B. (1984). Two person bargaining problems with incomplete information. *Econometrica* 52(2): 461–88.
- Myerson, R. B. and Satterthwaite, M. A. (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory* 29(2): 265–81.
- Nagel, R. (1995). Unraveling in guessing games: an experimental study. *American Economic Review* 85(5): 1313–26.
- Nagel, R. and Tang, F. F. (1988). Experimental results on the centipede game in normal form: an investigation on learning. *Journal of Mathematical Psychology* 42(2–3): 356–84.
- Nash, J. F. (1950a). Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences of the United States of America* 36(1): 48–9.
- Nash, J. F. (1950b). The bargaining problem. *Econometrica* 18(2): 155–62.
- Nash, J. F. (1953). Two-person cooperative games. *Econometrica* 21: 128–40.
- Neelin, J., Sonnenschein, H., and Spiegel, M. (1988). A further test of noncooperative bargaining theory: comment. *American Economic Review* 78(4): 824–36.
- Neral, J. and Ochs, J. (1992). The sequential equilibrium theory of reputation building: a further test. *Econometrica* 60(5): 1151–69.
- Neugebauer, T. and Selten, R. (2006). Individual behavior and first-price auctions: the importance of information feedback in computerized experimental markets. *Games and Economic Behavior* 54(1): 183–204.

- Nöldeke, G. and Samuelson, L. (1994). Learning to signal in markets. University of Bonn. Working Paper 271.
- Nydegger, R. V. and Owen, G. (1975). Two person bargaining: an experimental of the Nash axioms. *International Journal of Game Theory* 3: 239–49.
- Ochs, J. (1995). Games with unique, mixed strategy equilibria: an experimental study. *Games and Economic Behavior* 10(1): 202–17.
- Ochs, J. and Roth, A. E. (1989). An experimental study of sequential bargaining. *American Economic Review* 79(3): 355–84.
- Okano, Y. (2013). Minimax play by teams. *Games and Economic Behavior* 77(1): 168–80.
- O'Neill, B. (1987). Nonparametric test of the minimax theory of two-person zerosum games. *Proceedings of the National Academy of Sciences of the United States of America* 84(7): 2106–9.
- Palacios-Huerta, I. (2003). Professionals play minimax. *Review of Economic Studies* 70(2): 395–415.
- Palacios-Huerta I. and Volij, O. (2008). Experientia docet: professionals play minimax in laboratory experiments. *Econometrica* 76(1): 71–115.
- Payne, J. W., Bettman, J. R., and Johnson, E. J. (1993). *The Adaptive Decision Maker*. New York: Cambridge University Press.
- Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52(4): 1029–50.
- Penczynski, S. P. (2012). Persuasion: an experimental study of team decision making. University of Mannheim, Department of Economics. Working Paper, Unpublished.
- Plott, C. R. (2008). Properties of disequilibrium adjustment in double auction markets. In C. R. Plott and V. L. Smith (eds.), *Handbook of Experimental Economics Results, Volume 1, Part 1*. Amsterdam: North-Holland, pp. 16–21.
- Plott, C. R. and Smith, V. L. (eds.) (2008). *Handbook of Experimental Economics Results, Volume 1*. Amsterdam: North-Holland.
- Polak, B. (1999). Epistemic conditions for Nash equilibrium, and common knowledge of rationality. *Econometrica* 67(3): 673–6.
- Ponti, G. (2000). Cycles of learning in the centipede game. *Games and Economic Behavior* 30(1): 115–41.
- Potters, J., Sefton, M., and Vesterlund, L. (2007). Leading-by-example and signaling in voluntary contribution games: an experimental study. *Economic Theory* 33(1): 169–82.
- Potters, J. and Suetens, S. (2009). Cooperation in experimental games of strategic complements and substitutes. *Review of Economic Studies* 76(3): 1125–47.
- Quattrone, G. A., and Tversky, A. (1984). Causal versus diagnostic contingencies: on self deception and the voter's illusion. *Journal of Personality and Social Psychology* 46(2): 237–48.
- Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review* 83(5): 1281–302.
- Rabin, M. (1994). A model of pre-game communication. *Journal of Economic Theory* 63(2): 370–91.
- Radner, R. and Schotter, A. (1989). The sealed-bid mechanism: an experimental study. *Journal of Economic Theory* 48(1): 179–220.
- Raiffa, H. (1953). Arbitration schemes for generalized two-person games. In H. W. Kuhn and A. W. Tucker (eds.), *Contributions to the Theory of Games, II, Annals of Mathematics Studies*, 28. Princeton, NJ: Princeton University Press, pp. 361–87.
- Rapoport, A. (1966). *Two-Person Game Theory*. Ann Arbor, MI: Michigan University Press.
- Rapoport, A. (1988). Experiments with N-person social traps I: prisoners' dilemma, weak prisoners' dilemma, volunteers' dilemma, and largest number. *Journal of Conflict Resolution* 32(3): 457–72.
- Rapoport, A. and Amaldoss, W. (2000). Mixed strategies and iterative elimination of strongly dominated strategies: an experimental investigation of states of knowledge. *Journal of Economic Behavior and Organization* 42(4): 483–521.
- Rapoport, A. and Amaldoss, W. (2004). Mixed strategy play in single-stage all-pay auctions with symmetric players. *Journal of Economic Behavior and Organization* 54(4): 585–607.
- Rapoport, A. and Amaldoss, W. (2008). In search of experimental support for an asymmetric equilibria solution in symmetric investment games. *Journal of Economic Behavior and Organization* 67(1): 200–3.
- Rapoport, A. and Boebel, R. B. (1992). Mixed strategies in strictly competitive games: a further test in the minimax hypothesis. *Games and Economic Behavior* 4(2): 261–83.
- Rapoport, A. and Budescu, D. V. (1992). Generation of random series in two person strictly competitive games: an experimental study. *Economic Theory* 33(1): 169–82.

- itive games. *Journal of Experimental Psychology: General* 121(3): 352–63.
- Rapoport, A. and Erev, I. (1998). Coordination, “magic” and reinforcement learning in a market entry game. *Games and Economic Behavior* 23(2): 146–75.
- Rapoport, A., Erev, I., and Zwick, R. (1995). An experimental study of buyer-seller negotiation with one-sided incomplete information and time discounting. *Management Science* 41(3): 377–94.
- Rapoport, A., Frenkel, O., and Perner, J. (1977). Experiments with cooperative 2×2 games. *Theory and Decision* 8(1): 67–92.
- Rapoport, A., and Fuller, M. A. (1995). Bidding strategies in a bilateral monopoly with two-sided incomplete information. *Journal of Mathematical Psychology* 39(2): 179–96.
- Rapoport, A., Seale, D. A., Erev, I., and Sundali, J. A. (1998). Equilibrium play in large group market entry games. *Management Science* 44(1): 129–41.
- Rapoport, A., Seale, D. A., and Winter, E. (2002). Coordination and learning behavior in large groups with asymmetric players. *Games and Economic Behavior* 39(1): 111–36.
- Rapoport, A., Stein, W. E., Parco, J. E., and Nicholas, T. E. (2003). Equilibrium play and adaptive learning in a three-person centipede game. *Games and Economic Behavior* 43(2): 239–65.
- Reis, R. (2006). Inattentive consumers. *Journal of Monetary Economics* 53(8): 1761–800.
- Riley, J. G. and Samuelson, W. F. (1981). Optimal auctions. *American Economic Review* 71(3): 381–92.
- Robbins, J. M. and Krueger, J. I. (2005). Social projection to ingroups and outgroups: a review and meta-analysis. *Personality and Social Psychology Review* 9(1): 32–47.
- Rogers, B. W., Palfrey, T. R., and Camerer, C. F. (2009). Heterogeneous quantal response equilibrium and cognitive hierarchies. *Journal of Economic Theory* 144(4): 1440–67.
- Romer, D. (2006). *Advanced Macroeconomics*. New York: McGraw-Hill/Irwin.
- Rosenthal, R. W. (1981). Games of perfect information, predatory pricing and chain store paradox. *Journal of Economic Theory* 25(1): 92–100.
- Rosenthal, R. W., Shachat, J., and Walker, M. (2003). Hide and seek in Arizona. *International Journal of Game Theory* 32: 273–93.
- Ross, L., Greene, D., and House, P. (1977). The “false consensus effect”: an egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology* 13(3): 279–301.
- Roth, A. E. (1977). Independence of irrelevant alternatives and solutions to Nash’s bargaining problem. *Journal of Economic Theory* 16(2): 247–51.
- Roth, A. E. (1978). The Nash solution and the utility of bargaining. *Econometrica* 46(3–4): 587–94, 983.
- Roth, A. E. (1985). Toward a focal-point theory of bargaining. In A. E. Roth (ed.), *Game-Theoretic Models of Bargaining*. Cambridge: Cambridge University Press, pp. 259–68.
- Roth, A. E. (1995). Bargaining experiments. In J. H. Kagel and A. E. Roth (eds.), *Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press, pp. 253–347.
- Roth, A. E. and Malouf, M. W. K. (1979). Game theoretic models and the role of information in bargaining. *Psychological Review* 86(6): 574–94.
- Roth, A. E., Malouf, M. W. K., and Murnighan, J. K. (1981). Sociological versus strategic factors in bargaining. *Journal of Economic Behavior and Organization* 2(2): 153–77.
- Roth, A. E. and Murnighan, J. K. (1982). The role of information in bargaining: an experimental study. *Econometrica* 50: 1123–42.
- Roth, A. E., Murnighan, J. K., and Schoumaker, F. (1988). The deadline effect in bargaining: some experimental evidence. *American Economic Review* 78(4): 806–23.
- Roth, A. E., Prasnikar, V., Okuno-Fujiwara, M., and Zamir, S. (1991). Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: an experimental study. *American Economic Review* 81(5): 1068–95.
- Rubinstein, A. (1982). Perfect equilibrium in a bargaining model. *Econometrica* 50(1): 97–109.
- Rubinstein, A. (1999). Experience from a course in game theory: pre and post-class problem sets as a didactic device. *Games and Economic Behavior* 28(1): 155–70.
- Rubinstein, A., Tversky, A., and Heller, D. (1996). Naïve strategies in competitive games. In W. Albers, W. Güth, P. Hammerstein, B. Moldovanu, and E. van Damme (eds.), *Understanding Strategic Interaction—Essays in Honor of Reinhard Selten*. Berlin: Springer-Verlag, pp. 394–402.

- Samuelson, W. F. and Bazerman, M. H. (1985). The winner's curse in bilateral negotiations. In V. L. Smith (ed.), *Research in Experimental Economics, Volume 3*. Greenwich, CT: JAI Press.
- Schelling, T. C. (1960). *The Strategy of Conflict*. Cambridge, MA: Harvard University Press.
- Schelling, T. C. (1978). *Micromotives and Macrobehavior*. New York: W.W. Norton.
- Schotter, A. and Sopher, B. (2000). Social learning and coordination conventions in inter-generational games: an experiment in Lamarckian evolutionary dynamics. C. V. Starr Center for Applied Economics, Working Paper 00-01.
- Schotter, A. and Sopher, B. (2003). Social learning and coordination conventions in inter-generational games: an experimental study. *Journal of Political Economy* 111(3): 498-529.
- Schotter, A. and Sopher, B. (2007). Advice and behavior in intergenerational ultimatum games: an experimental approach. *Games and Economic Behavior* 58(2): 365-93.
- Schotter, A., Weigelt, K., and Wilson, C. (1994). A laboratory investigation of multi-person rationality and presentation effects. *Games and Economic Behavior* 6(3): 445-68.
- Scroggin, S. (2007). Exploitable actions of believers in the "law of small numbers" in repeated constant-sum games. *Journal of Economic Theory* 133(1): 219-235.
- Seely, B., Van Huyck, J. B., and Battalio, R. C. (2005). Credible assignments can improve efficiency in laboratory public goods games. *Journal of Public Economics* 89(8): 1437-55.
- Sefton, M. and Yavas, A. (1996). Abreu-Matsushima mechanisms: experimental evidence. *Games and Economic Behavior* 16(2): 280-302.
- Selten, R. (1975). Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* 4(1): 25-55.
- Selten, R. (1978). The chain store paradox. *Theory and Decision* 9(2): 127-59.
- Selten, R. (1998). Features of experimentally observed bounded rationality. *European Economic Review* 42(3-5): 413-36.
- Selten, R. and Güth, W. (1982). Equilibrium point selection in a class of market entry games. In M. Diestler, E. Furst, and G. Schwodiauer (eds.), *Games, Economic Dynamics, and Time Series Analysis*. Wien-Wurzburg: Physica-Verlag.
- Shachat, J. M. (2002). Mixed strategy play and the minimax hypothesis. *Journal of Economic Theory* 104(1): 189-226.
- Shachat, J. M. and Swarthout, J. T. (2004). Do we detect and exploit mixed strategy play by opponents? *Mathematical Methods of Operations Research* 59(3): 359-73.
- Shachat, J. M., and Swarthout, J. T. (2008). How do people play against Nash opponents in games which have a mixed strategy equilibrium? Georgia State University, Experimental Economics Center, Andrew Young School of Policy Studies. Working Paper 2008-07.
- Shafir, E., Diamond, P. A., and Tversky, A. (1997). Money illusion. *Quarterly Journal of Economics* 112(2): 341-74.
- Shaked, A. (1982). Existence and computation of mixed strategy Nash equilibrium for 3-firms location problem. *Journal of Industrial Economics* 31(1-2): 93-97.
- Shalev, J. (2000). Loss aversion equilibrium. *International Journal of Game Theory* 29(2): 269-87.
- Shiller, R. J. (2000). *Irrational Exuberance*. Princeton, NJ: Princeton University Press.
- Siegel, S. and Fouraker, L. E. (1960). *Bargaining and Group Decision Making*. New York: McGraw-Hill.
- Signorino, C. S. (1999). Strategic interaction and the statistical analysis of international conflict. *American Political Science Review* 93(2): 279-97.
- Slonim, R. L., Roth A. E., and Erev, I. (2003). Regularities in play of zero-sum games with unique mixed-strategy equilibria. Mimeo, Case Western Reserve University.
- Smith, V. L. (1962). An experimental study of competitive market behavior. *Journal of Political Economy* 70(2): 111-37.
- Smith, V. L. (1964). Effect of market organization on competitive equilibrium. *Quarterly Journal of Economics* 78(2): 181-201.
- Smith, V. L. and Sørensen, P. (2000). Pathological outcomes of observational learning. *Econometrica* 68(2): 371-98.
- Smith, V. L., Suchanek, G. L., and Williams, A. W. (1988). Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica* 56(5): 1119-51.

- Soetevent, A. R. (2005). Anonymity in giving in a natural context—a field experiment in 30 churches. *Journal of Public Economics* 89(11–12): 2301–23.
- Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics* 87(3): 355–74.
- Stahl, D. O. and Wilson, P. W. (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior and Organization* 25(3): 309–27.
- Stahl, D. O. and Wilson, P. W. (1995). On players' models of other players: theory and experimental evidence. *Games and Economic Behavior* 10(1): 218–54.
- Ståhl, I. (1972). *Bargaining Theory*. Stockholm: Stockholm School of Economics.
- Straub, P. G. (1995). Risk dominance and coordination failure in static games. *The Quarterly Review of Economics and Finance* 35(4): 339–63.
- Sugden, R. (1993). Thinking as a team: towards an explanation of non-selfish behavior. *Social Philosophy and Policy* 10(1): 69–89.
- Sugden, R. (1995). A theory of focal points. *Economic Journal* 105(430): 533–50.
- Sugden, R. (2003). Reference-dependent subjective expected utility. *Journal of Economic Theory* 111(2): 172–91.
- Sundali, J. A., Rapoport, A., Seale, D. A. (1995). Coordination in market entry games with symmetric players. *Organizational Behavior and Human Decision Processes* 64(2): 203–18.
- Tesfatsion, L. and Judd, K. L. (eds.) (2006). *Handbook of Computational Economics*, Vol. 2: *Agent-Based Computational Economics*. Amsterdam: North-Holland.
- Thaler, R. H. (1988). Anomalies: the winner's curse. *Journal of Economic Perspectives* 2(1): 191–202.
- Topkis, D. M. (1998). *Supermodularity and Complementarity*. Princeton, NJ: Princeton University Press.
- Tor, A. and Bazerman, M. H. (2003). Focusing failures in competitive environments: explaining decision errors in the Monty Hall game, the acquiring a company game, and multiparty ultimatum. *Journal of Behavioral Decision Making* 16(5): 353–74.
- Tsay, C.-J. and Bazerman, M. H. (2009). A decision-making perspective to negotiation: a review of the past and a look into the future. *Negotiation Journal* 25(4): 467–80.
- Valley, K., Thompson, L., Gibbons, R., and Bazerman, M. H. (2002). How communication improves efficiency in bargaining games. *Games and Economic Behavior* 38(1): 127–55.
- Van Boening, M. V. and Wilcox, N. T. (1996). Avoidable cost: ride a double auction roller coaster. *American Economic Review* 86(3): 461–77.
- Van Boven, L. and Loewenstein, G. (2005). Cross-situational projection. In M. Alicke, J. Krueger, and D. Dunning (eds.), *The Self in Social Judgment*. New York: Psychology Press, pp. 43–64.
- Van Damme, E. (1987). *Stability and Perfection of Nash Equilibria*. Berlin: Springer Verlag.
- Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *American Economic Review* 80(1): 234–48.
- Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1991). Strategic uncertainty, equilibrium selection, and coordination failure in average opinion games. *Quarterly Journal of Economics* 106(3): 885–910.
- Van Huyck, J. B., Battalio, R. C., and Beil, R. O. (1993). Asset markets as an equilibrium selection mechanism: coordination failure, game form auctions, and tacit communication. *Games and Economic Behavior* 5(3): 485–504.
- Van Huyck, J. B., Cook, J. P., and Battalio, R. C. (1997). Adaptive behavior and coordination failure. *Journal of Economic Behavior and Organization* 32(4): 483–503.
- Van Huyck, J. B., Gillette, A. B., and Battalio, R. C. (1992). Credible assignments in coordination games. *Games and Economic Behavior* 4(4): 606–26.
- Varian, H. (1992). *Microeconomic Analysis*. New York: W.W. Norton.
- Vesterlund, L. (2003). The informational value of sequential fundraising. *Journal of Public Economics* 87(3–4): 627–57.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16(1): 8–37.
- von Neumann, J. (1928). Zur theorie der gesellschaftsspiele. *Mathematische Annalen* 100(1): 295–320.
- Walker, M. and Wooders, J. (2001). Minimax play at Wimbledon. *American Economic Review* 91(5): 1521–38.

- Weber, R. A. (2006). Managing growth to achieve efficient coordination in large groups. *American Economic Review* 96(1): 114–26.
- Weg, E., Rapoport, A., and Felsenthal, D. S. (1990). Two-person bargaining behavior in fixed discounting factor games with infinite horizon. *Games and Economic Behavior* 2(2): 76–95.
- Wooders, J. (2010). Does experience teach? Professionals and minimax play in the lab. *Econometrica* 78(3): 1143–54.
- Yildiz, M. (2007). Wishful thinking in strategic environments. *Review of Economic Studies* 74(1): 319–44.
- Zauner, K. G. (1999). A payoff uncertainty explanation of results in experimental centipede games. *Games and Economic Behavior* 26(1): 157–85.
- Zhong, C.-B., Loewenstein, J., and Murnighan, J. K. (2007). Speaking the same language: the cooperative effects of labeling in the prisoner's dilemma. *Journal of Conflict Resolution* 51(3): 431–56.
- Zwick, R., Rapoport, A., and Howard, J. C. (1992). Two-person sequential bargaining behavior with exogenous breakdown. *Theory and Decision* 32(3): 241–68.

PART 5

Behavioral Models of Learning

INTRODUCTION TO PART 5

In neoclassical economics, one takes preferences as the primitives. Conditional on these preferences, economic actors choose the set of best elements from some well-defined choice set. For instance, consumers may engage in constrained utility maximization, and firms may engage in constrained profit maximization. However, this leaves open the question of the source of these preferences. It is also typical in economics to focus, unlike in biology and physics, mainly on equilibrium analysis. But equilibrium analysis often does not address the question of the time path to an equilibrium. Indeed the temporal learning behavior of economic agents may offer illuminating insights on what sort of equilibria, if any, are likely to be achieved. In Part 5 of the book we address these questions.

Part 5 is split into three chapters. In Chapter 14, we introduce *evolutionary game theory*, and use it to explore the source of human preferences with a focus on *other-regarding preferences* that is the subject matter of Part 2 of the book. In these models, learning typically takes place by random mutations at the level of the gene, as a response to changes in the environment.¹ In some sense, this may be considered the simplest model of learning that is relevant for human choice.

In evolutionary game theory, the animal's genes, or its *genotype*, determine its behavior, or its *phenotype*. In any given environment, the competing phenotypes have different degrees of success. In a Darwinian selection mechanism, more successful phenotypes have greater *fitness*, or reproductive success. Hence, they are more likely to proliferate in the next period. The currently dominant phenotype in the population may be termed as the *resident phenotype*. The basic question in evolutionary game theory is: Can the resident phenotype be invaded by any *mutant phenotype*? If not, then the resident phenotype is known as an *evolutionary stable strategy* (ESS). We provide the motivation for an ESS, and give the necessary criteria, the primary criterion and the secondary criterion, to check if a strategy is an ESS. We also show how the concept of ESS can be extended to asymmetric information about the types of players.

The concept of an ESS is silent about the underlying dynamics that give rise to the equilibrium. An important advance, the *replicator dynamics*, clarified the dynamic process underlying an ESS. The replicator dynamics captures, in a simple model, the idea that fitter phenotypes will have a higher growth rate in the population. We explore the properties of replicator dynamics, and derive the following nested structure. The steady states of replicator dynamics nest the set of Nash equilibria that in turn nest, sequentially, asymptotically stable states of replicator dynamics and ESS.

Within the framework of evolutionary game theory, we ask some fundamental questions. These questions include the following: Why do we observe prosocial human behavior such as altruism, and the desire to cooperate, and accept norms? Why do we observe reciprocity, and other-regarding preferences? If selfishness has higher fitness, how can altruism survive? How is *culture* transmitted across generations? Can self-regarding preferences be a stable outcome of an evolutionary social system?

¹ We omit a discussion of the relation between *molecular genetics* and economics, which supposes that the genetic endowment of individuals can explain the observed individual-level economic outcomes; for a recent survey see Benjamin et al. (2012).

We discuss the inter-generational transmission of culture through a model of *gene-culture co-evolution*. We also consider a range of explanations for prosocial behavior. These include: *kin selection* and *inclusive fitness*, *group selection* and *multilevel selection*, *reciprocal altruism* or *direct reciprocity*, *indirect reciprocity*, *repeated game considerations*, and the *indirect evolutionary approach*. Models of *social imitation* in which humans consciously choose their strategies based on simple principles of *social imitation* and *emulation*, can produce replicator-like dynamics. However, these results are not very robust to changes in the mechanics of social imitation.

One limitation of models of evolutionary game theory is that they cannot explain the rapid rate of learning in experiments because evolution, by its very nature, is a slow and conservative process. In Chapter 15 we explore the theory and empirics of human learning that is much faster relative to learning at the level of the gene. We shall focus on deterministic, and typically stationary, environments. This chapter also allows us to explore the claim that a Nash equilibrium is the product of some underlying learning mechanism. The results of the experiments provide, at best, mixed support for Nash equilibrium. However, in a large number of cases, too large for Nash equilibrium to be taken as the dominant equilibrium concept, play does not conform to a Nash equilibrium. We shall explore exciting new theoretical developments that have made an attempt to organize the learning data.

The most basic learning model that we consider is the *reinforcement model of learning*, also known as *stimulus learning* and *rote learning*. In this model, players begin with some *propensities* to play their pure strategies. They then observe the payoffs from their pure strategies, and update their propensities over time. Propensities are converted into probabilities of play of various strategies. The rationality requirements in this model are rather limited. Strategies that have led to higher payoffs in the past are played with a higher probability, i.e., these are *better replies* rather than *best replies* to the available information. The informational requirements are fairly limited too. Players do not need to know the structure of the underlying game, or even the number of other players, and their strategies.

We consider a wide variety of models of reinforcement learning. Some perform better than others in empirical tests. The main result from the empirical literature is that while reinforcement learning is good at predicting the direction of change, the predicted speed of learning is too sluggish. If the speed of adjustment is increased, then the model fails in other respects. Factors that are not predicted to affect individual choices under reinforcement learning, such as information on the payoffs of others, are often found to be important in experiments. Reinforcement learning is probably most effective when players have a limited understanding of the underlying structure of the game, and they have relatively low information.

A second class of models, models of *belief-based learning*, impose relatively greater rationality. In these models, players form beliefs about others and about the environment around them, and how these beliefs evolve over time. In an important subclass of this model, *fictitious play*, players assign some initial weights to the strategies of opponents, and update the weights as they temporally observe the strategies of the opponents. The relative weights are used to compute the expected probabilities of play of the opponent. The player then plays a best response to these expectations. A special case is *Cournot dynamics*. For 2×2 games, the limiting distribution of play in these models corresponds to a Nash equilibrium. But these results do not hold for 3×3 games.

Existing data does not produce a clear winner in the choice between belief-based models and reinforcement learning models. A range of other variables, not predicted by these theories, such as the optimization premium, and information on market conditions, or about other players, turn out to be important in experiments. Further, a variety of other learning mechanisms, such as imitate the best, learning direction theory, and trial and error learning are also empirically

important. Thus, fictitious play, by itself, is too simple to explain all relevant aspects of human learning.

It turns out that reinforcement learning and fictitious play are special cases of an overarching model, the *experience weighted attraction* (EWA) model. The EWA allows for the incorporation of regret from forgone payoffs, and the depreciation of past experience on account of limited memory, or forgetfulness. EWA also nests a range of other known models such as the Cournot model, *cumulative reinforcement model*, *weighted fictitious play*, and the *average reinforcement model*. The performance of the EWA model is superior to reinforcement learning, and fictitious play, even after taking account of its larger number of parameters.

A more parsimonious alternative to the EWA model is the one-parameter *functional* EWA model or fEWA or *EWA-lite*. The parameters in this model are *self-tuning* in the sense that they respond to changes in the environment. In one comprehensive comparison of the alternative models, based on seven different games, EWA and fEWA do better than reinforcement learning, weighted fictitious play, and quantal response equilibrium.

EWA and fEWA are models of *adaptive play* in the sense that players specify a rule that transforms the history of play into current actions. However, sophisticated players may wish to take account of the information about other players in a way that allows them to outguess others, and improve their payoffs. Such players may also wish to teach other players, and shepherd them in the direction of better payoffs. These are models of *anticipatory learning* and they are supported by the empirical evidence.

Learning direction theory aims to qualitatively predict the direction of change, rather than specify the probabilities with which players play their strategies. The basic theory predicts well the high degree of observed cooperation in repeated prisoner's dilemma game experiments and the direction of bid changes in auction experiments. The theory has been extended to allow for quantitative predictions of the probabilities of play. We consider a range of these extensions, including *impulse matching learning*, *impulse balance equilibrium*, *action-sample learning*, and *payoff-sampling learning*. The emerging evidence makes a strong case for taking these theories seriously, and in many cases, they outperform other standard theories of learning. As our final model in this class, we consider the explanation of the experimental results from *order statistic games*.

Individuals may have to learn in situations in which nature probabilistically draws a game from a set of games, and the *complexity of the learning environment* varies. In order to cope with complexity, players may *categorize games* to reduce the cognitive requirements. For each category, players may apply to each game the average beliefs, or average payoffs for the category, and then play a best response. Greater complexity (larger set of games and lower information) reduces the probability of playing a Nash equilibrium. One also observes learning spillovers across games. Transparency of information on payoffs rather than knowledge of the past play of the opponent is more conducive to the Nash equilibrium being played.

The conditions required for the limiting frequencies of play to be a Nash equilibrium in most plausible learning models are strong. In the presence of enough ex-ante uncertainty about the payoffs of others, there is no guarantee that long-run play is a Nash equilibrium. In more realistic and cognitively plausible models that involve the idea of *regret matching*, the limiting play approaches the set of *coarse correlated equilibria*. This set of equilibria is typically much larger than the set of Nash equilibria.

Chapters 14 and 15 explore various aspects of learning by using deterministic dynamics. In these dynamics, the main object of interest, an *asymptotically stable steady state*, has two properties. First, the outcome is *history dependent*. Second, once the system reaches the basin of attraction of any steady state, it stays there forever. In the final chapter in Part 5, Chapter 16, we consider

stochastic dynamics by introducing a source of *persistent randomness*. The outcome in any period is uncertain; however, we may be able to predict where the stochastic dynamic system is likely to spend most of its time in the long run. The resulting object of interest, a *stochastic steady state*, is able to relax both the steady state properties of deterministic dynamics, namely, history dependence, and the locked-in effect in steady states.

A stochastic steady state may be *ergodic*, i.e., there is a high likelihood that the system spends most of the time in that steady state, irrespective of the history of the system. This gives rise to a powerful argument for choosing among steady states of a deterministic system. Furthermore, we may have *punctuated equilibria*, i.e., once a steady state is achieved, small mutations may accumulate, and may eventually push the system into the basin of attraction of another steady state. The point at which this occurs is called the *tipping point*. A stochastically stable steady state is one in which the basin of attraction is hardest to escape from.

The underlying environment in models of stochastic dynamics has relatively low rationality. Our main vehicle of analysis will be *perturbed Markov chains* that essentially specify adaptive play with memory length T , sample size τ , and error rate ε . Players play a best response to a sample of past observations with a high probability, and with some small probability, ε , they make a mistake, and play randomly. This framework can be used to explain the evolution of institutions and norms as the result of gradual, small, adjustments arising from adaptively learning individuals. The low-rationality adaptive-play environment of *perturbed dynamics* often generates a Nash equilibrium without requiring the stringent epistemic conditions for a Nash equilibrium.

Stochastically stable states may allow for the play of strictly dominated strategies when players use *imitate the best*, rather than play best replies to past observations. For 2×2 games, the stochastically stable state is also the risk dominant state. However, many of the results from 2×2 coordination games do not generalize to 3×3 , or higher dimensional games. A strict risk dominant equilibrium for 3×3 games may not exist, and may not coincide with stochastic stability. However, several of the results generalize for a class of games called *acyclic games*. Stochastic dynamics allow us to pick among alternative norms. The chosen norm may be inefficient, but relatively fair in its distributive implications. This has the potential to explain adherence to moral standards of behavior.

Finally, we briefly consider games played on graphs with boundedly rational players whose direct interaction is local, but indirect interaction is global. We characterize the features of the stochastically stable states, and apply the theory to agrarian contracts.

The applications of stochastic dynamics are interesting and important in their own right. However, it is not clear how one might use this framework to explain the results of experiments in a variety of commonly used games in behavioral economics. For a better evaluation of this line of research, we need to pit these models in a horse race against other models. In contrast to the material that we cover in Chapter 15, insufficient progress has been made in this direction.

A range of Appendices may be useful to interdisciplinary audiences. The material here simply states the relevant results, and is not a substitute for consulting the original sources, and gaining a deeper understanding of the techniques. Appendix A outlines the basics of ordinary differential equations. We consider, briefly, the existence, uniqueness, and stability conditions for linear, and non-linear ordinary differential equations, and systems of such equations. Appendix B considers a model of stochastic utility, and shows that when the preference shocks have an extreme value distribution, the resulting choice of strategies by the players follows a logit model. This material also underpins some of the models considered in Part 4, such as the quantal response equilibrium. Appendix C gives a very condensed guide to the relevant results in unperturbed and perturbed Markov chains.

CHAPTER 14

Evolutionary Game Theory

14.1 Introduction

Evolutionary game theory is based on Darwin's theory of natural selection. The basic idea is that each animal is represented by a *genotype*, or its exact genetic structure. We are using the word genotype in the sense of *alleles*, i.e., any number of alternative forms of the same gene, so there is *heterogeneity* among the genotypes. Each genotype induces a particular behavior that is its phenotype. Hence, animals are *hardwired* to follow certain strategies. By contrast, in classical game theory, strategies, at least as applied to humans, are the product of conscious deliberation. In this chapter, each genotype maps into a unique phenotype, but more generally, a particular genotype may induce a probability distribution of phenotypes.¹

In any period of play, distinct phenotypes have different levels of *fitness*. Fitness is taken to mean the *reproductive success* of the animal. Nature undertakes the following process of *natural selection*. In any period, phenotypes that had higher fitness in the previous period *proliferate* relatively more in numbers. At the same time, phenotypes with lower fitness in the previous period *decay*, or proliferate at a lower rate. The relative fitness of a phenotype depends on (i) its suitability for the current environment, and (ii) the fitness of other phenotypes in the population. New phenotypes arise randomly through a process of *mutation* of the underlying genotype. This is the sense in which we may conceive of learning in evolutionary game theory.

A central question in evolutionary game theory is whether the existing phenotypes, also known as *resident phenotypes*, can survive invasion by other mutant phenotypes. If they can, then they are said to be *evolutionary stable strategies* (ESS). If in an ESS natural selection picks a unique surviving genotype, then the resulting population is said to be *monomorphic*. If there are multiple genotypes in an ESS, then the resulting population is said to be *polymorphic*.

Evolution is an extremely slow and conservative process, occasionally punctuated by the random event of mutation. We also briefly consider the literature on *group selection*, an idea that was initially proposed by Wynne-Edwards (1962). In this view, an individual's reproductive success may be curtailed if this benefits the overall fitness of the group; this idea has not been without controversy.

The speed of learning under evolutionary game theory is slower relative to other learning contexts, such as human learning, in experimental settings (Camerer, 2003). Does this mean that evolutionary game theory is of limited relevance for behavioral economics? This is not the case,

¹ For instance, genes determine that humans have fingerprints. However, the exact fingerprint is unique to each human.

especially if one is interested in exploring the source of the primitives in most economic problems, i.e., preferences and beliefs. For instance, human emotions, visceral influences, the presence of fairness and reciprocity concerns, the nature of time discounting, all have evolutionary explanations.

Section 14.2 reviews some preliminaries, and describes the basic notation.

In Section 14.3, we derive and define an ESS from first principles (Hammerstein and Selten, 1994). We apply the ESS concept to symmetric two-player games in which the players may belong to identical, or different populations. We then derive the *primary* and *secondary* criteria for checking that a strategy is an ESS. These criteria are sometimes also known as *equilibrium* and *stability* conditions, respectively, and considerably simplify the problem of finding an ESS.

In Section 14.4, we establish the relation between ESS and other equilibrium concepts in game theory. For instance, we show that every ESS is a symmetric Nash equilibrium. But the converse does not hold unless the symmetric Nash equilibrium is strict, or we have a symmetric Nash equilibrium in completely mixed strategies (i.e., all pure strategies are played with a strictly positive probability). Hence, ESS is a refinement of Nash equilibrium. The ESS is also a perfect equilibrium, so, in an ESS, no weakly dominated strategies are played.

The concept of an ESS is agnostic about the underlying dynamic process. An important advance, the *replicator dynamics* (Taylor and Jonker, 1978), provides the necessary foundations for an ESS. It is considered in Section 14.5. We describe the rationale for replicator dynamics, and describe its various properties such as proportionality, quotient dynamics, and additive invariance of payoffs.² We also show how the replicator dynamics can be extended to the *asymmetric populations case*. We derive two main properties of the replicator dynamics. First, any asymptotically stable state of the replicator dynamics is a Nash equilibrium. Second, any ESS is an asymptotically stable state of the replicator dynamics.

Section 14.6 considers some applications of evolutionary game theory. These applications make the following main points. First, in the context of a prisoner's dilemma game, we show that ESS can enable us to choose among multiple *monomorphic* steady states of the replicator dynamics. The chosen steady state in ESS typically depends on the initial conditions. Second, while strategies are hardwired in evolutionary game theory, these strategies could be quite complex history-dependent strategies. We examine the evolutionary stability of the *tit-for-tat strategy* for the twice repeated prisoner's dilemma game. Third, in the context of a 2×2 game of chicken, we show that the evolutionary equilibrium could be *polymorphic*, in which both strategies coexist. Fourth, we show that, in the rock-scissor-papers game, there may not be an asymptotically stable steady state even when we allow players to play totally mixed strategies. Fifth, we demonstrate an application of replicator dynamics, and study asymptotic stability using a model of competing and *heterogenous multiple populations*.

In many cases, it is difficult to isolate pairwise interactions among animals for a prize. Rather, it appears that players are *playing the field*, rather than being in direct competition with a particular opponent. The relevant notion of a game in this case is a *population game*. Section 14.7 considers these issues. We define a population game, and define the appropriate notion of an ESS in such a game.

Section 14.8 considers *gene-culture coevolution* (Gintis, 2009b, 2011; Gintis and Helbing, 2014). Environmental changes that are neither too slow, nor too rapid, may be transmitted from one generation to another through *epigenetic transmission*. Epigenetic changes do not alter the structure

² For formal derivations of the replicator dynamics, see Hofbauer and Sigmund (1988) and Gintis (2009a).

of DNA, but they are physiological changes. Examples include beaver dams, and social structures in bee hives; a phenomenon known as *niche construction* (Dawkins, 2006).

In humans, an important form of epigenetic change occurs through the *transmission of culture*; culture may be directly encoded into the brain as distinct from being encoded into the DNA (Gintis, 2011). The physiology of communication in humans is a good example of gene–culture coevolution (Gintis, 2011). As the advantages of communications (a cultural change) became reflected in the greater fitness of the communicators, there were genetic changes that lowered the voice box in the larynx, hence, facilitating even better communication. Other examples that we consider include *evolutionary political economy* (Gintis et al., 2014), the *Westermarck effect* (Wilson, 2012), and *evolutionary prosociality* (Gintis, 2011). We also examine a formal model of cultural transmission in the form of *intergenerational transmission of altruism* despite its lower biological fitness relative to selfishness (Gintis and Helbing, 2014).

Section 14.9 considers the *evolutionary stability of other-regarding preferences* that play a central role in Part 2 of the book. Section 14.9.1 considers models of *kin selection* and *inclusive fitness* (Hamilton, 1964; Maynard Smith, 1964). The idea is that humans may take altruistic actions that are personally costly, say, in term of their own fitness levels, provided the fitness of other kin members, who carry similar genes, can be enhanced. This results in an enduring mathematical inequality, termed as *Hamilton's rule*, which specifies the conditions under which individuals will be prosocial. We work out the implications for evolutionary stability in a prisoner's dilemma game (Nowak, 2006). One drawback of this approach is that it does not explain altruistic human behavior towards genealogically unrelated other humans. However, Hamilton's inequality survives, in a suitably reinterpreted manner, in most of the alternative explanations that we consider.³

Section 14.9.2 considers the conflict between *individual and group selection* (Hamilton, 1964; Price, 1972); these models are also known as models of *multilevel selection*. Consider several competing groups. Within each group, altruism has lower fitness relative to selfishness. However, in competition across groups for scarce resources, groups that have more prosocial traits are likely to be more successful relative to groups of predominantly self-regarding individuals. If this group-selection effect can dominate the first effect (relatively greater fitness of selfishness over altruism in groups), then prosocial preferences may be evolutionarily stable. The benefits of altruism can be enhanced by *assortative matching* of individuals that makes it more likely that altruistic individuals are more likely to meet each other.

Section 14.9.3 considers the notion of *reciprocal altruism* (Trivers, 1971). The idea is that altruism may evolve because individuals reciprocate the previously altruistic choices of others. It was shown that in computer tournaments using the prisoner's dilemma game, the *tit-for-tat strategy* (begin by cooperating, and cooperate in period t if the opponent cooperated in period $t-1$, otherwise defect forever) was the winning strategy relative to pure defection (Axelrod, 1984). However, reciprocal altruism is hard to find in animals and mostly found in humans, and in close relatives, such as chimps. Further the evolutionary stability of strategies such as tit-for-tat is quite sensitive to behavioral errors by players and the size of the group. It also does not explain altruism towards strangers in one-shot games. Empirical evidence for these models, based on the historical record, is not supportive either (Bowles and Gintis, 2011).

³ For an important re-examination of Hamilton's inequality, see Gintis (2014).

Section 14.9.4 considers *altruism* and *spite* towards third parties, based on the indirect transmission of information through, say, gossip and language; this is known as *indirect reciprocity* (Alexander, 1987; Nowak and Sigmund, 1998).

Section 14.9.5 summarizes the objections to using standard repeated game arguments in classical game theory to justify observed prosocial behavior (Bowles and Gintis, 2011; Gintis, 2009b). Indeed, explanations for prosocial behavior based on the folk theorem are so problematic, despite this being offered as the main explanation for cooperation in classical game theory, that Bowles and Gintis (2011) term them as *evolutionary irrelevant equilibria*. Problems include the lack of a specification of the dynamic process towards equilibrium, lack of predictions for the realistic case in which, with a high probability, players receive private signals, lack of foundations for the play of mixed strategies under private signals, and the inability of the purification agenda to apply to many important classes of games including repeated games.

Section 14.9.6 outlines the *indirect evolutionary approach* (Güth and Yaari, 1992; Güth, 1995). In this approach, the preferences of individuals are selected by evolutionary pressures. However, as in neoclassical economics, individuals maximize their utility at each instant in time. When individuals are matched to play a game, they are assumed to recognize the preferences of others, and then play a best response. However, one criticism of this approach is that the recognition of preferences of others, and the associated cues, ought themselves to be endogenous to the evolutionary process (Robson and Samuelson, 2011).

Within the indirect evolutionary approach, we examine the evolutionary stability of the textbook self-regarding preferences (Alger and Weibull, 2013). It turns out that selfishness is wiped out by a form of preferences that they call *homo-moralis*. Indeed, a subset of these preferences, *homo-hamiltonensis*, whose degree of morality is linked to the *degree of assortativity* of matchings, turns out to be the unique evolutionary stable form of preferences. This work illustrates the power and usefulness of the evolutionary approach. We already know from the empirical evidence in Part 2 of the book that prosociality is critical in humans. A demonstration of the evolutionary instability of selfishness is not only consistent with the empirical evidence, it more decisively debunks the neoclassical preference to give center-stage to sociopaths (individuals who have exclusively self-regarding preferences).

In Section 14.9.7, we briefly comment on some assorted issues in reciprocity.

Section 14.10 asks if *models of social learning*, in which individuals are not hardwired to play particular strategies, can produce dynamic processes similar to replicator dynamics. In Section 14.10.1, we consider a two-population model of *social imitation* in which players engage in bilateral meetings. Each period, some players get an opportunity to revise their strategy, and they select the strategy of a randomly chosen player from their population. The dynamics arising from such a process of imitation are also known as *emulation dynamics*. One feature of such models is that they might disproportionately select the median payoff in the population (median enhancing emulation dynamics). In Section 14.10.2, we derive the conditions under which the emulation dynamics are identical to the replicator dynamics. However, the replicator dynamics are not robust to changes in the form of social imitation.

Yet another formulation of the basic Darwinian idea is based on the concept of *payoff monotonicity* that we consider in Section 14.10.3. This requires that strategies that have a higher payoff also have a higher growth rate. Payoff monotone dynamics turn out to have many of the same properties as replicator dynamics. In addition, these dynamics ensure that dominated pure strategies are never played.

14.2 Some preliminaries

In this section, we introduce the notation and review some of the preliminaries for subsequent analysis.⁴ Consider a population of a large number of identical players. In every period, two randomly drawn players from the population play a symmetric, two-player, normal form game.

The set of players is $N = \{1, 2\}$. The pure strategy set of any player is $S = \{s_1, s_2, \dots, s_n\}$. A *pure strategy profile* in any strategic interaction between the two players is $s = (s_i, s_j) \in S \times S$, where s_i is the pure strategy chosen by the first player, and s_j is the pure strategy chosen by the second player.⁵

Player i 's *mixed strategy*, σ_i , is a probability distribution over his pure strategy set, S . We denote by $\sigma_i(s_k) \in [0, 1]$, the probability with which player i plays the pure strategy $s_k \in S$. Sometimes we may need to use the slightly more cumbersome notation $\sigma_i(s_{ik})$ for the probability with which player i plays his k th pure strategy $s_{ik} \in S$. Any pure strategy that is played with a non-zero probability under σ_i is said to belong to its *support*. In particular, a pure strategy $s_k \in S$ for player i , simply corresponds to a degenerate mixed strategy, $\sigma_i(s_k) = 1$. The set of all mixed strategies of player i is denoted by Σ_i . Formally,

$$\Sigma_i = \left\{ \sigma_i = (\sigma_i(s_1), \sigma_i(s_2), \dots, \sigma_i(s_n)) : \sigma_i(s_j) \in [0, 1] \ \forall j \text{ and } \sum_{j=1}^n \sigma_i(s_j) = 1 \right\}. \quad (14.1)$$

A *completely mixed strategy*, σ_i , gives a strictly positive probability weight to all pure strategies, i.e., $\sigma_i(s_j) > 0$ for all $s_j \in S$. The mixed strategy profile of the game $\sigma = (\sigma_i, \sigma_j)$ is an element of the set $\Sigma_i \times \Sigma_j$, $i \neq j$. It will sometimes be convenient to express the mixed strategy profile by $\sigma = (\sigma_i, \sigma_{-i})$, where σ_{-i} is the mixed strategy profile of the player other than player i . The payoff of any player, $i = 1, 2$, is given by the mapping $\pi_i : \Sigma_i \times \Sigma_j \rightarrow \mathbb{R}_+$, $i \neq j$.

Recalling that pure strategies are simply degenerate mixed strategies, the restriction of *symmetric* payoffs is that $\pi_i(s_i, s_j) = \pi_j(s_j, s_i)$ (so player identities are unimportant).⁶ A two-player *normal form*, or *strategic form*, game is denoted by $G(2, S, \pi)$, where $\pi = (\pi_i, \pi_j)$, and each of π_i, π_j is a symmetric payoff function.

It is sometimes convenient to denote the $n \times n$ matrix of payoffs of player i by Π_i . A typical element of this matrix is denoted by $a_{ij} \equiv \pi_i(s_i, s_j)$. The k th row of the payoff matrix Π_i , denoted by Π_{ik} , gives the payoffs to player i when he follows a fixed pure strategy $s_k \in S$, and the opponent successively follows his pure strategies, $s_1, \dots, s_n$.

$$\Pi_{ik} = (a_{k1}, a_{k2}, \dots, a_{kn}). \quad (14.2)$$

⁴ For textbook treatments of evolutionary game theory, see Maynard Smith (1982), Weibull (1995), Vega-Redondo (1996), Samuelson (1997), Hofbauer and Sigmund (1998), Gintis (2009b), and Sandholm (2010).

⁵ While our focus is on random meetings between players, in actual practice one may meet disproportionately many players who play a similar strategy. For instance, altruists might have a network in which they are more likely to meet other altruists. These considerations can be dealt with by using evolutionary graph theory (Lieberman et al., 2005).

⁶ Notice, for instance, that despite the evidence on gender differences that we review in several parts of the book, we do not draw any distinction between strategies played by males and females. Indeed, the dynamic processes in this chapter, such as replicator dynamics, will be typically based on asexual reproduction.

	C	D
C	2, 2	0, 3
D	3, 0	1, 1

Figure 14.1 The two-player prisoner's dilemma game.

We shall assume that a mixed strategy of player i , σ_i , is a *column vector*, and its transpose is denoted by σ_i^T . Then we can write the payoff of player i compactly as

$$\pi_i(\sigma_i, \sigma_j) = \sigma_i^T \Pi_i \sigma_j, i \neq j. \quad (14.3)$$

Example 14.1 Consider the two-player, symmetric, prisoner's dilemma game shown in Figure 14.1. In this case, $S = \{C, D\}$. Denote the row player by 1, and the column player by 2. The payoff matrix of player 1 is given by

$$\Pi_1 = \begin{bmatrix} 2 & 0 \\ 3 & 1 \end{bmatrix}.$$

The payoff matrices are symmetric, so, for instance, $\pi_1(C, D) = \pi_2(D, C) = 0$. Let the corresponding mixed strategies of the players be given by the column vectors

$$\sigma_1 = \begin{pmatrix} p \\ 1-p \end{pmatrix}, \sigma_2 = \begin{pmatrix} q \\ 1-q \end{pmatrix},$$

where p and q are, respectively, the probabilities with which player 1 and player 2 play C. The payoff of player 1 from the mixed strategy profile (σ_1, σ_2) can be written as

$$\pi_1(\sigma_1, \sigma_2) = 2pq + 0 \times p(1-q) + 3(1-p)q + (1-p)(1-q).$$

which can be written in more compact matrix notation as $\pi_1(\sigma_1, \sigma_2) = \sigma_1^T \Pi_1 \sigma_2$.

The next lemma contains two useful results for later application.

Lemma 14.1 Consider the game $G(2, S, \pi)$. Suppose that σ_j is a linear combination of two mixed strategies, σ_j^a and σ_j^b , such that $\sigma_j = \alpha \sigma_j^a + (1-\alpha) \sigma_j^b$, $\alpha \in [0, 1]$.

(a) Then, (14.3) can be written as

$$\pi_i(\sigma_i, \sigma_j) = \alpha \pi_i(\sigma_i, \sigma_j^a) + (1-\alpha) \pi_i(\sigma_i, \sigma_j^b).$$

(b) Suppose that $\sigma_i = \sigma_j$. In this case

$$\pi_i(\sigma_j, \sigma_j) = \alpha \pi_i(\sigma_j^a, \sigma_j) + (1-\alpha) \pi_i(\sigma_j^b, \sigma_j).$$

Proof: The proof follows from a simple application of the distributive laws of matrices. From (14.3)

$$\begin{aligned} \pi_i(\sigma_i, \sigma_j) &= \sigma_i^T \Pi_i (\alpha \sigma_j^a + (1-\alpha) \sigma_j^b) \\ &= \alpha \sigma_i^T \Pi_i \sigma_j^a + (1-\alpha) \sigma_i^T \Pi_i \sigma_j^b \\ &= \alpha \pi_i(\sigma_i, \sigma_j^a) + (1-\alpha) \pi_i(\sigma_i, \sigma_j^b) \end{aligned}$$

which proves part (a). To prove (b) note that

$$\begin{aligned}\pi_i(\sigma_j, \sigma_j) &= \left(\alpha \sigma_j^a + (1 - \alpha) \sigma_j^b \right)^T \Pi_i \sigma_j \\ &= \alpha \left(\sigma_j^a \right)^T \Pi_i \sigma_j + (1 - \alpha) \left(\sigma_j^b \right)^T \Pi_i \sigma_j \\ &= \alpha \pi_i(\sigma_j^a, \sigma_j) + (1 - \alpha) \pi_i(\sigma_j^b, \sigma_j).\end{aligned}$$

This proves part (b). ■

14.3 Evolutionary stable strategies in two-player games

An *evolutionary stable strategy* (ESS) is the standard equilibrium concept in evolutionary game theory. It was introduced in the seminal paper by Maynard Smith and Price (1973). A strategy, or a *phenotype*, is an ESS if starting from a state where it is played by all players in the population, it cannot be invaded by a small fraction of mutants who play an alternative strategy. We apply the definition of ESS only to symmetric two-player normal form games $G(2, S, \pi)$ (see Section 14.2). For an extension to extensive form games, see Section 7 in Hammerstein and Selten (1994).

Suppose that there is an established mixed strategy in the population, $\sigma \in \Sigma$; this could be a pure strategy if σ is degenerate. Now consider the random event of *mutation*. An exogenous fraction, $\lambda > 0$, of the population mutates, and plays a mixed strategy, $\hat{\sigma} \neq \sigma$. The remaining proportion of the population, $1 - \lambda$, continues to play the established strategy σ . Define the average strategy in the population, $\bar{\sigma}$, by

$$\bar{\sigma} = \lambda \hat{\sigma} + (1 - \lambda) \sigma. \quad (14.4)$$

The mixing operation in (14.4) makes use of the usual compounding rule of probabilities.⁷ When random, pairwise, meetings between the players take place, then, on average, one meets an opponent who plays the strategy $\bar{\sigma}$. We would like to see how each of the strategies, $\sigma, \hat{\sigma}$, fare against $\bar{\sigma}$.

Definition 14.1 (*Fitness*): The fitness of the established strategy, σ , and the mutant strategy, $\hat{\sigma}$, is defined as its payoff against the average strategy $\bar{\sigma}$. Hence, the respective fitness levels of the strategies, $\sigma, \hat{\sigma}$, are given by $\pi(\sigma, \bar{\sigma})$ and $\pi(\hat{\sigma}, \bar{\sigma})$. The fitness of the average strategy is given by $\pi(\bar{\sigma}, \bar{\sigma})$.

Measuring biological payoffs in the field is not an easy task. As Hammerstein and Selten (1994, p. 975) write: “Biological critics of evolutionary game theory have argued that it is almost impossible to get empirical access to game-theoretic payoffs [see the open peer commentaries of a survey paper by Maynard Smith (1984)].” See also the discussion following this quote for some attempts to measure biological payoffs and the difficulties in doing so. We will assume away these important problems in our discussion below.

ESS is a static concept, but a better understanding of it can be achieved by considering the time evolution of the mutant strategy. Let time $t = 0, 1, 2, \dots$ be discrete, so that at time t , the

⁷ For instance, suppose that there are only two pure strategies (s_1, s_2) . Let the mixed strategies be $\sigma = (0.6, 0.4)$ and $\hat{\sigma} = (0.2, 0.8)$. If $\lambda = 0.1$, then the average strategy is $\bar{\sigma} = (0.56, 0.44)$.

respective fractions of the mutant and non-mutant populations are λ_t and $1 - \lambda_t$. There are several methods of motivating an ESS; we follow the method in Hammerstein and Selten (1994).

Definition 14.2 (*Growth factors*): The growth factors of the mutant and non-mutant populations are given, respectively, by $\frac{\lambda_{t+1}}{\lambda_t}$ and $\frac{1-\lambda_{t+1}}{1-\lambda_t}$.

Consistent with the Darwinian view of evolution, we assume that the growth rates of phenotypes depend on their respective fitness levels, $\pi(\hat{\sigma}, \bar{\sigma})$, $\pi(\sigma, \bar{\sigma})$, relative to the fitness of the average strategy, $\pi(\bar{\sigma}, \bar{\sigma})$. Hence,

$$\frac{\lambda_{t+1}}{\lambda_t} = \frac{\pi(\hat{\sigma}, \bar{\sigma})}{\pi(\bar{\sigma}, \bar{\sigma})} \text{ and } \frac{1-\lambda_{t+1}}{1-\lambda_t} = \frac{\pi(\sigma, \bar{\sigma})}{\pi(\bar{\sigma}, \bar{\sigma})}. \quad (14.5)$$

Thus, the relative growth rate of the two phenotypes is given by

$$\frac{\left(\frac{\lambda_{t+1}}{\lambda_t}\right)}{\left(\frac{1-\lambda_{t+1}}{1-\lambda_t}\right)} = \frac{\pi(\hat{\sigma}, \bar{\sigma})}{\pi(\sigma, \bar{\sigma})}. \quad (14.6)$$

The growth rate of mutants is greater if they do relatively better against the average strategy. To discover a potential ESS, we are interested in conditions under which the proportion of mutants disappears over time, i.e., $\lambda_t \rightarrow 0$ as $t \rightarrow \infty$. This is the sense in which the existing population strategy cannot be successfully invaded by mutants. We have

$$\lambda_t - \lambda_{t+1} = \left(1 - \frac{\lambda_{t+1}}{\lambda_t}\right) \lambda_t. \quad (14.7)$$

Substituting (14.5) in (14.7), we get

$$\lambda_t - \lambda_{t+1} = \frac{\pi(\bar{\sigma}, \bar{\sigma}) - \pi(\hat{\sigma}, \bar{\sigma})}{\pi(\bar{\sigma}, \bar{\sigma})} \lambda_t. \quad (14.8)$$

Using (14.4) and Lemma 14.1, we get

$$\pi(\bar{\sigma}, \bar{\sigma}) = \pi(\lambda_t \hat{\sigma} + (1 - \lambda_t) \sigma, \bar{\sigma}) = \lambda_t \pi(\hat{\sigma}, \bar{\sigma}) + (1 - \lambda_t) \pi(\sigma, \bar{\sigma}), \quad (14.9)$$

or,

$$\pi(\bar{\sigma}, \bar{\sigma}) - \pi(\hat{\sigma}, \bar{\sigma}) = (1 - \lambda_t) [\pi(\sigma, \bar{\sigma}) - \pi(\hat{\sigma}, \bar{\sigma})]. \quad (14.10)$$

Substituting (14.10) in (14.8) we get

$$\lambda_t - \lambda_{t+1} = \frac{(1 - \lambda_t) [\pi(\sigma, \bar{\sigma}) - \pi(\hat{\sigma}, \bar{\sigma})]}{\pi(\bar{\sigma}, \bar{\sigma})} \lambda_t \quad (14.11)$$

Since payoffs are assumed to be positive (see Section 14.2), $\pi(\bar{\sigma}, \bar{\sigma}) \geq 0$. Phenotype fractions are non-negative, so $\lambda_t \geq 0$. Thus, from (14.11), we get that

$$\lambda_t \geq \lambda_{t+1} \Leftrightarrow \pi(\sigma, \bar{\sigma}) \geq \pi(\hat{\sigma}, \bar{\sigma}). \quad (14.12)$$

Hence, the mutant phenotype grows over time ($\lambda_t < \lambda_{t+1}$) if $\pi(\sigma, \bar{\sigma}) < \pi(\hat{\sigma}, \bar{\sigma})$, i.e., the mutant does relatively better against the average strategy as compared to the dominant phenotype. Conversely, for the mutant to disappear over time, a necessary condition is that $\lambda_t > \lambda_{t+1}$ for all t , i.e., in each period, the dominant phenotype outperforms the mutant against the average strategy. This can be summarized in the following definition.

Definition 14.3 (*Evolutionary stable strategy-I; Taylor and Jonker, 1978*): Consider the symmetric two-player game $G(2, S, \pi)$. Then, a mixed strategy $\sigma \in \Sigma$ is said to be an ESS, if for all mutant strategies $\hat{\sigma} \in \Sigma$, $\hat{\sigma} \neq \sigma$, there exists some “uniform invasion barrier” $0 < \bar{\lambda} < 1$ such that if the mutant population is small enough in the sense that $0 < \lambda < \bar{\lambda}$, then

$$\pi(\sigma, \bar{\sigma}) > \pi(\hat{\sigma}, \bar{\sigma}),$$

where $\bar{\sigma} = \lambda\hat{\sigma} + (1 - \lambda)\sigma$.⁸

A useful way of restating the definition of an ESS arises by eliminating $\bar{\sigma}$ from (14.12). Using Lemma 14.1, and (14.4), we get:

$$\pi(\sigma, \bar{\sigma}) = \lambda_t \pi(\sigma, \hat{\sigma}) + (1 - \lambda_t) \pi(\sigma, \sigma). \quad (14.13)$$

$$\pi(\hat{\sigma}, \bar{\sigma}) = \lambda_t \pi(\hat{\sigma}, \hat{\sigma}) + (1 - \lambda_t) \pi(\hat{\sigma}, \sigma). \quad (14.14)$$

From (14.13), (14.14) we get

$$\pi(\sigma, \bar{\sigma}) - \pi(\hat{\sigma}, \bar{\sigma}) = \lambda_t [\pi(\sigma, \hat{\sigma}) - \pi(\hat{\sigma}, \hat{\sigma})] + (1 - \lambda_t) [\pi(\sigma, \sigma) - \pi(\hat{\sigma}, \sigma)]. \quad (14.15)$$

Substituting (14.15) in (14.12) we get:

$$\lambda_t \geq \lambda_{t+1} \Leftrightarrow \lambda_t [\pi(\sigma, \hat{\sigma}) - \pi(\hat{\sigma}, \hat{\sigma})] + (1 - \lambda_t) [\pi(\sigma, \sigma) - \pi(\hat{\sigma}, \sigma)] \geq 0. \quad (14.16)$$

This gives rise to the following useful observations.

1. Suppose that the fraction of mutants is extremely small, so $\lambda_t \rightarrow 0$. In this case, (14.16) implies that

$$\lambda_t \geq \lambda_{t+1} \Leftrightarrow \pi(\sigma, \sigma) \geq \pi(\hat{\sigma}, \sigma). \quad (14.17)$$

This is called the *primary criterion* for determining an ESS. In particular, the existing strategy σ is an ESS if it does at least as good against itself, relative to how well a mutant strategy does against it.

2. Suppose now that $\pi(\sigma, \sigma) = \pi(\hat{\sigma}, \sigma)$, so the dominant strategy and the mutant strategy are both best replies to the dominant strategy. In this case, (14.16) implies that

$$\lambda_t \geq \lambda_{t+1} \Leftrightarrow \pi(\sigma, \hat{\sigma}) \geq \pi(\hat{\sigma}, \hat{\sigma}). \quad (14.18)$$

⁸ The result on the uniform invasion barrier only holds for two-player games (Vickers and Cannings, 1987).

This is called the *secondary criterion* for determining an ESS. In particular, if $\pi(\sigma, \sigma) = \pi(\hat{\sigma}, \sigma)$, a necessary condition for the mutant to eventually vanish is that in a population that comprises predominantly mutants, strategy σ does better against the mutants relative to how well the mutants do against themselves.

We summarize these observations for future reference.

Definition 14.4 (*Evolutionary stable strategy-II; Maynard Smith and Price, 1973*): Consider a symmetric two-player game $G(2, S, \pi)$. A mixed strategy σ is said to be an ESS, if for any mutant strategy, $\sigma \neq \hat{\sigma}$, any of the following two conditions are satisfied.

- (a) *Equilibrium condition or primary criterion*: $\pi(\sigma, \sigma) > \pi(\hat{\sigma}, \sigma)$.
- (b) *Stability condition or secondary criterion*: If $\pi(\sigma, \sigma) = \pi(\hat{\sigma}, \sigma)$, then $\pi(\sigma, \hat{\sigma}) > \pi(\hat{\sigma}, \hat{\sigma})$.

From Definition 14.4, the primary criterion reveals that an ESS is a Nash equilibrium. However, the secondary criterion proposes a strengthening of Nash equilibrium. Hence, the concept of an ESS is more stringent as compared to a Nash equilibrium.

In 2×2 games, an ESS is guaranteed under fairly mild conditions as we see next, but this result does not generalize to higher dimensional games.

Proposition 14.1 (*van Damme, 1991*): Consider a symmetric two-player game $G(2, S, \pi)$ in which each player has two pure strategies, so $S = (s_1, s_2)$. Suppose that $\pi(s_1, s_1) \neq \pi(s_2, s_1)$ and $\pi(s_2, s_2) \neq \pi(s_1, s_2)$. Then G is guaranteed to have an ESS.

The ESS is a static equilibrium concept and it does not specify the path to the equilibrium. One starts with a candidate ESS strategy and checks that it cannot be successfully invaded by a mutant strategy. Although the ESS is not tied to any dynamic process, the dynamic underpinnings of this concept are considered in Section 14.5 below when we discuss replicator dynamics.

The ESS takes mutations as a rare event. Hence, the evolutionary stability of the existing strategy is checked against invasion by one mutant strategy at a time. It is assumed that once a mutation takes place, the system settles down to a stable state before the next mutation occurs. There is no presumption that an ESS would survive against several simultaneous mutations. In actual practice, however, one may expect mutations to take place on a continual basis, particularly when the mutations are not biological, but, say, cultural. These issues require the introduction of stochastic dynamics, where the relevant equilibrium concept is *stochastic stability*; see Chapter 16. In applications of deterministic evolutionary dynamics, it is a good idea to allow for a large number of initial pure strategies. Extinct strategies stay extinct forever. This will not be a feature of the framework under stochastic dynamics.

An ESS might not exist in certain games, or there could be a multiplicity of ESS for the same game. In the latter case, the history of the game determines which of the two equilibria is stable. This allows evolutionary game theory to provide a powerful equilibrium selection argument in choosing among alternative Nash equilibria. Perhaps surprisingly, the introduction of noise into the underlying dynamics, in the form of persistent mutation, in Chapter 16, allows us to shed this history-dependence. Indeed, with persistent randomness, we shall speak of a limiting ergodic distribution of strategies, irrespective of the history of the system.

While the concept of an ESS is underpinned by the idea of natural selection, it need not maximize the *welfare* of the species. Ex-ante, it is not even clear how one might define the welfare of a species. For instance, is welfare associated with the number of offspring, or the long-term probability of the survival of the species, or its aggregate biomass?

Sometimes a mixed strategy ESS is given two alternative interpretations. First, it could be that players play the same mixed strategy, which is possibly degenerate. In this case, an ESS is said to be played by a *monomorphic population* of players. An alternative interpretation is that all individual players play pure strategies. In a large population, the probability of playing the pure strategy $s_i \in S$ is the fraction of the population that plays the strategy s_i . In this case, an ESS is said to be played by a *polymorphic population* of players. Since the setting in an ESS is static, players who play a mixed strategy have no incentive to mix their pure strategies in the equilibrium proportions. On these grounds, the polymorphic population interpretation is superior. The classic example is the male to female sex ratio of $\frac{1}{2} : \frac{1}{2}$ in most animals with sexual reproduction. In this case, it is clearly polymorphic (or, actually, dimorphic).

14.4 Relation of ESS with other equilibrium concepts

Suppose that in the game $G(2, S, \pi)$, the established strategy in the population is $\sigma \in \Sigma$. The profile (σ, σ) is a Nash equilibrium, if

$$\pi(\sigma, \sigma) \geq \pi(\hat{\sigma}, \sigma) \quad \forall \hat{\sigma} \in \Sigma. \quad (14.19)$$

The degenerate strategy profile (σ, σ) is a strict Nash equilibrium if

$$\pi(\sigma, \sigma) > \pi(\hat{\sigma}, \sigma) \quad \forall \hat{\sigma} \in \Sigma, \quad (14.20)$$

which is simply the primary criterion in an ESS (Definition 14.4). Strict Nash equilibrium arises only in degenerate mixed strategies. However, it is possible that $\pi(\sigma, \sigma) = \pi(\hat{\sigma}, \sigma)$, in which case the secondary criterion in Definition 14.4 may have to be invoked, and we may get a polymorphic equilibrium.

Proposition 14.2

- (a) Every ESS is a Nash equilibrium.
- (b) Any strict Nash equilibrium is an ESS.
- (c) In general, the set of ESS is a strict subset of the set of Nash equilibria. In this sense, ESS is a refinement of Nash equilibrium.

Proof: (a) In an ESS, $\pi(\sigma, \sigma) \geq \pi(\hat{\sigma}, \sigma)$ for all $\hat{\sigma}$, which is also the condition for a Nash equilibrium. Hence, an ESS is a Nash equilibrium.

(b) If we have a strict Nash equilibrium, then for degenerate mixed strategies, we have $\pi(\sigma, \sigma) > \pi(\hat{\sigma}, \sigma)$ for all $\hat{\sigma}$. This is also the primary criterion for an ESS, so σ is also an ESS. Thus, ESS and strict Nash are equivalent.

(c) From (a) and (b), the set of ESS is a subset of the set of Nash equilibria. ■

Proposition 14.3 Suppose that $\sigma \in \Sigma$ is weakly dominated, then it is not an ESS.

Proof: Suppose to the contrary that σ is an ESS, despite being weakly dominated. Then, the primary criterion of ESS: $\pi(\sigma, \sigma) \geq \pi(\hat{\sigma}, \sigma)$ for all $\hat{\sigma}$, is violated. ■

Proposition 14.4 Suppose that $\sigma \in \Sigma$ is a completely mixed ESS. Then it is the unique Nash equilibrium of the game, and $\pi(\sigma, \hat{\sigma}) > \pi(\hat{\sigma}, \hat{\sigma}) \quad \forall \hat{\sigma} \in \Sigma$ such that $\sigma \neq \hat{\sigma}$.

Proof: Since $\sigma \in \Sigma$ is completely mixed, so $\sigma(s_j) > 0$ for all $s_j \in S$. Thus, all pure strategies lie in the support of σ . By the definition of a mixed strategy, we must have $\pi(s_j, \sigma) = \pi(\sigma, \sigma)$. Now pick any other mixed strategy $\hat{\sigma} \in \Sigma$. It follows that

$$\pi(\sigma, \sigma) = \pi(\hat{\sigma}, \sigma). \quad (14.21)$$

By assumption, σ is an ESS. Using, Definition 14.4 and (14.21) we need to invoke the secondary criterion for an ESS, i.e., for all $\forall \hat{\sigma} \in \Sigma, \sigma \neq \hat{\sigma}, \pi(\sigma, \hat{\sigma}) > \pi(\hat{\sigma}, \hat{\sigma})$. It also follows that σ is the unique Nash equilibrium. ■

Consider a strategy $\sigma \in \Sigma$, and denote a ball of radius $\epsilon > 0$ around σ by $B_\epsilon(\sigma)$. The next proposition shows that an ESS is an isolated Nash equilibrium because there is no other strategy in its neighborhood that is a Nash equilibrium.

Proposition 14.5 (Hofbauer and Sigmund, 1998): *A strategy $\sigma \in \Sigma$ is an ESS if there exists some $\epsilon > 0$ such that $\pi(\sigma, \hat{\sigma}) > \pi(\hat{\sigma}, \hat{\sigma})$ for all $\hat{\sigma} \in B_\epsilon(\sigma) \setminus \sigma$.*

Proposition 14.6 (Haigh, 1975): *The game $G(2, S, \pi)$ with a finite strategy set, S , has at most finitely many ESS.*

14.5 Replicator dynamics

Recall that an ESS is a static concept. It would be useful to consider if an ESS arises as a steady state of some dynamic process. There are several dynamic processes, which incorporate the Darwinian idea that more successful strategies in the current period proliferate in the next period, while less successful strategies decay. In this section, we consider a popular adjustment process, the *replicator dynamics*, introduced by Taylor and Jonker (1978).⁹

Suppose that there is a single population. Let us assume that only pure strategies can be inherited from one period to the next. The extension to mixed strategies is non-trivial; see, for instance, Vega-Redondo (1996). If an individual follows a mixed strategy, then he receives the same payoff from all pure strategies in the support of the mixed strategy. Hence, there appears little incentive for him to learn by altering the distribution of his pure strategies. The pure strategy set is given by $S = \{s_1, s_2, \dots, s_n\}$. Suppose that at some instant in time, t , a fraction $x_t(s_i)$ of the population plays the pure strategy s_i . While each individual plays a pure strategy, the state of the system represents frequencies of play of each pure strategy that sum up to one.

Definition 14.5 (State of the system): *At any instant in time, t , the state of the system is denoted by $x_t(s)$. It is a record of the frequency with which each of the pure strategies in the set S is being played, i.e.,*

$$x_t(s) = (x_t(s_1), x_t(s_2), \dots, x_t(s_n)) : x_t(s_j) \geq 0, \forall s_j \in S \text{ and } \sum_{j=1}^n x_t(s_j) = 1.$$

For large populations, the state of the system can be viewed as the mixed strategy profile that a player faces when he is paired against a randomly chosen opponent from the population.

The replicator dynamics specifies how the state of the system evolves over time, and is influenced by the fitness of various pure strategies. The replicator dynamics can be derived in several different ways, but we proceed as follows.¹⁰ Using (14.8), and setting λ_t equal to x_t , we have

⁹ For other adjustment processes, see Hofbauer and Sigmund (1998) and Samuelson (1997).

¹⁰ See, Gintis (2009a) for three alternative methods of deriving the replicator dynamics. We give some microfoundations of replicator dynamics in Section 14.10.2.

$$x_{t+\Delta}(s_i) - x_t(s_i) = \frac{\pi(s_i, \bar{\sigma}) - \pi(\bar{\sigma}, \bar{\sigma})}{\pi(\bar{\sigma}, \bar{\sigma})} x_t(s_i); s_i \in S, \quad (14.22)$$

where the average strategy $\bar{\sigma}$ is defined as the population weighted average strategy,

$$\bar{\sigma} = \sum_{j=1}^n x_t(s_j) s_j.$$

It is more convenient to work with the continuous time analogue of (14.22). Suppose that the current time is $t \in [0, T]$, where T has the interpretation of terminal time. We use the following standard notation, $\dot{x} = \frac{dx}{dt}$. Let Δ be the interval of time between two successive periods. Hence, we can rewrite (14.22) as follows

$$\frac{x_{t+\Delta}(s_i) - x_t(s_i)}{\Delta} = \frac{[\pi(s_i, \bar{\sigma}) - \pi(\bar{\sigma}, \bar{\sigma})] \Delta}{\pi(\bar{\sigma}, \bar{\sigma}) \Delta}; s_i \in S, 0 \leq t \leq T. \quad (14.23)$$

The payoff, π , is now the payoff at time t , while the payoff $\pi \Delta$ arises over the interval Δ . The strategies $s_i, \bar{\sigma}$, give, respectively, the i th pure strategy, and the average strategy at time t . Divide both sides of (14.23) by Δ , and take the limits as $\Delta \rightarrow 0$ to get

$$\dot{x}_t(s_i) = \left(\frac{\pi(s_i, \bar{\sigma}) - \pi(\bar{\sigma}, \bar{\sigma})}{\pi(\bar{\sigma}, \bar{\sigma})} \right) x_t(s_i); s_i \in S, 0 \leq t \leq T. \quad (14.24)$$

The denominator, $\pi(\bar{\sigma}, \bar{\sigma})$, is identical for all pure strategies, hence, it does not affect the relative rates of change of the various pure strategies. Thus, in continuous time, we may consider, instead, the following differential equation.

$$\dot{x}_t(s_i) = [\pi(s_i, \bar{\sigma}) - \pi(\bar{\sigma}, \bar{\sigma})] x_t(s_i); s_i \in S, 0 \leq t \leq T. \quad (14.25)$$

The *replicator dynamics* are given by the system of n equations in (14.25). At each instant in time, an increase in the number of offspring that are hardwired to play a particular strategy, takes the form of asexual reproduction of *exact replica copies* of itself.¹¹ We use replicator dynamics to model the time path of different strategies; these dynamics can also be used to model the time paths of norms, conventions, and cultural forms.

We now consider some of the properties of the replicator dynamics.

1. (*Proportionality*) The rate of change of the frequency with which a pure strategy is played is *proportional* to its fitness, $\pi(s_i, \bar{\sigma})$, relative to the fitness of the average strategy against itself, $\pi(\bar{\sigma}, \bar{\sigma})$. Strategies that perform worse than the average, decay over time.
2. (*Relation to best response dynamics*) There is no presumption that a strategy which does better than the average is also a *best response*. Thus, the condition $\pi(s_i, \bar{\sigma}) > \pi(\bar{\sigma}, \bar{\sigma})$ implies that s_i is a *better response*, but not necessarily a best response to an opponent who plays $\bar{\sigma}$. This is an important difference between the *replicator dynamics* and the *best response dynamics*. Under replicator dynamics, strategies that are not in the best response set might still grow over an interval of time.

¹¹ This is clearly a very simple view of the inheritance properties, say, of genes, from a parent to an offspring. More realistic inheritance properties comprise the subject matter of *population genetics*; the interested reader may wish to pursue Section 5.3 in Hammerstein and Selten (1994).

3. (*Extinct strategies*) If $x_t(s_i) = 0$ for some $s_i \in S$, then $x_{t'}(s_i) = 0, \forall t' > t$. Thus, once a strategy is extinct, it is extinct forever.
4. (*Time behavior of average payoffs*) Replicator dynamics do not imply that the average payoff in the population will grow over time. In the prisoner's dilemma game (see Figure 14.1), suppose that the established strategy is C (cooperation), but a mutation arises in which the mutants play the strategy D (defection). The payoff to the mutants from D will be higher than the average, so D will grow over time. However, the average payoff will continue to fall, relative to the initial state in which all play C.
5. (*Quotient dynamics*) Suppose that s_i, s_j belong to the support of the average strategy, $\bar{\sigma}$. Using (14.25) we get,

$$\frac{\dot{x}_t(s_i)}{x_t(s_i)} - \frac{\dot{x}_t(s_j)}{x_t(s_j)} = \pi(s_i, \bar{\sigma}) - \pi(s_j, \bar{\sigma}). \quad (14.26)$$

Such dynamics in which the difference in the growth rates of two strategies equals the differences in their respective payoffs, against the average strategy, are referred to as *quotient dynamics*.

6. (*Additive invariance of payoffs*) Recall the payoff matrix for player i , Π_i . Consider row i of this payoff matrix, $\Pi_{i\cdot}$. Suppose that we now consider another payoff matrix for player i , $\hat{\Pi}_i$, in which the payoffs in row i , $\hat{\Pi}_{i\cdot}$, are given by $a_{ij} + c_j$. Simple calculations, that we leave for the reader, give rise to the following expressions.

$$\hat{\pi}_i(s_i, \bar{\sigma}) = \pi_i(s_i, \bar{\sigma}) + \sum_j x(s_j) c_j. \quad (14.27)$$

$$\hat{\pi}_i(\bar{\sigma}, \bar{\sigma}) = \pi_i(\bar{\sigma}, \bar{\sigma}) + \sum_j x(s_j) c_j. \quad (14.28)$$

The equation for the replicator dynamics in (14.25) when the payoff matrix is $\hat{\Pi}_i$, is given by

$$\dot{x}_t(s_i) = [\hat{\pi}(s_i, \bar{\sigma}) - \hat{\pi}(\bar{\sigma}, \bar{\sigma})] x_t(s_i); s_i \in S, 0 \leq t \leq T. \quad (14.29)$$

Substituting (14.27), (14.28) in (14.29) we get

$$\dot{x}_t(s_i) = [\pi(s_i, \bar{\sigma}) - \pi_i(\bar{\sigma}, \bar{\sigma})] x_t(s_i); s_i \in S, 0 \leq t \leq T, \quad (14.30)$$

which is the replicator dynamics with respect to the original payoff matrix Π_i . Hence, the replicator dynamics exhibit *additive invariance of payoffs*.

We now consider the dynamic properties of the replicator dynamics when there is a single population of players. We continue to assume that only pure strategies are inherited from one generation to the next. We use the notation

$$x_t(s) = \sigma \Leftrightarrow (x_t(s_1), \dots, x_t(s_n)) = (\sigma_1, \dots, \sigma_n), \quad (14.31)$$

i.e., σ_i is the proportion of the population at time t , $x_t(s_i)$, that plays the pure strategy $s_i \in S$.




Figure 14.2 Relation between replicator dynamics, evolutionary stable strategies, and Nash equilibria.

Definition 14.6 (*Steady state*): A steady state of the system $\bar{x}_t(s) = (\bar{x}_t(s_1), \bar{x}_t(s_2), \dots, \bar{x}_t(s_n))$ is a solution to $\dot{x}_t(s_i) = 0$ for $i = 1, \dots, n$; it is also sometimes referred to as a “dynamic equilibrium” or a “fixed point” of the replicator dynamics.

Proposition 14.7 Suppose that there exists some symmetric Nash equilibrium, $\sigma = (\sigma_1, \dots, \sigma_n)$, of the game $G(2, S, \pi)$, where $x_t(s_i) = \sigma_i$ for $s_i \in S$. Consider the population state $x_t(s)$, defined in (14.31). Then, the population state $x_t(s)$ is a steady state of the replicator dynamics.

Proof: The average strategy is given by $\bar{\sigma} = \sum_{i=1}^n x_t(s_i) s_i$; if s_i is not in the support of σ then $x_t(s_i) = 0$. For any s_i, s_j in the support of a mixed strategy Nash equilibrium, σ , we have

$$\pi(s_i, \bar{\sigma}) = \pi(s_j, \bar{\sigma}). \quad (14.32)$$

In particular, we also have

$$\pi(s_i, \bar{\sigma}) = \pi(\bar{\sigma}, \bar{\sigma}). \quad (14.33)$$

From the definition of replicator dynamics in (14.25), and using (14.32), (14.33) we get that

$$\dot{x}_t(s_i) = 0, i = 1, \dots, n. \quad (14.34)$$

Using Definition 14.6 and (14.34) we get that $x_t(s) = \sigma$ is a steady state of the replicator dynamics. ■

Proposition 14.7 shows that the set of Nash equilibria is a subset of the set of steady states of the replicator dynamics. However, there are steady states of the replicator dynamics that may not be a Nash equilibrium, so the converse need not be true.

We illustrate the relation between the various concepts in Figure 14.2. We have already shown in Proposition 14.7 that the Nash equilibrium is a subset of the steady states of the replicator dynamics (RD). We will now show below that (i) the *asymptotically stable steady states* of the RD are a subset of the set of Nash equilibria, and (ii) the set of *evolutionary stable* strategies is a subset of the set of asymptotically stable states. Finally, if we were to allow for the inheritance of mixed strategies in RD (we only allow the inheritance of pure strategies), then the asymptotically stable states of RD, and the set of ESS, coincide.

Proposition 14.8 Suppose that there is some asymptotically stable state (see Definition 16.41 in Appendix A) of the replicator dynamics given by

$$x^* = (x^*(s_1), x^*(s_2), \dots, x^*(s_n)).$$

Now define a mixed strategy

$$\sigma^* = x^* : \sigma_i^* = x^*(s_i); i = 1, 2, \dots, n.$$

Then σ^* is a Nash equilibrium of the game.

Proof: We briefly sketch the arguments in the proof, assuming quotient dynamics, which follow from replicator dynamics (see above). Suppose that x^* is an asymptotically stable steady state of the replicator dynamics. However, suppose to the contrary, that the corresponding strategy profile $\sigma^* = x^*$ is not a Nash equilibrium.

Given the continuity of payoffs, there exists¹² (1) some $s_i \in \text{supp}(\sigma^*)$, (2) some pure strategy $s_j \notin \text{supp}(\sigma^*)$, and (3) numbers $\delta > 0$ and $\varepsilon > 0$ such that

$$\pi(s_j, \hat{x}) > \pi(s_i, \hat{x}) + \varepsilon,$$

for all $\hat{x}$ in a δ neighborhood of x^* .

Now consider a path that remains within a δ neighborhood of x^* . Using the quotient dynamics in (14.26) we get that

$$\begin{aligned} \frac{\dot{x}_t(s_j)}{x_t(s_j)} - \frac{\dot{x}_t(s_i)}{x_t(s_i)} &= \pi(s_j, \hat{x}) - \pi(s_i, \hat{x}) > \varepsilon. \\ \Rightarrow \frac{\dot{x}_t(s_j)}{x_t(s_j)} - \frac{\dot{x}_t(s_i)}{x_t(s_i)} &> 0. \end{aligned}$$

Thus, the strategy s_j grows faster than the strategy s_i in the whole neighborhood around x^* . This contradicts the original supposition of an asymptotically stable steady state. ■

Proposition 14.8 shows that the asymptotically stable states of the replicator dynamics provide a refinement of the standard Nash concept. The next proposition gives the relation between asymptotically stable states of the replicator dynamics and the set of ESS. This result can be found in Taylor and Jonker (1978), Hofbauer et al. (1979), and Zeeman (1980).

Proposition 14.9 Suppose that an ESS of the game $G(2, S, \pi)$ is given by $\sigma^* = (\sigma_1^*, \sigma_2^*, \dots, \sigma_n^*)$. Then the population state x^* such that

$$x^* = \sigma^* : \sigma_i^* = x^*(s_i); i = 1, 2, \dots, n,$$

is an asymptotically stable state of the replicator dynamics.

We have shown above that the asymptotically stable steady states of the replicator dynamics are a subset of Nash equilibria of the game and an ESS is a subset of the set of Nash equilibria; this underpins Figure 14.2. Since any strategy that does not survive the iterated deletion of strongly dominated strategies cannot be a part of a Nash equilibrium, the next result is fairly intuitive.

Proposition 14.10 Suppose that $x_t = (x_t(s_1), x_t(s_2), \dots, x_t(s_n))$ is a “completely mixed” trajectory of the replicator dynamics. Suppose that pure strategy s_k is eliminated by iterated deletion of strongly dominated strategies. Then, in the replicator dynamics, $\lim_{t \rightarrow \infty} x_t(s_k) = 0$.

¹² The notation $\text{supp}(\sigma^*)$ means the ‘support of the distribution σ^* ’.

The situation is nuanced for weakly dominated strategies. The next result is stated as Theorem 12.5 in Gintis (2009a).

Proposition 14.11 *Suppose that $x_t = (x_t(s_1), x_t(s_2), \dots, x_t(s_n))$ is a “completely mixed” trajectory of the replicator dynamics. Suppose that some pure strategy s_i is weakly dominated by the pure strategy s_k , so $\pi(s_k, s_j) > \pi(s_i, s_j)$ for some $s_j \in S$. In the replicator dynamics, if $\lim_{t \rightarrow \infty} x_t(s_j) > 0$, then $\lim_{t \rightarrow \infty} x_t(s_i) = 0$.*

REPLICATOR DYNAMICS FOR THE ASYMMETRIC POPULATION CASE

Suppose that there are several populations $p = 1, 2, \dots, P$. Each population, p , has n_p pure strategies. An important issue in setting up the replicator dynamics is the relative sizes of the different populations. Suppose, for instance, that $P = 2$, and the second population is twice the size of the first population. Then, in random pairwise meetings between the two populations, players from the first population will have twice as many matchings as compared to players from the second population. This will clearly have an effect on the population dynamics of each population. We abstract from these issues by assuming that each population is of the same size.

Subscript variables with a p denote the p th population. Suppose that s_{pi} is the i th pure strategy in the p th population, $\bar{\sigma}_p$ is the average strategy in the p th population, and $\bar{\sigma}_{-p}$ is the average mixed strategy profile of all the other populations other than the p th. Then, we can write the replicator dynamics (14.30) as follows.

$$\dot{x}_t(s_{pi}) = [\pi(s_{pi}, \bar{\sigma}_{-p}) - \pi(\bar{\sigma}_p, \bar{\sigma}_{-p})] x_t(s_{pi}); i = 1, \dots, n_p; p = 1, 2, \dots, P; t \geq 0. \quad (14.35)$$

Most of the results for the single population case survive to the multiple populations case. However, an ESS in the asymmetric populations case must be a strict Nash equilibrium; see Gintis (2009a) for a discussion. Section 14.6.4 considers a concrete example for the case of $P = 2$.

14.6 Applications of evolutionary game theory

We now consider several applications of evolutionary game theory in increasingly complicated problems.

14.6.1 Prisoner's dilemma game: monomorphic equilibrium

We consider separately a static and a dynamic version of the game.

STATIC PRISONER'S DILEMMA GAME

Consider the two-player, symmetric, prisoner's dilemma game shown in Figure 14.1, which has a strictly dominated strategy, C . The unique Nash equilibrium of this game is (D, D) . In contrast to classical game theory, individuals in evolutionary game theory are *hardwired* to play either C or D . Suppose that a fraction x of the population plays C , and the remaining fraction $1 - x$ plays D . Thus, the average strategy in the population is given by $\bar{\sigma} = xC + (1 - x)D$. We assume that the population is large, so in pairwise random meetings, it is *as if* any player meets an opponent who plays the mixed strategy $\bar{\sigma}$. The payoff from the respective strategies, C and D , is given by:

$$\pi(C, \bar{\sigma}) = 2x, \pi(D, \bar{\sigma}) = 1 + 2x. \quad (14.36)$$




Figure 14.3 Graph of the replicator dynamics for the PD game.

The payoff of the average strategy against itself is given by

$$\pi(\bar{\sigma}, \bar{\sigma}) = x\pi(C, \bar{\sigma}) + (1-x)\pi(D, \bar{\sigma}) = 1+x. \quad (14.37)$$

The replicator dynamics in (14.25) imply that the time evolution of the fraction of the population that plays strategy C is given by

$$\dot{x} = x[\pi(C, \bar{\sigma}) - \pi(\bar{\sigma}, \bar{\sigma})]. \quad (14.38)$$

Substituting (14.36), (14.37) in (14.38), we get a first order non-linear differential equation.

$$\dot{x} = f(x) = x(x-1). \quad (14.39)$$

The steady state, $\bar{x}$, is found by solving $\dot{x} = 0 \Leftrightarrow x(x-1) = 0$, which has two solutions shown in Figure 14.3 (recall that x lies between 0 and 1)

$$\bar{x} = 0, \bar{x} = 1. \quad (14.40)$$

The steady state $\bar{x} = 0$ corresponds to an equilibrium in which everyone plays D, while in the steady state $\bar{x} = 1$ everyone plays C.

Recall that the set of ESS is a subset of the set of steady states of the replicator dynamics (see Figure 14.2). We wish to check if either of the two steady states (C, C) or (D, D) is an ESS. Using the *primary criterion* for ESS in Definition 14.4, D is ESS if $\pi(D, D) > \pi(C, D)$, which is true (see the payoff matrix in Figure 14.1), hence, D is an ESS. On the other hand, the state C is neither a Nash equilibrium nor an ESS.

What is the relation of ESS to the asymptotically stable states of the replicator dynamics in this example? Figure 14.3 shows a plot of the replicator dynamics in (14.39). Using the fact that $\frac{1}{x(x-1)} = \frac{-1}{x} + \frac{1}{(x-1)}$, the differential equation in (14.39) can be written in the *variable separable form*.

$$\int \left[\frac{-1}{x} + \frac{1}{(x-1)} \right] dx = \int dt.$$

A simple calculation shows that the solution is given by

$$x(t) = \left[1 - e^t \left(1 - \frac{1}{x(0)} \right) \right]^{-1}. \quad (14.41)$$

Taking the limits on both sides in (14.41) we get that

$$\lim_{t \rightarrow \infty} x(t) = \begin{cases} 0 & \text{if } x(0) > 1 \\ 1 & \text{if } x(0) = 1 \\ 0 & \text{if } x(0) < 1 \end{cases}. \quad (14.42)$$

Equation (14.42) brings out the important role of *history*, or initial conditions, in determining the steady state that the replicator dynamics gravitates towards. Unless all individuals begin with the play of strategy *C* (i.e., $x(0) = 1$), the entire population always drifts off to $x(t) = 0$, i.e., all play *D*. Thus, strategy *C* is an unstable steady state. Alternatively, one may check for stability in the following direct manner. From (14.39), $f'(x) = 2x - 1$, hence, $f'(0) = -1 < 0$, and $f'(1) = 1 > 0$. Using standard stability arguments (see Appendix A), the ESS, $\bar{x} = 0$, is also the stable equilibrium, while $\bar{x} = 1$ is an unstable steady state.

Multiple equilibria are a pervasive problem in standard game theory. Evolutionary game theory provides a *history-dependent refinement* of Nash equilibrium.

DYNAMIC PRISONER'S DILEMMA GAME

In evolutionary game theory, individuals are hardwired to play certain strategies. However, the hardwired strategies can be fairly sophisticated. Furthermore, in repeated games, hardwired strategies, of any complexity, depend on the history of play so far. In order to illustrate this point in a simple example, consider a twice repeated version of the prisoner's dilemma. As usual, in each of the two periods, two players from the population are randomly matched to play the game.

Three possible strategies in the twice repeated game are: always cooperate (*C*), always defect (*D*), or tit-for-tat (*T*); these were examined in a well-known study by Axelrod and Hamilton (1981). An individual who uses strategy *T*, cooperates in the first period, but in the second period, chooses the following history-contingent action: cooperate if the opponent cooperated in period 1, and defect if the opponent defected in period 1 (tit-for-tat). The payoff matrix, shown in Figure 14.4, adapts the single period payoffs in Figure 14.1, to two periods. Intertemporal payoffs are simply the sum of payoffs over the two periods (no discounting is applied).

When two players who are hardwired to play *C* (for both periods) meet, they each get a payoff of $2 + 2 = 4$. Similarly, when two *D* players meet, they each get a payoff of $1 + 1 = 2$. When a *C* player meets a *D* player, the former gets $0 + 0 = 0$, and the latter gets $3 + 3 = 6$. When a *T*

	C	D	T
C	4, 4	0, 6	4, 4
D	6, 0	2, 2	4, 1
T	4, 4	1, 4	4, 4

Figure 14.4 Payoff matrix for the two-period tit-for-tat game.

player meets a C player, then both cooperate in both periods, so each gets $2 + 2 = 4$. When a T player meets an opponent who plays D , then, the T player plays C in the first period, and D in the second period (for a payoff of $0 + 1 = 1$), while the D player simply plays D in both periods for a payoff of $3 + 1 = 4$.

The strategy C is weakly dominated by strategies D , and T , hence, it cannot be an ESS (see Proposition 14.3). Let us focus on the remaining strategies D and T ; the relevant payoffs are given in italics in the symmetric, lower right, 2×2 submatrix in Figure 14.4. There are two Nash equilibria, (D, D) and (T, T) .

We wish to check if T is an ESS, starting from a population in which all play T . Applying the primary criterion for ESS (Definition 14.4), T is an ESS if $\pi(T, T) > \pi(D, T)$. But $\pi(T, T) = \pi(D, T) = 4$, hence, the primary criterion cannot distinguish between the two strategies. The secondary criterion (Definition 14.4) requires $\pi(T, D) > \pi(D, D)$. However, $\pi(T, D) = 1$, $\pi(D, D) = 2$, hence, the strategy T is not an ESS. By contrast, starting from a population in which all play D , let us introduce a small fraction of mutants who play T . In this case, $\pi(D, D) > \pi(T, D)$, hence, the strategy D is the only ESS in the game. Thus, ESS allows us to choose among multiple Nash equilibria.

Consider the replicator dynamics when there are only two strategies, D and T . Suppose that the established strategy is D , and a fraction x of the population mutates to play the strategy T . Thus, the average strategy is $\bar{\sigma} = xT + (1 - x)D$. Individuals are randomly paired, once, to play the sequential two-period prisoner's dilemma game. Players play pure strategies only, either T or D . The payoffs from playing the respective strategies, D and T , against the average strategy are given by,

$$\pi(D, \bar{\sigma}) = 2(1 - x) + 4x = 2(1 + x), \quad (14.43)$$

$$\pi(T, \bar{\sigma}) = (1 - x) + 4x = 1 + 3x. \quad (14.44)$$

The payoff of the average strategy against itself is given by

$$\pi(\bar{\sigma}, \bar{\sigma}) = x\pi(T, \bar{\sigma}) + (1 - x)\pi(D, \bar{\sigma}) = x^2 + x + 2. \quad (14.45)$$

The replicator dynamics are given by

$$\dot{x} = x[\pi(T, \bar{\sigma}) - \pi(\bar{\sigma}, \bar{\sigma})]. \quad (14.46)$$

Substituting (14.44), (14.45), in (14.46) we have

$$\dot{x} = f(x) = x(1 - x)(x - 1). \quad (14.47)$$

There are two steady states: $\bar{x} = 0$ and $\bar{x} = 1$. Let us now check for the asymptotic stability of these steady states. From (14.47), $f'(x) = -3x^2 + 4x - 1$, hence, $f'(0) = -1 < 0$, and $f'(1) = 0$. Thus, only $\bar{x} = 0$ (where none play T) satisfies the condition for asymptotic stability. The dynamics are shown in Figure 14.5.

In the static and dynamic prisoner's dilemma games, an ESS is a *monomorphic equilibrium*. In our next example, we consider the possibility of a *polymorphic equilibrium*.




Figure 14.5 Dynamics of the tit-for-tat strategy in a repeated prisoner's dilemma game.

	C	M
C	0, 0	-2, 2
M	2, -2	-4, -4

Figure 14.6 Payoff matrix for the game of chicken.

14.6.2 *The game of chicken: polymorphic equilibria*

We now consider the game of chicken. Several versions of the game are given in the literature. In the most common version, two drivers race simultaneously towards a cliff. The one who stops first (strategy C, which is short for “chicken”) loses, and the one who stops last (strategy M, for “macho”) wins.

The symmetric normal form is given in Figure 14.6. There are two pure strategy Nash equilibria, (C, M) and (M, C) , and a mixed strategy Nash equilibrium in which each of the pure strategies is played with probability one half. None of the steady states, (C, C) and (M, M) , satisfy the conditions for an ESS because the set of ESS is a subset of the set of Nash equilibria.

Let us now use the replicator dynamics to look at the evolution of each of the pure strategies in the population. Suppose that initially, the prevalent strategy in the population is C. Consider an invasion by the mutant strategy, M, which is played by a fraction, x , of the population. The average strategy in the population is given by $\bar{\sigma} = xM + (1 - x)C$. We can now compute the following payoffs against the average strategy:

$$\pi(C, \bar{\sigma}) = -2x; \pi(M, \bar{\sigma}) = 2(1 - 3x). \quad (14.48)$$

$$\pi(\bar{\sigma}, \bar{\sigma}) = (1 - x)\pi(C, \bar{\sigma}) + x\pi(M, \bar{\sigma}) = -4x^2. \quad (14.49)$$

Using the replicator dynamics, we get

$$\dot{x} = x[\pi(M, \bar{\sigma}) - \pi(\bar{\sigma}, \bar{\sigma})]. \quad (14.50)$$

Substituting (14.48), (14.49) in (14.50), we get

$$\dot{x} = f(x) = x(4x^2 - 6x + 2). \quad (14.51)$$




Figure 14.7 A plot of the dynamic equation of motion in the game of chicken.

A plot of $f(x)$ is shown in Figure 14.7. There are three steady states, $\bar{x} = 0$, $\bar{x} = 0.5$, and $\bar{x} = 1$. To check for stability, note that $f'(x) = 12x(x - 1) + 2$. Hence, $f'(0) > 0$, and $f'(1) > 0$, but $f'(0.5) < 0$, thus, the only stable steady state is $\bar{x} = 0.5$.

The asymptotic solution is a polymorphic equilibrium in which half the individuals play C , while the other half play M . The monomorphic equilibrium in which all play C ($\bar{x} = 0$) is unstable because a small number of mutants who play M will be very successful; they get $2 > 0$. Analogously, the monomorphic equilibrium in which everyone is macho is unstable. In this case, there are too many people who are headed for the cliff, and each gets a payoff of -4 ; an encroachment by a small number of mutants who play C will ensure a payoff $-2 > -4$. Thus, the only stable equilibrium is an interior one.

So far, we have restricted attention only to pure strategies. Now consider the mixed strategy W that puts equal weight on the pure strategies C and M , so $W = \frac{1}{2}C + \frac{1}{2}M$. We now wish to check if W is an ESS, when W is the initially prevalent strategy. First check if W can be invaded, sequentially, by a mutant population that plays the pure strategy M or C ; the fraction of the mutant population at time t is x_t . A simple calculation gives the following payoffs.

$$\pi(W, W) = \frac{1}{2} \left\{ \frac{1}{2}\pi(C, C) + \frac{1}{2}\pi(C, M) \right\} + \frac{1}{2} \left\{ \frac{1}{2}\pi(M, C) + \frac{1}{2}\pi(M, M) \right\} = -1. \quad (14.52)$$

$$\pi(W, M) = \frac{1}{2}\pi(C, M) + \frac{1}{2}\pi(M, M) = -3. \quad (14.53)$$

$$\pi(W, C) = \frac{1}{2}\pi(M, C) + \frac{1}{2}\pi(C, C) = 1.$$

$$\pi(M, W) = \frac{1}{2}\pi(M, C) + \frac{1}{2}\pi(M, M) = -1. \quad (14.54)$$

$$\pi(C, W) = \frac{1}{2}\pi(C, C) + \frac{1}{2}\pi(C, M) = -1. \quad (14.55)$$

Let us first check if W is an ESS. W is an ESS against M if $\pi(W, W) > \pi(M, W)$. But $\pi(W, W) = \pi(M, W)$, hence, the primary criterion is not decisive. A similar issue arises when

the mutants play C because $\pi(W, W) = \pi(C, W)$. Thus, we need to resort to the secondary criterion. For mutations involving pure strategies M, C , one at a time, we can check the following.

$$W \text{ against } M : \pi(W, M) > \pi(M, M) \Leftrightarrow -3 > -4.$$

$$W \text{ against } C : \pi(W, C) > \pi(C, C) \Leftrightarrow 1 > 0.$$

Based on the secondary criterion, W is an ESS against successive mutations to strategies M and C .

14.6.3 Non-existence of an asymptotically stable equilibrium

In evolutionary game theory, an asymptotically stable steady state may not exist if there are more than two strategies, even when players play *completely mixed strategies* (i.e., all pure strategies are played with a strictly positive probability). Consider, for instance, a two-player zero-sum game, the *rock-scissors-paper game*. The normal form is shown in Figure 14.8.

There are three strategies, rock (R), paper (P), and scissors (S). When players who play identical strategies meet, the payoffs are $(0, 0)$; see the payoffs along the diagonal. When strategies are different, one strategy wins, for a payoff of 1, and the other loses, for a payoff of -1 . R wins against S , S wins against P , and P wins against R .

There are no pure strategy Nash equilibria in this game. Denote the probability of play of R, S , and P , respectively, by σ_1, σ_2 , and $1 - \sigma_1 - \sigma_2$. In the unique Nash equilibrium in mixed strategies, each of the three strategies is played with equal probability, so

$$\sigma^N = (\sigma_1, \sigma_2, 1 - \sigma_1 - \sigma_2) = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right).$$

We have already shown that asymptotically stable steady states of the replicator dynamics and the ESS are both subsets of Nash. Thus, the unique Nash equilibrium, σ^N , is the only candidate strategy profile that could be both an ESS and an asymptotically stable state. Let the population proportions playing the three strategies, R, S , and P , at time t be, respectively $x_t(R)$, $x_t(S)$, and $1 - x_t(R) - x_t(S)$. The average population strategy, $\bar{\sigma}$, is given by

$$\bar{\sigma} = x_t(R)R + x_t(S)S + (1 - x_t(R) - x_t(S))P.$$

Since we have a zero-sum game, for any pair of strategies, the payoff from the average strategy against the average strategy is zero, i.e., $\pi(\bar{\sigma}, \bar{\sigma}) = 0$. Also note that

$$\pi(R, \bar{\sigma}) = x_t(S) - [1 - x_t(R) - x_t(S)], \quad (14.56)$$

	R	S	P
R	0, 0	1, -1	-1, 1
S	-1, 1	0, 0	1, -1
P	1, -1	-1, 1	0, 0

Figure 14.8 The rock-scissors-paper game.

$$\pi(S, \bar{\sigma}) = -x_t(R) + [1 - x_t(R) - x_t(S)]. \quad (14.57)$$

The replicator dynamics is given by the following system of two differential equations.¹³

$$\dot{x}_t(R) = x_t(R) [\pi(R, \bar{\sigma}) - \pi(\bar{\sigma}, \bar{\sigma})]. \quad (14.58)$$

$$\dot{x}_t(S) = x_t(S) [\pi(S, \bar{\sigma}) - \pi(\bar{\sigma}, \bar{\sigma})]. \quad (14.59)$$

Using $\pi(\bar{\sigma}, \bar{\sigma}) = 0$, and (14.56), (14.57), we can rewrite (14.58), (14.59) as follows.

$$\dot{x}_t(R) = f_1(x_t(R), x_t(S)) = x_t(R) [2x_t(S) + x_t(R) - 1]. \quad (14.60)$$

$$\dot{x}_t(S) = f_2(x_t(R), x_t(S)) = x_t(S) [-2x_t(R) - x_t(S) + 1]. \quad (14.61)$$

To find the interior steady state of the system, set $\dot{x}_t(R) = \dot{x}_t(S) = 0$. This gives us two equations in two unknowns, $2x_t(S) + x_t(R) - 1 = 0$ and $-2x_t(R) - x_t(S) + 1 = 0$. The solution to these equations gives the steady state:

$$\mathbf{x}^* = (x^*(R), x^*(S), x^*(P)) = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right).$$

(14.60), (14.61) is a system of two non-linear differential equations. To check for the stability of the steady state, we need to evaluate the Jacobian of $\mathbf{f} = (f_1, f_2)$ at the steady state $\mathbf{x}^*$:

$$D\mathbf{f}(\mathbf{x}^*) = \begin{bmatrix} \frac{\partial f_1(\mathbf{x}^*)}{\partial x_t(R)} & \frac{\partial f_1(\mathbf{x}^*)}{\partial x_t(S)} \\ \frac{\partial f_2(\mathbf{x}^*)}{\partial x_t(R)} & \frac{\partial f_2(\mathbf{x}^*)}{\partial x_t(S)} \end{bmatrix} = \begin{bmatrix} \frac{1}{3} & \frac{2}{3} \\ -\frac{2}{3} & -\frac{1}{3} \end{bmatrix}.$$

To find the eigenvalues, solve for $|D\mathbf{f}(\mathbf{x}^*) - \lambda I| = 0$. This gives the equation $9\lambda^2 + 3 = 0$. The two complex roots are $0 \pm i\sqrt{\frac{1}{3}}$. When the real part of the eigenvalue is zero, the steady state of the system is given by a neutrally stable center. Furthermore, the center is surrounded by closed orbits. This is shown in Figure 14.9. So, unless the dynamical system starts off in the steady state, it will never approach the steady state. Thus, there are no asymptotically stable steady states. Since the set of ESS is a subset of these states, an ESS does not exist in this case. It turns out that a small amount of noise in the payoffs can alter these results. One of the exercises asks you to solve such a problem.

14.6.4 A two-population model of coordination

Consider the coordination game shown in Figure 14.10. Two players, the row player (player 1) and the column player (player 2), can choose either action A or D . If the players coordinate on the same action, they have a payoff of 1, otherwise they get a payoff of 0. There are two Nash equilibria in the game, (A, A) and (D, D) . There is also a mixed strategy Nash equilibrium in which both players play A and D with equal probability

¹³ The time path of the evolution of $x_t(P)$ can be found residually from $x_t(P) = 1 - x_t(R) - x_t(S)$.




Figure 14.9 Orbits for the rock-scissors-paper game.

	A	D
A	1, 1	0, 0
D	0, 0	1, 1

Figure 14.10 A two-player coordination game.

Suppose that there are two separate populations of players 1 and 2, who are randomly matched each period. The respective fractions of populations 1 and 2 who play A are x and y . Thus, the average strategies of the two populations are $\bar{\sigma}_1, \bar{\sigma}_2$:

$$\bar{\sigma}_1 = xA + (1-x)D, \quad \bar{\sigma}_2 = yA + (1-y)D.$$

Let π_1 and π_2 denote the payoffs in populations 1 and 2. Using (14.35), the replicator dynamics for the case of two populations is written as:

$$\dot{x}_t = x_t [\pi_1(A, \bar{\sigma}_2) - \pi_1(\bar{\sigma}_1, \bar{\sigma}_2)]. \quad (14.62)$$

$$\dot{y}_t = y_t [\pi_2(A, \bar{\sigma}_1) - \pi_2(\bar{\sigma}_2, \bar{\sigma}_1)]. \quad (14.63)$$

We can now compute the following payoffs in both populations.

$$\pi_1(A, \bar{\sigma}_2) = y, \quad \pi_1(D, \bar{\sigma}_2) = 1 - y, \quad \pi_2(A, \bar{\sigma}_1) = x, \quad \pi_2(D, \bar{\sigma}_1) = 1 - x. \quad (14.64)$$

$$\pi_1(\bar{\sigma}_1, \bar{\sigma}_2) = x\pi_1(A, \bar{\sigma}_2) + (1-x)\pi_1(D, \bar{\sigma}_2) = 1 - x - y(1 - 2x). \quad (14.65)$$

$$\pi_2(\bar{\sigma}_2, \bar{\sigma}_1) = y\pi_2(A, \bar{\sigma}_1) + (1-y)\pi_2(D, \bar{\sigma}_1) = 1 - x - y(1 - 2x). \quad (14.66)$$

Substituting (14.64)–(14.66) in (14.62) and (14.63), we get.

$$\dot{x}_t = f_1(x_t, y_t) = x_t(2y_t - 1)(1 - x_t). \quad (14.67)$$

$$\dot{y}_t = f_2(x_t, y_t) = y_t(2x_t - 1)(1 - y_t). \quad (14.68)$$




Figure 14.11 Coordination game in the two-population case.

Setting $\dot{x}_t = \dot{y}_t = 0$, we can find the following steady states of the system.

$$(\bar{x}, \bar{y}) = (0, 0), (1, 0), (0, 1), (1, 1), \left(\frac{1}{2}, \frac{1}{2}\right).$$

To determine the asymptotic stability of the steady states, we first form the Jacobian of the appropriate matrix.

$$Df(\bar{x}, \bar{y}) = \begin{bmatrix} \frac{\partial f_1(\bar{x}, \bar{y})}{\partial x_t} & \frac{\partial f_1(\bar{x}, \bar{y})}{\partial y_t} \\ \frac{\partial f_2(\bar{x}, \bar{y})}{\partial x_t} & \frac{\partial f_2(\bar{x}, \bar{y})}{\partial y_t} \end{bmatrix} = \begin{bmatrix} (2y_t - 1)(1 - 2x_t) & 2x_t(1 - x_t) \\ 2y_t(1 - y_t) & (2x_t - 1)(1 - 2y_t) \end{bmatrix}.$$

We now check the stability of each of the steady states, using the relevant stability conditions.

1. $(\bar{x}, \bar{y}) = (0, 0)$: The two repeated eigenvalues are $r_1 = -1$ and $r_2 = -1$, hence, this is an asymptotically stable steady state.
2. $(\bar{x}, \bar{y}) = (1, 0)$: The two repeated eigenvalues, $r_1 = 1$ and $r_2 = 1$, are positive. Thus, the state $(1, 0)$ is asymptotically unstable.
3. $(\bar{x}, \bar{y}) = (0, 1)$: The two repeated eigenvalues are $r_1 = 1$ and $r_2 = 1$, so the state $(1, 0)$ is asymptotically unstable.
4. $(\bar{x}, \bar{y}) = (1, 1)$: The two repeated eigenvalues are $r_1 = -1$ and $r_2 = -1$, hence, this is an asymptotically stable steady state.
5. $(\bar{x}, \bar{y}) = (\frac{1}{2}, \frac{1}{2})$: The two eigenvalues are $r_1 = \frac{1}{2}$ and $r_2 = -\frac{1}{2}$. One of the eigenvalues is positive, and the other is negative, so we have a saddle point.

Figure 14.11 shows some representative trajectories of the dynamic system. The $\dot{x} = 0$ locus is shown as the horizontal line at $y = 1/2$, and the $\dot{y} = 0$ locus is shown as the vertical line at $x = 1/2$. The direction of motion is shown by arrows in each of the resulting four regions. If the proportion of any population that plays A exceeds $1/2$, then the proportion of the other population that plays A is also growing. Hence, in the region $x > 1/2$ and $y > 1/2$, the proportion of each population playing A is growing. In this case, one converges to the stable steady state $(1, 1)$. In the opposite case, $x < 1/2$ and $y < 1/2$, one converges to the stable steady state $(0, 0)$. Finally, the unique saddlepath is shown as the SS curve, which converges to the steady state $(\frac{1}{2}, \frac{1}{2})$.

14.7 Playing the field

In some animal interactions, it is not possible to isolate pairwise interactions, and it appears that animals are playing against all other animals in the population. Animals are then said to be *playing the field*. Our brief formalization here follows Hammerstein and Selten (1994). Suppose that Σ is the convex hull of finitely many points. Elements of Σ may be mixed strategies (as in this chapter), but more generally they could also be interpreted as vectors of biological parameters. For consistency with the rest of the chapter, we refer to elements of Σ as strategies.

Suppose that we are interested in the monomorphic stability of some strategy $p \in \Sigma$. Let there be an invasion by a fraction $\varepsilon \in [0, 1]$ of mutants who play another strategy, $q \in \Sigma$. Consider a *bimorphic population state* $z = (1 - \varepsilon)p + \varepsilon q \in \Sigma$.

Definition 14.7 (*Population game*): A population game $G = (n, \Sigma, \pi)$ is a triplet that specifies the number of players in the population n , the set Σ , and a payoff function π with the following properties.

- (i) The set Σ is the convex hull of finitely many points.
- (ii) The payoff function $\pi : \mathbb{R}^m \times \mathbb{R}^m \rightarrow \mathbb{R}$. $\pi(r, z)$ is the payoff to a player who follows strategy $r \in \Sigma$ and plays against a field where the average strategy is $z = (1 - \varepsilon)p + \varepsilon q \in \Sigma$, $\varepsilon \in [0, 1]$. We require that $\pi(r, z)$ be defined for all $r, z \in \Sigma$.
- (iii) For $\varepsilon = 0$ we have $z = p$. We require that $\pi(r, p)$ does not depend on q .

The strategy $r \in \Sigma$ is a best response to $p \in \Sigma$ if $\pi(r, p) \geq \pi(q, p)$ for all $q \in \Sigma$. If r is the unique best response, then it is a strict best reply to p .

Definition 14.8 (*ESS in the population game; Hammerstein and Selten, 1994*): Consider the population game $G = (n, \Sigma, \pi)$ (Definition 14.7). A strategy $p \in \Sigma$ is an evolutionary stable strategy if:

- (a) *Equilibrium condition or primary criterion*: $\pi(p, p) > \pi(q, p)$ for all $q \in \Sigma$.
- (b) *Stability condition or secondary criterion*: Consider the set of best replies to p that are different from p , $BR(p) = \{r \neq p : \pi(p, p) = \pi(r, p), r, p \in \Sigma\}$. For any $r \in BR(p)$ there exists ε_r such that

$$\pi(p, (1 - \varepsilon)p + \varepsilon r) > \pi(r, (1 - \varepsilon)p + \varepsilon r) \text{ for all } \varepsilon \in (0, \varepsilon_r).$$

The ESS result in Definition 14.8 can be stated for a uniform invasion barrier in which ε_r is independent of $r \in BR(p)$. This requires suitable restrictions on the payoff function.¹⁴ For a further discussion of these issues the reader may consult Hammerstein and Selten (1994).

14.8 Gene-culture coevolution

Part 2 of the book considered a range of human prosocial behaviors. What accounts for such prosocial concerns among humans? Why are *character virtues* possessed by many humans, and valued by most? Two sorts of contrasting answers were given to these questions in two different behavioral sciences. Economics provided an account based on models of self-interest that form

¹⁴ For instance, $\pi(r, p)$ is required to be linear in the first argument; see Crawford (1990a,b).

the basis of neoclassical economics. There is now compelling evidence against the purely self-interested approach that is reviewed in Part 2.

At the other extreme, sociology advocated the *cultural hegemony model* in which individuals passively accepted the dominant culture that is the source of the relevant norms associated with various social roles; the dominant culture may have been put into place by the ruling or dominant class. The predictions of this model too are not consistent with the evidence. There are many instances in which the dominant cultures are replaced by minority counter cultures, or even assimilated by them. For instance, African Americans in the era of the civil rights movement, or the failed attempts to indoctrinate people into socialist morality in the former Soviet Union; see Gintis et al. (2014).

An alternative is the model of *gene-culture coevolution*, which is the study of the transmission of genes and culture from one generation to the next, as well as the interaction between them.

14.8.1 Basics of gene-culture coevolution

Bowles and Gintis (2011, p. 13) define *culture* as “the ensemble of preferences and beliefs that are acquired by means other than genetic transmission.” Wilson (2012, p. 213) defines culture thus: “culture is the combination of traits that distinguishes one group from another. A cultural trait is a behavior that is either first invented within a group, or else learned from another group, then transmitted among members of the group.” These definitions can be used to identify the presence of cultures in other animals as well. Culture is typically transmitted by imitation. Examples include: Chimpanzees that use a branch of a tree, stripped of its leaves, poke it into a termite mound, and then eat the soldier termites that stick to the branch; or the fishing practices of dolphins.

Interestingly, the Neanderthals exhibited very little culture in the form of art and ornaments, in contrast to their close relatives, *Homo sapiens*, despite having a larger brain size. Yet, it appears that *Homo sapiens* overpowered, and probably eliminated the Neanderthals (Wilson, 2012). What accounts for the success of *Homo sapiens* relative to other related species, and their unprecedented increase in brain size over time? Several factors have been suggested.

An influential line of research, associated mainly with the names of J. Tooby, L. Cosmides, and S. Pinker, advances the *cognitive niche hypothesis*.¹⁵ In this view, non-human animals typically possess *domain-specific* or *dedicated intelligence* that is suitable for adapting to the particular environment that the animal lives in. In contrast, humans possess *improvisational intelligence*, which implies sufficient learning flexibility in the brain that allows them to adapt quickly to new and strange environments. Such adaptation, a form of culture, is then quickly shared with others who avoid the information acquisition costs of adaptation.

Another important element behind the success of *Homo sapiens* may well be their willingness and instinct to engage in cooperative activities and shared goals. For instance, bipedalism, the development of fire, establishments of campsites, agriculture, and the development of *eusociality*. This has been termed as the *cultural intelligence hypothesis* in which the development of language played a key role; see, for instance, Tomasello (1999).¹⁶

Boyd et al. (2011) do not deny the importance of superior human cognitive abilities embodied in the cognitive niche hypothesis. However, they argue that by itself this cannot provide an adequate explanation of human adaptability. They argue instead for the *cultural niche hypothesis* in

¹⁵ See, for instance, Barrett et al. (2007), Tooby and Cosmides (1992), and Pinker (2010).

¹⁶ For a fascinating account of the development of language, see Wilson (2012, ch. 22).

which humans learn over many generations, and evolve culturally transmitted bodies of knowledge in which the causal link between variables is often unknown. Consider the following two examples.

Children are programmed to learn by imitation. In experiments, even when learning involves irrelevant steps that are clearly identified as unnecessary by visual inspection, children copy all steps, including the irrelevant steps, but chimpanzees omit the irrelevant steps. Humans have also learned to use chilli peppers in meat, despite their unpalatable taste, to protect themselves against food-borne pathogens. However, this occurred prior to an understanding of the causal mechanism. Boyd et al. (2011, p. 10924) close with the quote: “So it does not make sense to ask, does culture overcome biology? The right question to ask is, how do genetic and cultural inheritance interact to produce the observed patterns of human psychology and behavior?”

There is now a compelling case for *gene–culture coevolution*.¹⁷ To motivate the theory, one may distinguish between three kinds of changes in the environment, and our response to them.

1. Adaptability to extremely slow changes in the environment are directly encoded into our DNA.
2. It is not efficient to code extremely rapid changes in the environment directly into the DNA. Genetic changes are slow and conservative. In humans, the ability to respond to rapid environmental change is incorporated into brain structures that allow for a flexible capacity to deal with novel situations as a form of problem solving.
3. For environmental changes that are neither very rapid nor very slow, a third form of response is optimal. This takes the form of non-genetic or *epigenetic changes*. These are physiological developments that are not genetically hardwired; neither are they involuntary reflexes such as eyeblinks; for a review see Gintis (2009b).

Wilson (2012, p. 194) gives a cogent introduction to epigenetic changes: “The behaviors created by epigenetic rules are not hardwired like reflexes. It is the epigenetic rules instead that are hardwired, and hence compose the true core of human nature. These behaviors are learned, but the process is what psychologists call “prepared.” In prepared learning, we are innately predisposed to learn, and thereby reinforce one option over another.” For instance, children learn a fear of snakes very rapidly, but do not a fear pigeons or rabbits. *Prepared learning* allows us to find a landscape as beautiful, but we would not classify an open waste pit as beautiful.¹⁸ This view is in contrast to the *blank slate view of the human mind* in which language and culture are learned through the environment.

A prominent example of an epigenetic transmission mechanism is *culture*. Humans, over their evolutionary history, are *predisposed* to learn various aspects of culture, e.g., language, through three channels. The *vertical channel* (parents to children), the *oblique channel* (from non-parent elders to younger people), and the *horizontal channel* (between peers). This is a form of epigenetic transmission because no changes of the human DNA are involved. For instance, no gene that is exclusively responsible for cooperative behavior has been discovered. Bowles and Gintis (2011, p. 16) write: “the idea of one to one mapping between genes and behavior is unlikely given what is now known about gene expression, and is implausible in light of the complexity

¹⁷ For early, and pioneering, work on gene–culture coevolution, see Lumsden and Wilson (1980, 1981). For a fascinating, authoritative, and non-technical, introduction, see Wilson (2012). For a focus on cultural evolution, see Cavalli-Sforza and Feldman, (1982), Boyd and Richerson (1985), Dunbar (1993), and Richerson and Boyd (2004).

¹⁸ Lumsden and Wilson (1980, 1981) stressed the importance of the human instinct, as arising through prepared learning.

and cultural variation of cooperative behaviors.” For these reasons, we interpret culture at the level of the observed phenotype rather than the unobserved genotype.

The evidence suggests that culture is directly encoded in the brain (Gintis, 2011, pp. 379–80). For instance, the prefrontal cortex is associated with human morality. Individuals with a damaged prefrontal cortex engage in several kinds of antisocial, and even sociopathic behaviors. Many of the brain structures associated with culture are either unique to humans, or are best developed in them relative to close primate relatives.

The ideas are best conveyed by examples.

Example 14.2 (*Niche construction*) *Niche construction refers to the ability of an organism to influence the environment around it, and to epigenetically transmit this information from one generation to the next. As examples, Dawkins (2006) noted environmental constructs such as beaver dams, transmission of the social structure of bee hives across generations of bees, and species-specific hunting/mating practices of animals. An example of niche construction in humans is lactose tolerance (Wilson, 2012, ch. 20). In most of human evolutionary history, the production of the enzyme lactase that allows the digestion of milk, took place only among infants. However, with the development of herding, 9000 to 3000 years ago, a mutation that allowed for lactose tolerance in adults took place, and rapidly spread through the human population. In this example, the niche was the advent of domestication of cattle.*

Example 14.3 (*Cooperation and prosociality; Gintis, 2011*) *Unrelated or loosely related individuals in hunter-gatherer societies typically cooperated with each other in hunting, defence, and territorial acquisition. As Gintis (2011, p. 879) notes: “While some contemporary biological and economic theorists have attempted to show that such cooperation can be supported by self-regarding rational agents,¹⁹ the conditions under which their models work are implausible even for small groups.²⁰ Rather, the social environment of early humans was conducive to the development of prosocial traits, such as empathy, shame, pride, embarrassment, and reciprocity, without which social cooperation would be impossible.²¹” Such cooperative behavior is an epigenetic change that appears to have improved the fitness of the human species in that environment. This conclusion is supported by neuroscientific studies that identify particular brain regions that are responsible for these functions (such as empathy and prosociality).*

This view of human sociality is in contrast to the view that human altruism is a maladaptation in which humans mistakenly apply repeated game heuristics to static situations and/or mistake anonymous interactions for publicly observed interactions.²² There are two issues here. One is the evidence on the accuracy of this historical account. On this, Gintis (2011, p. 882) writes: “However, our understanding of contemporary hunter-gatherer societies indicates the lack of credibility of this argument, as do analyses of trade and migration patterns in Pleistocene.”²³ The second issue is the evidence from experimental games, which shows that human behavior in static and repeated games is quite different. Part 2 of the book explores this evidence in some detail. Furthermore, the evidence from the experiments of Andreoni

¹⁹ See, for instance, Alexander (1987), Fudenberg et al., (1994), and Trivers (1971).

²⁰ See, for instance, Boyd and Richerson (1988) and Gintis (2009b).

²¹ See, for instance, Sterelny (2011).

²² See, for instance, Haley and Fessler (2005), and Hagen and Hammerstein (2006).

²³ See, for instance, Bowles and Gintis (2011) and Fehr and Joseph (2004).

and Miller (2002) (see Part 2) shows that altruistic giving in dictator games does not arise from confusion; economic agents vary the amount of giving as the price varies.²⁴

Example 14.4 (*The physiology of communication; Gintis, 2009b, 2011*): In one of the most dramatic examples of gene–culture coevolution, the cultural development of oral communication, and the resulting fitness advantages, were supported by genetic changes that facilitated richer communication. Gintis (2011, pp. 380–1) outlines the relevant changes. These include: Expansion in the motor cortex region to facilitate speech production; development of brain regions such as Brocas and Wernicke’s areas, unique to humans, that facilitate grammatical speech; a gradual lowering of the larynx to increase the range and resonance of sounds; anatomical changes and developments in the tongue muscles that facilitate more precise speech; and changes in facial nerves and musculature that allowed a richer range of emotions to be displayed in a complementary manner with speech.

Example 14.5 (*The Westermarck effect; Wilson, 2012*) Humans instinctively avoid incest with others that they have been in contact with in early life, a phenomenon known as the Westermarck effect. Human children raised together up to the age of about 30 months avoid partnering to produce offspring. This tendency is shared with primate relatives such as chimpanzees and baboons. In humans, this tendency arises even if one is raised in early life with non-biological siblings. Wilson (2012, p. 202) writes: “it is evident that the human brain is programmed to follow a simple rule of thumb: Have no sexual interest in those whom you knew intimately during the earliest years of your life.” This is a cultural practice, and it is reinforced by taboos and social sanctions. Yet there are well-known genetic advantages from avoiding inbreeding. Hence, this is an example of gene–culture coevolution.

Example 14.6 (*Evolutionary political economy; Gintis et al., 2014*) Gintis et al. (2014) give a fascinating discussion of the gene–culture coevolution of dominance seeking, coalition formation, and the role of leadership in humans.²⁵ We encourage the interested reader to pursue the relevant references in their paper; what follows is a flavor of some of these fascinating issues.

The authors describe shared human evolutionary history with related primates, and demonstrate that humans have a genetic predisposition towards dominance seeking, and coalition formation. Most primate societies, except our closest relatives, the chimpanzees, do not engage in organized collective action, so they do not need leaders to coordinate their actions. Rather, in most primate societies, leadership is based on physical prowess, and the leader’s main role is to keep the peace between warring factions in the group.

Chimpanzees engage in two kinds of coalition building, rank changing and leveling. Rank changing coalitions arise when an individual takes over the leadership of a group, not necessarily because he is the strongest, but perhaps because he has a comparative advantage in coalition building. Leveling coalition formation occurs when males who are situated lower down in the chimp hierarchy, form a coalition to counter the resource appropriating tendency of the current leader. This is also known as reverse dominance hierarchy. For instance, one observes that female chimps in captivity make an attempt to neutralize “alpha-male bullies.” At least two factors played a role in the development of leadership in humans that was based not on physical prowess, but rather on the ability to motivate and persuade others. These

²⁴ However, the evidence from dictator games also shows that humans may be reluctant altruists and are less likely to give if they are guaranteed anonymity. We refer the reader to Part 2 for the details.

²⁵ Another lucid and fascinating source is Wilson (2012).

factors were: (1) The emergence and formation of leveling coalitions. (2) Humans developed lethal weapons so that there was a real threat to the life of the dominant male from other non-dominant males. By contrast, even in chimps, it is very difficult, without lethal weapons to kill the leader. Gintis et al. (2014, p. 24) give the determinants of the quality of persuasion among leaders: “Persuasion depends on clear logic, analytical abilities, a high degree of social cognition (knowing how to form coalitions and motivate others), and linguistic facility.” These leadership qualities (despite lacking superior physical prowess) are absent in other animals, although chimps seem to possess them in a more rudimentary form.

14.8.2 A formal model of cultural evolution

We now consider a formal model of cultural evolution (Gintis and Helbing, 2014).²⁶ Consider a single population of individuals, and a single cultural norm, A (short for *altruism*), that can be adopted, or not adopted. Suppose that at time t , the population is split into two categories. A fraction x_t adopts the norm A, and has fitness $1 - s$, $s \in (0, 1)$, while the remaining population fraction, $1 - x_t$, follows the alternative norm B (short for *selfishness*, or not A) and has fitness 1. Thus, the fitness of the selfish trait is higher than the fitness of the altruistic trait. The puzzle is to explain why, despite this fact, one observes the altruistic trait in humans.

In the initial period, there are n males and n females, so denote the total population size at time t by

$$P_t = 2n. \quad (14.69)$$

Males and females are randomly matched to produce offspring, so, depending on the categories of the matched pairs, these matching can be of the form AA, AB, BA, or BB.

Cultural transmission takes the following form.

1. In the *vertical channel* (parents to children), the offspring of AA parents are of type A, and the offspring of BB parents are type B. When parents are of mixed types (AB or BA), then half the offspring are type A, and half are type B.
2. In the *oblique channel*, a fraction $\gamma \in (0, 1)$ of type B children (of type AB, BA, and BB parents) are persuaded by *institutions*, or *other elders* who promote altruistic norms, to convert to type A. Anecdotal evidence supports the existence and strength of this channel. For instance, people typically encourage children to engage in moral, not immoral acts; to help others, not to steal from others; to tell the truth, and not lie. Institutions reward acts of bravery, not cowardice; the tax system gives tax deductions on charitable contributions, but imposes penalties for tax evasion. This channel is critical to the results below.
3. We also allow some players of type A to switch to type B in any period, on account of the higher fitness of the selfish trait (see below).

Recalling that the fractions of types A and B in the population are, respectively, x_t and $1 - x_t$, the number of distinct unions of each type are shown in Table 14.1.

²⁶ There is a rich literature on cultural transmission that typically uses a different transmission mechanism to the one we use here. For a model of cultural transmission motivated by an explanation of the inter-marriage rates and socialization among Protestants, Catholics, and Jews, see Bisin et al. (2004). For a survey of similar themes, see Bisin and Verdier (2011). Mengel (2008) considers cultural transmission of a norm of cooperation in a repeated prisoner's dilemma game in which not everybody adheres to a norm.

Table 14.1 Number of conjugal unions of each type.

	A	B
A	nx_t^2	$nx_t(1 - x_t)$
B	$nx_t(1 - x_t)$	$n(1 - x_t)^2$

Table 14.2 Number of offspring for each conjugal pair.

	A	B
A	$nx_t^2(1 - s)^2\beta$	$nx_t(1 - x_t)(1 - s)\beta$
B	$nx_t(1 - x_t)(1 - s)\beta$	$n(1 - x_t)^2\beta$

Each conjugal pair in Table 14.1 produces a new generation of offspring. Once a new generation is born at the end of the period, the previous generation dies. Let β be a parameter to be determined later, which ensures that the population grows at some rate $g(x_t)$ that depends on the fraction of altruists, x_t , at time t . Thus, the total population in the next time period, $t + 1$, equals

$$P_{t+1} = 2n(1 + g(x_t)).$$

In keeping with the idea of Darwinian fitness, the number of offspring of each conjugal pair in Table 14.1 is proportional to the product of the biological fitness of each pair of parents; see Table 14.2. For instance, the product of biological fitness of (i) AA parents is $(1 - s)^2$, (ii) AB or BA parents is $(1 - s)$, (iii) BB parents is 1. The multiplication by the parameter β allows one to capture the one-period population growth rate in each category.

The total number of offspring at time t , which equals the population at time $t + 1$ (because the previous generation dies off at the end of the period), is given by

$$P_{t+1} = nx_t^2(1 - s)^2\beta + 2nx_t(1 - x_t)(1 - s)\beta + n(1 - x_t)^2\beta. \quad (14.70)$$

The parameter β can now be found from the definition of the growth rate of population between periods t and $t + 1$

$$g(x_t) = \frac{P_{t+1} - P_t}{P_t}. \quad (14.71)$$

Substituting (14.69) and (14.70) into (14.71), we get

$$\beta = \frac{2(g(x_t) + 1)}{(1 - sx_t)^2}. \quad (14.72)$$

We can now find the fraction of offspring of each conjugal pair in Table 14.2. For instance, the proportion of offspring of AA parents is $\lambda_{AA} = \frac{nx_t^2(1-s)^2\beta}{P_{t+1}} = \frac{nx_t^2(1-s)^2\beta}{2n(1+g(x_t))}$. Substituting the value of β from (14.72), we get $\lambda_{AA} = \frac{x_t^2(1-s)^2}{(1-sx_t)^2}$. Proceeding in this manner, a simple calculation shows that the fractions of offspring of all types of conjugal pairs are given as follows.

$$\lambda_{AA} = \frac{x_t^2(1-s)^2}{(1-sx_t)^2}; \lambda_{AB} + \lambda_{BA} = \frac{2x_t(1-x_t)(1-s)}{(1-sx_t)^2}; \lambda_{BB} = \frac{(1-x_t)^2}{(1-sx_t)^2}. \quad (14.73)$$

We now wish to determine the fraction of altruists at time $t + 1$. We do this in a series of steps.

1. The vertical transmission channel ensures that the fraction λ_{AA} is comprised of all altruists because both parents are altruistic.
2. Consider now the offsprings of parents of types AB and BA. The fraction $\frac{1}{2}(\lambda_{AB} + \lambda_{BA})$ of offspring is altruistic by birth (vertical channel) and the remaining fraction, $\frac{1}{2}(\lambda_{AB} + \lambda_{BA})$ is selfish by birth (vertical channel). However, of the latter fraction that is selfish by birth, a fraction $(\gamma x_t) \frac{1}{2}(\lambda_{AB} + \lambda_{BA})$ is converted to the altruistic type on account of the oblique channel, where $\gamma \in (0, 1)$ is time invariant. Thus, the greater the fraction of altruists, x_t , in the population, the higher is the rate of conversion.
3. Finally, the fraction of individuals λ_{BB} are all selfish because both parents are selfish (vertical channel). Of this, a fraction $(\gamma x_t) \lambda_{BB}$ are converted to altruistic types on account of the oblique channel; the rate of conversion depends on the fraction of altruists, x_t , in the population.

Thus, the change in the fraction of altruists between time t and $t + 1$ is

$$x_{t+1} - x_t = \left[\lambda_{AA} + \frac{1 + \gamma x_t}{2} (\lambda_{AB} + \lambda_{BA}) + \gamma x_t \lambda_{BB} \right] - x_t. \quad (14.74)$$

Substituting (14.73) in (14.74), we get

$$x_{t+1} - x_t = \frac{x_t(1-x_t)(\gamma-s)}{(1-sx_t)}. \quad (14.75)$$

While it has been pedagogically simpler to outline the model in discrete time, it is easier to work in continuous time once the dynamic equation has been derived. The continuous time analog of (14.75) is given by the first order, non-linear, differential equation.

$$\dot{x}_t = \frac{x_t(1-x_t)(\gamma-s)}{(1-sx_t)}. \quad (14.76)$$

We now consider our second transmission channel. The fitness level of the selfish types is greater than the altruistic types. So should not an individual of type A wish to imitate the selfish trait of an individual of type B? To take account of this possibility, suppose that each individual i randomly samples the fitness of another member j , and switches to j 's type if the fitness of j is higher. Thus, only altruistic types switch to selfish types because their fitness is relatively lower ($1 - s < 1$).

There might be *informational* or *observational frictions* about the payoffs of others. Thus, the probability that any member of type A switches to type B, denoted by p , is directly proportion to the fitness difference between the two types, so $p = \sigma s$, where $\sigma \in (0, 1)$. The leakage of some A types to B types, through this channel, reduces the growth rate of altruistic types. In a large population, the probability that a type A player randomly meets a type B player is $x_t(1-x_t)$. Such a meeting results in a conversion of type A to B with probability p . Hence, once this imitation channel works itself out, the change in the altruistic population is given by a modification of (14.76)

$$\dot{x}_t = \frac{x_t(1-x_t)(\gamma-s)}{(1-sx_t)} - \sigma sx_t(1-x_t). \quad (14.77)$$

The two terms on the RHS in (14.77) represent two channels through which the altruistic types evolve over time. The first is the *cultural transmission mechanism* (vertical and oblique) that favors the altruistic norm. The second is the *standard replicator dynamics* that favors the relatively fitter cultural trait of selfishness. Let us rewrite (14.77) in a manner that makes the steady states fairly transparent.

$$\dot{x}_t = f(x_t) = x_t(1-x_t) \left[\frac{(\gamma-s)}{(1-sx_t)} - \sigma s \right]. \quad (14.78)$$

Proposition 14.12 Consider the first order non-linear differential equation given in (14.78).

(a) The three steady states are

$$\bar{x} = 0, \bar{x} = 1, \text{ and } \bar{x} = \frac{1}{s^2\sigma} (s + s\sigma - \gamma). \quad (14.79)$$

(b) For $s < s_c = \frac{\gamma}{1+\sigma}$, the unique stable steady state is $\bar{x} = 1$ where all individuals are altruistic.

(c) For intermediate levels of s : $s_c < s < s_-$, where s_- is defined in (14.80)

$$s_- = \frac{(1+\sigma) - \sqrt{(1+\sigma)^2 - 4\sigma\gamma}}{2\sigma}, \quad (14.80)$$

both $\bar{x} = 0$, and $\bar{x} = 1$ (all play A and all play B) are locally stable. In this case, $\bar{x} = \frac{1}{s^2\sigma} (s + s\sigma - \gamma)$ is an unstable equilibrium that separates the basins of attraction of the two stable equilibria.

(d) If $s_- < s$, then $\bar{x} = 0$ is the unique stable steady state.

Proof: (a) Set $f(x_t) = 0$ in (14.78) and solve out for the three roots to get (14.79).

(b) Differentiating $f(x_t)$ in (14.78), we get

$$f'(x_t) = \frac{-1}{(sx-1)^2} (s(1+\sigma)-\gamma + 2x(\gamma-s(1+\sigma+s\sigma))+sx^2(s+4s\sigma-\gamma+s^2\sigma)-2s^3x^3\sigma).$$

Let us evaluate f' at the three steady states in (14.79).

$$f'(0) = \gamma - s(1+\sigma). \quad (14.81)$$

$$f'(1) = \frac{-1}{(1-s)} (s^2\sigma - s(1+\sigma) + \gamma). \quad (14.82)$$

$$f'\left(\frac{1}{s^2\sigma} (s(1+\sigma) - \gamma)\right) = \frac{-1}{s(s-\gamma)} (\gamma - s(1+\sigma)) (\gamma - s(1+\sigma) + s^2\sigma). \quad (14.83)$$

We now sequentially check the stability condition, $f'(\bar{x}) < 0$, for the three steady states.

(bi) Steady state $\bar{x} = 0$: First note that

$$\gamma - s(1 + \sigma) \begin{matrix} \geq \\ \leq \end{matrix} 0 \text{ if } s_c = \frac{\gamma}{1 + \sigma} \begin{matrix} \geq \\ \leq \end{matrix} s. \quad (14.84)$$

Using (14.81) and (14.84), we get

$$\begin{cases} s < s_c & \Rightarrow \bar{x} = 0 \text{ is unstable.} \\ s > s_c & \Rightarrow \bar{x} = 0 \text{ is stable.} \end{cases}$$

(bii) Steady state $\bar{x} = 1$: The roots of the quadratic equation in s

$$s^2\sigma - s(1 + \sigma) + \gamma = 0, \quad (14.85)$$

are given by $s = \frac{(1+\sigma) \pm \sqrt{(1+\sigma)^2 - 4\sigma\gamma}}{2\sigma}$. Since $\sigma, \gamma \in (0, 1)$, we have $(1 + \sigma)^2 - 4\sigma\gamma > 0$, thus, both roots are real. Denote the smaller root by s_- ; this is given in (14.80). The larger root $s_+ = \frac{(1+\sigma) + \sqrt{(1+\sigma)^2 - 4\sigma\gamma}}{2\sigma} > 1$ is not feasible because $s \in (0, 1)$, so we shall ignore it. One can show that $s_- > s_c$. Since, the coefficient of s^2 in (14.85), $\sigma > 0$, so for any $s < s_c$ we have $s^2\sigma - s(1 + \sigma) + \gamma > 0$. Using (14.82) we have $f'(1) < 0$. Thus, $\bar{x} = 1$ is a stable steady state for $s < s_c$.

(biii) Steady state $\bar{x} = \frac{1}{s^2\sigma}(s + s\sigma - \gamma)$. One can check that in this case, for any $s < s_c$, we have $\bar{x} < 0$, which is not feasible.

Putting together (bi)–(biii), the only stable steady state when $s < s_c$ is $\bar{x} = 1$.

(c) Simple calculations show in this case that $f'(0) < 0$, and $f'(1) < 0$, so the intermediate steady state must satisfy $f'\left(\frac{1}{s^2\sigma}(s(1 + \sigma) - \gamma)\right) > 1$. Thus, the two extreme steady states are locally stable, while the intermediate steady state is unstable, and it separates the basins of attraction of the two stable equilibria.

(d) In this case $\bar{x} = \frac{1}{s^2\sigma}(s + s\sigma - \gamma) > 1$ is not feasible, and $f'(0) < 0$, but $f'(1) > 0$. Thus, $\bar{x} = 0$ is the only stable steady state. ■

Proposition 14.13 (*Fundamental Theorem of Sociology; Gintis and Helbing, 2014*) *In the model described above, altruistic norms persist in a cultural equilibrium only if there is a strictly positive rate of cultural transmission of altruism via social institutions, i.e., $\gamma > 0$.*

Proof: Suppose that there is an absence of cultural transmission of altruism through the oblique channel, so $\gamma = 0$. Substitute $\gamma = 0$ in (14.81), (14.82), and (14.83), and we have $f'(0) < 0$, but $f'(1) > 0$, and $f'\left(\frac{1}{s^2\sigma}(s(1 + \sigma))\right) > 0$. Thus, $\bar{x} = 0$ (all play B) is the unique stable steady state in which all individuals are selfish. ■

Proposition 14.12 reflects the opposing forces of altruistic cultural transmission, and the standard replicator dynamics that favors selfishness because it is relatively fitter. When the fitness advantages of the selfish trait are weak in the sense of low enough s , then everyone is altruistic (Proposition 14.12(b)). At the other extreme, when s is high enough, then the effect arising through replicator dynamics is so strong that cultural transmission is unable to have any effect in the stable steady state (Proposition 14.12(d)). However, in the absence of any cultural transmission $\gamma = 0$, everyone is selfish (Proposition 14.13). Thus, we have the important result that *cultural transmission is a necessary, but not sufficient condition for altruism to emerge in society.*

In this book we have, in many instances, reported the effect of cultural factors on empirical results. For instance, there are quite dramatic cultural effects in the ultimatum game results when it is played in diverse cultures (Henrich et al., 2001). The formal model of cultural transmission in this section is admittedly simple in that it treats two polar extremes of altruism and

selfishness. By modifying this model to consider a range of altruistic strategies, say, $A_1, \dots, A_m$, and allowing for a range of parameter values, one may be able to explain the cultural richness of the experimental results arising from various games. This model can also be used to underpin the emergence of social norms (see Part 4 of the book and Chapter 16) and various forms of other-regarding preferences and character virtues (see Part 2 of the book).

14.9 The evolution of altruism and reciprocity

Bowles and Gintis (2011) provide an authoritative evolutionary treatment of human cooperation. Their thesis rests on three main arguments, all backed by good evidence. First, humans have evolved institutions and cultural practices to (i) protect altruistic members from being exploited by selfish members, and (ii) to promote altruism. In addition to the role of leveling coalitions that we mention above, other methods include ostracism and punishment of deviants and free-riders who violate the cooperative norm, tax deductibility of charitable giving, awards for altruism and other character virtues. Second humans have internalized the norm of cooperation; the mechanism is similar to the evolutionary stability of the culturally transmitted altruistic strategy in Section 14.8.2 above. Third, competition for scarce resources favors groups who are relatively better at coordinating the acquisition of resources, e.g., territorial conquests.

Now we examine the evolutionary origins of cooperation, using a taxonomy suggested by Nowak (2006). Many other areas of behavioral economics are not as well underpinned by evolutionary foundations, hence, we omit a discussion of these. However, for non-exponential time preferences, see Sozou (1998) and Dasgupta and Maskin (2005), although only the latter can explain hyperbolic discounting. For a nice discussion of these issues, see Robson and Samuelson (2011).

The usage of the terms direct reciprocity, indirect reciprocity, and reciprocal altruism in this section follows the standard usage in biology. This usage is sometimes different from that in Part 2 in order to maintain coherence with the relevant literature.

14.9.1 *Kin selection and inclusive fitness*

The idea behind *inclusive fitness*, introduced by Hamilton (1963, 1964), comes from observations on the behavior of animals in which they provide costly assistance to other animals of the same species. For instance, a member of Belding's ground squirrels will typically emit a cry to warn other animals in the group about the presence of a predator, but in the process, put itself into extreme danger by identifying itself to the predator. Maynard Smith (1964) referred to this idea as *kin selection*, which is a narrower concept than inclusive fitness because it refers only to interaction between genealogical relatives. We may take inclusive fitness as a broader concept that deals with the observed prosocial behavior among humans, even towards those who are not related by kin. Sometimes in common usage, the related ideas of kin selection and inclusive fitness are used interchangeably.

Suppose that an individual confers an altruistic benefit b to another individual at some personal cost c , and the *degree of relatedness* (or probability of sharing a gene) of the other individual is r . For instance, $r = 1/2$ for siblings, $r = 1/4$ for nieces/nephews, $r = 1/8$ for cousins, and so on. Benefits b and cost c are measured in units of fitness. Individuals may be willing to engage in altruistic acts towards relatives and this willingness may be increasing in the degree of relatedness. This is nicely captured in a quote attributed to J.B.S. Haldane: "I will jump into the river to save two

	A	B
A	$b-c, b-c$	$-c, b$
B	$b, -c$	$0, 0$

Figure 14.12 A prisoner's dilemma game.

	A	B
A	$(b-c)(1+r), (b-c)(1+r)$	$br-c, b-rc$
B	$b-rc, br-c$	$0, 0$

Figure 14.13 Prisoner's dilemma game, modified to take account of inclusive fitness.

Source: From Nowak, M. A. (2006). "Five rules for the evolution of cooperation." *Science* 314(5805): 1560–3. Copyright © 2006, American Association for the Advancement of Science. Reprinted with permission from AAAS. Readers may view, browse, and/or download material for temporary copying purposes only, provided these uses are for noncommercial personal purposes. Except as provided by law, this material may not be further reproduced, distributed, transmitted, modified, adapted, performed, displayed, published, or sold in whole or in part, without prior written permission from the publisher.

brothers, or eight cousins." In a seminal contribution, Hamilton (1964) proposed the following inequality for altruism to be favored by natural selection

$$rb > c \Leftrightarrow r > \frac{c}{b}. \quad (14.86)$$

Rewriting (14.86) as $\frac{b}{c} > \frac{1}{r}$, the required ratio of benefits to cost can be quite high as the degree of relatedness decreases. Gintis (2014) shows that Hamilton's original inequality is consistent with both prosocial and antisocial behavior, and proposes important amendments to this inequality.

Consider the one-shot prisoner's dilemma game shown in Figure 14.12. Suppose that there are two traits in the population, A (altruism) and B (selfishness). Suppose that there are random binary meetings between members of the population. An individual with trait A confers a benefit b to the other in the binary meeting, at a cost c , such that $b > c > 0$. However, individuals with trait B neither incur the cost nor impart benefits to others.

In the presence of inclusive fitness, we may modify the strategic form game in Figure 14.12 as follows. Suppose that the average degree of relatedness in the population is $r \in (0, 1)$. The main implication of inclusive fitness is that one internalizes the payoffs of relatives (benefits and costs) in a proportion that depends on r . This is reflected in the modified strategic normal form game shown in Figure 14.13.

We leave the proof of the next proposition as an exercise. It shows that Hamilton's inequality in (14.86) has two desirable properties.

Proposition 14.14 (Nowak, 2006): *Consider the strategic normal form game shown in Figure 14.13. Cooperation (trait A) is an evolutionary stable strategy as well as a risk dominant strategy if (14.86) holds.*²⁷

²⁷ See the Appendix on game theory for the definition of a risk dominant strategy.

14.9.2 Group selection or multilevel selection

Hamilton's inequality shows that the explanation of prosocial behavior towards non-kin members (very low r) requires an unrealistically high benefit to cost ratio of prosocial behavior. In the quest to explain cooperation, one is then led to the idea of *group selection*, sometimes known as *multilevel selection*. The basic idea here is that over human evolutionary history (Wilson, 2012): (i) people have lived in small groups, (ii) there was inter-group competition for scarce resources, and (iii) cooperation among more prosocial groups results in higher success in inter-group competition.

Within each group, certain prosocial traits such as altruism/cooperation may have lower fitness relative to selfishness. So if each group lived in isolation, prosocial traits may never have been evolutionarily stable. However, in the struggle between groups, groups with more cooperative traits are likely to prevail over others in the competition for scarce resources. This imparts higher group-level fitness to such prosocial traits. Whether prosociality survives, depends on the relative strengths of these two opposing effects. Haldane (1932) proposed an initial transmission mechanism for how group selection may operate. Price (1972) gave a more formal account of such a mechanism. We give a flavor of the argument; the interested reader can consult Bowles and Gintis (2011) for the details.

Consider the prisoner's dilemma game shown in Figure 14.12. The population is divided into several groups. Denote the fraction of individuals in group j by λ_j . Let i index an individual in group j and let the expected fitness of this individual be denoted by π_{ij} . Let d_{ij} be a dummy variable that equals 1 if the individual has trait A, otherwise $d_{ij} = 0$. Denote the fraction of individuals in the entire population with trait A at time t by x_t and the fraction with trait A in group j by x_{jt} . Suppose that fitness takes the following form.

$$\pi_{ij} = \beta_0 + \beta_g x_{jt} + \beta_i d_{ij}. \quad (14.87)$$

In (14.87), fitness, π_{ij} , depends, in an additively separable manner, on the fraction of other altruists in the same group, x_{jt} , and whether the individual is altruistic or not, d_{ij} . β_g and β_i measure, respectively, the marginal effects of these variables and β_0 considers the effect of omitted variables. The *Price equation* can then be shown to take the following form²⁸

$$\Delta x_t = x_{t+1} - x_t = (b - c) \text{var}(x_{jt}) - c \overline{\text{var}}(d_{ij}), \quad (14.88)$$

where Δx_t is the change in the fraction of altruists in the population, $\text{var}(x_{jt})$ is the between-groups genetic variance of the trait A, and $\overline{\text{var}}(d_{ij}) = \sum_j \lambda_j \text{var}(d_{ij})$ is the average within-groups variation of the trait A. A higher variation between-groups suggests that there is greater sorting of groups into high A and low A types. In high type A groups, the fitness of the altruistic strategy is high because altruistic types are more likely to run into other altruistic types, hence, the payoff, $b - c$ (see Figure 14.12), is high; this is reflected in the higher reproductive success of these altruistic types. In binary meetings across groups, however, the altruistic trait has a lower fitness compared to the selfish trait (second term in (14.88)).

Let $P(A | A)$ and $P(A | B)$ denote the probability that individuals with traits A and B, respectively, encounter an individual with trait A. For the prisoner's dilemma game in Figure 14.12, Bowles and Gintis (2011) use this framework to derive a critical condition for the steady state

²⁸ For a derivation, see Appendix A10 in Bowles and Gintis (2011).

$$\Delta x_t \begin{matrix} \geq \\ \leq \end{matrix} 0 \Leftrightarrow P(A | A) - P(A | B) \begin{matrix} \geq \\ \leq \end{matrix} \frac{c}{b}. \quad (14.89)$$

Notice the similarity of (14.89) with Hamilton's rule in (14.86) if we interpret $P(A | A) - P(A | B) = r$. Thus, the fraction of altruists is expected to increase if the probability that altruists meet other altruists is high, the cost of altruism is low, and the benefits to the other party of altruism are high.

Readers may wonder if these sorts of models pass the empirical test of describing the long sweep of human evolutionary history. We leave the reader to consult Chapter 6 in Bowles and Gintis (2011) for the details. They find that, on average, r is unlikely to exceed 0.10. Thus, models of kin-based selection (see (14.86)) require b to exceed c by a factor of about 10, which is unlikely. However, for individual actions on which the survival of the group may be at stake, as could be the case in models of group or multilevel selection, it is possible that such a relationship holds. Hence, the claim of the group selection model is that it describes better the emergence of altruism.

More generally, group selection can also explain the development of human morality. While selfish traits enhance individual fitness, group fitness may be enhanced by moral actions toward others. Furthermore, unless one is a sociopath, the human mind appears predisposed to accept and imitate moral actions.

Nowak et al. (2010) argued that group selection is a more persuasive explanation for human evolution as compared to kin selection. Wilson (2012) gives a non-technical and wide-ranging account of this position. Kin selection or inclusive fitness was in turn defended against group selection by Abbot et al. (2011) in a paper that has 136 authors. Hence, the idea of group selection is still viewed with great suspicion by many evolutionary biologists. In the opinion pages of the *New York Times* on June 24, 2012, one of the main proponents of group selection, Edward O. Wilson, wrote:

But at no time have our mathematical and empirical arguments been refuted or even seriously challenged. Since that protest, the number of supporters of the multilevel selection approach has grown, to the extent that a similarly long list of signatories could be obtained. But such exercises are futile: science is not advanced by polling. If it were, we would still be releasing phlogiston to burn logs and navigating the sky with geocentric maps.

There is no consensus of views on these two opposing theories at this time.

14.9.3 Reciprocal altruism or direct reciprocity

The tit-for-tat strategy (see Section 14.6.1, above) is an exceedingly simple form of *reciprocal altruism*, sometimes referred to as *direct reciprocity* in the literature (Trivers, 1971). Following Bowles and Gintis (2011), consider the tit-for-tat strategy, T , for the infinitely repeated prisoner's dilemma game with the stage game given in Figure 14.12. We assume that individuals who use the tit-for-tat strategy, initially play the strategy A , and then in any subsequent period $t > 1$, copy the strategy of the opponent in period $t - 1$. This is sometimes known as the nice tit-for-tat because players initially play A .

The tit-for-tat strategy is a form of *conditional cooperation*, and we may refer to it as such. Suppose that the one-period discount factor is $\delta \in [0, 1]$ and players use exponential discounting. We are interested in checking if T is an ESS.

ESS is a subset of the set of Nash equilibria, and the two concepts coincide when the Nash equilibrium is strict. Thus, let us find the conditions for “all play T ” to be a strict Nash equilibrium. Conditional on all others playing their part of the Nash equilibrium, an individual prefers the strategy T over the strategy “defect” (i.e., play B every period) if²⁹

$$(b - c)[1 + \delta + \delta^2 + \dots] > b,$$

or $\delta > \frac{c}{b}$, which is similar to Hamilton’s inequality in (14.86) once the discount factor replaces the degree of relatedness, r .

Reciprocal altruism is, however, difficult to find in most animals, except close primate relatives because most animals have quite high discount rates, hence, they exhibit very low values of δ (Hammerstein, 2003). However, reciprocal altruism appears to play a role in human interaction. An example is the greater observed cooperation in the partner-treatment, as compared to the stranger-treatment in public goods games experiments (see Part 2 of the book). Indeed nothing in our derivation restricts such altruism to kin members only. In computer tournaments, Axelrod (1984) showed that the tit-for-tat was the winning strategy in the prisoner’s dilemma game shown in Figure 14.12.

Conditional cooperation of this sort can be an ESS under certain conditions, but not others. These conditions have to do with the size of groups, the number of interacting individuals, and behavioral errors. For instance, in larger groups, it will become more difficult to keep track of previous defectors, and for a variety of behavioral biases, individuals might not have perfect recall. A more detailed analysis of these issues can be found in Bowles and Gintis (2011) who also show that the empirical evidence for these sorts of models (e.g., reciprocal altruism) is not overwhelming. For these reasons, they advocate models of gene-culture coevolution as providing a more convincing explanation.

Suppose that there are behavioral errors, say, for whatever reason, players choose their strategies with some trembles. This increases the cost of playing T because one round of unintended defection by the opponent, induces endless defection under T , reducing the players’ payoffs; see, for instance, Selten and Hammerstein (1984).

This has led to a search for alternative strategies in repeated games that have been shown to perform better than T . These include: (i) the *generous tit-for-tat* in which an individual cooperates with some probability even if the opponent has defected in the previous period, and (ii) the *win-stay, lose-shift strategy* in which the previous period strategy is repeated if one is doing well (say, if both players have played A) and the previous period strategy is changed if one is not doing well (say, if the opponent has played B).³⁰

Colman et al. (2010) explore the possibility of learning under a noisy win-stay, lose-shift strategy when players have limited information on the strategic environment, as in models of reinforcement learning. They find that adaptive learning is very slow and players typically do not show awareness that they are participating in a strategic environment. Even relatively high incentives do not improve the speed of learning. Learning is slower in multiplayer, relative to dyadic, interactions, even when players are informed that they are engaged in strategic interaction.

Consider the one-period prisoner’s dilemma game shown in Figure 14.12. We wish to find how the payoff matrix is altered when we compare the strategy T against the alternative strategy that

²⁹ Payoffs are unaltered if one switches to A conditional on all others playing T .

³⁰ For a more detailed account of these strategies and their superiority to strategy T , see Fudenberg and Maskin (1990) and Nowak and Sigmund (1992, 1993).

	T	D
T	$\frac{(b-c)}{(1-p)}, \frac{(b-c)}{(1-p)}$	$-c, b$
D	$b, -c$	$0, 0$

Figure 14.14 Prisoner's dilemma game, modified to take account of direct reciprocity.

Source: From Nowak, M. A. (2006). "Five rules for the evolution of cooperation." *Science* 314(5805): 1560–3. Copyright © 2006, American Association for the Advancement of Science. Reprinted with permission from AAAS. Readers may view, browse, and/or download material for temporary copying purposes only, provided these uses are for noncommercial personal purposes. Except as provided by law, this material may not be further reproduced, distributed, transmitted, modified, adapted, performed, displayed, published, or sold in whole or in part, without prior written permission from the publisher.

plays B every period; call this the strategy, D , for defect. Suppose that the probability with which the same players meet in the next round is p . Thus, the expected number of rounds played between the two players is $1/(1-p)$. Thus, when one T player meets another T player, the expected payoff is $(b-c) + p(b-c) + p^2(b-c) + \dots = \frac{b-c}{1-p}$. When a T player meets a D player, then the current payoff is $-c$, followed by a payoff of zero in every future period. If a D player meets a T player, then he gets b in the current period, and zero thereafter, for a total intertemporal payoff of b . These payoffs are summarized in the payoff matrix in Figure 14.14.

The next proposition considers the evolutionary stability of strategy T against the strategy D ; we omit the proof and the interpretation and leave these as an exercise.

Proposition 14.15 (Nowak, 2006): Consider the strategic normal form game in Figure 14.14. The strategy T is an ESS if $p > \frac{c}{b}$ and it is risk dominant if $\frac{b}{c} > \frac{2-p}{p}$.

14.9.4 Indirect reciprocity

Direct reciprocity requires that individuals engage in repeated interactions with each other. In this case, individuals use *directly observed and experienced information* about past play to decide on the degree of cooperation. But we also do observe altruism and spite that is directed at complete strangers, where we have not had a chance to directly observe past play. Cooperative behavior with strangers and non-strangers is likely to be observed by others. Information on one's *reputation* and behavior, cooperative or non-cooperative, is then likely to spread among others through means such as *gossip* (Alexander, 1987; Nowak and Sigmund, 1998). The evolution of language must have played a critical role in this process. This gives rise to *indirect reciprocity*, i.e., conditioning one's current play on *indirectly observed past play of opponents*. Consistent with this view, evidence suggests that more helpful individuals are helped more in the future (Panchanathan and Boyd, 2004; Nowak and Sigmund, 2005). We may use indirect reciprocity not only to model the spread of altruism and cooperation, but also other social norms and even morality; for the implications for morality see Alexander (1987); and for social norms see Brandt and Sigmund (2004).

Let us consider indirect reciprocity when the stage game is the prisoner's dilemma game given in Figure 14.12. We need to specify two elements. The first element is an *action rule*, i.e., choose A or B depending on the *reputation* or *image score* of the other player. At the simplest level, the image score could be binary: 1 if the individual has a good reputation and 0 if the reputation is bad. The second element is an *assessment norm*, i.e., how does a third party assign image scores to others.

	C	D
C	$(b - c), (b - c)$	$-c(1 - q), b(1 - q)$
D	$b(1 - q), -c(1 - q)$	$0, 0$

Figure 14.15 Indirect reciprocity and the prisoner's dilemma game.

Source: From Nowak, M. A. (2006). "Five rules for the evolution of cooperation." *Science* 314(5805): 1560–3. Copyright © 2006, American Association for the Advancement of Science. Reprinted with permission from AAAS. Readers may view, browse, and/or download material for temporary copying purposes only, provided these uses are for noncommercial personal purposes. Except as provided by law, this material may not be further reproduced, distributed, transmitted, modified, adapted, performed, displayed, published, or sold in whole or in part, without prior written permission from the publisher.

In the context of the prisoner's dilemma game, consider two strategies. The temporal strategy, defect in every period, D , involves the play of B in the stage game, every period. The temporal strategy, C , denotes the cooperative strategy in which players play A unless they know that the reputation of the opponent is bad (corresponding to D). In the prisoner's dilemma game in Figure 14.12, players know the strategies of the opponents due to, say, observations on past reputations; and players are hardwired to play their strategies. Players who play C always cooperate when they meet other players who play C . However, when they meet a player with reputation to play D , they cooperate (i.e., confer a benefit b at a cost c) with probability $1 - q$, $q > 0$. These assumptions modify the original game to the one shown in Figure 14.15, which can now be considered to be a strategic normal form game.

The proof of the next proposition is left as an exercise for the reader. It shows that a form of Hamilton's rule is required to ensure the ESS of the cooperative strategy.

Proposition 14.16 (Nowak, 2006): Consider the strategic normal form game in Figure 14.15. The strategy C is (i) an ESS if $q > \frac{c}{b}$, and (ii) a risk dominant strategy if $\frac{q}{2-q} > \frac{c}{b}$.

14.9.5 Repeated games, cooperation, and classical game theory

Readers trained in classical game theory may believe that altruism can be explained using the theory of repeated games and self-regarding preferences. Gintis (2009b) and Bowles and Gintis (2011) strongly reject such an explanation and term Nash equilibria of repeated games as *evolutionary irrelevant equilibria*. We have already explored these arguments in Part 4 of the book. Here we simply note the salient problems with Nash equilibria in repeated games.

First, the folk theorem is an existence theorem that does not provide any satisfactory dynamic process for the adjustment to equilibrium. Second, repeated games typically involve public signals, i.e., all players receive the same signal, although the signal may be perfect or imperfect (e.g., defection of a particular player may be observed with some error). However, when players have private signals, the most relevant real-life case, then the existing results only apply to models where we have near-public signals. Third, there is no mechanism for coordinating on trigger strategies, particularly in large groups. Fourth, the equilibrium in repeated games, particularly with private signals, involves the play of mixed strategies. Since players are indifferent between the pure strategies in the support of the equilibrium mixed strategy, they have no incentive to play the pure strategies in the equilibrium proportions. This problem with mixed strategies is addressed in a simple class of games, through the purification agenda, but this class excludes repeated games.

14.9.6 The indirect evolutionary approach

In the *indirect evolutionary approach*, individuals are endowed with evolving preferences and they maximize their utility in each time period; see Güth and Yaari (1992) and Güth (1995). Over time, preferences evolve, depending on some underlying revision protocol, but conditional on the preferences, individuals continue to maximize their utility whenever they are called upon to make a choice.

THE BASIC CONCEPTS

Consider the following example that is taken from Robson and Samuelson (2011). The two-player strategic normal form game in Figure 14.16 has a unique Nash equilibrium (T, R) .

Suppose that there are two populations of players. The respective strategy sets in populations 1 and 2 are $S_1 = \{T, B\}$ and $S_2 = \{L, R\}$. In population 1, strategy T dominates strategy B . Players from the two populations are randomly matched to play the game in Figure 14.16. Players are hardwired with preferences for particular strategies. An important assumption is that players know the preferences of the opponent. The motivation is that humans possess subtle cues to recognize traits in others; often these traits are similar to their own, which is the *green-beard effect*. The material payoffs in Figure 14.16 are also the relevant fitness levels that determine reproductive success.

From the point of view of population 1, it is ideal if a player could pre-commit to playing the dominated strategy B . In standard game theory, this is ruled out because, in a static game, players cannot make such commitments. However, in the indirect evolutionary approach, credible commitments arise through hardwired strategies or preferences.

It might well be that some players in population 1 have the preference to play strategy T and others have the preference to play B . Since player 2 (from population 2) is assumed to observe these preferences, he will play R (respectively, L) when matched against an opponent who prefers to play T (respectively, B). Thus, in population 1, players who play B will get a higher material payoff, and achieve higher reproductive fitness, relative to players who play T . Thus, over time, as preferences evolve, players in population 1 who possess the preference to play B will overwhelm those who play T . In other words, B will eventually be an ESS, but T fails to be an ESS.

This framework has been used to provide a justification for the existence of departures from the findings of the standard textbook economic model. These departures include punishments in public goods games, altruism, and the endowment effect (Ostrum, 2000). However, there are several limitations of the indirect evolutionary approach. In the game in Figure 14.16, population 2 would likewise benefit by evolving a hardwired preference for playing R . Since such a preference will be observed by player 1, his best response is T , so that player 2 gets 4. It is not clear what outcome one would get by allowing a richer array of such preferences.

	L	R
T	6, 2	4, 4
B	5, 1	2, 0

Figure 14.16 Fitness levels from each strategy in two different populations.

Source: Reprinted from *Handbook of Social Economics*, Behabib, J., Bisin, A., and Jackson, O. (eds.), "The evolutionary foundations of preferences," Robson, A. and Samuelson, L. Vol. 1, pp. 221–310. Copyright © 2011, with permission from Elsevier.

The assumption that the preferences of others can be observed is a strong one; however, this has been relaxed. When preferences are unobservable, then the steady state in the indirect evolutionary approach must be a Nash equilibrium (Ely and Yilankaya, 2001; Ok and Vega-Redondo, 2001). But then, in Figure 14.16, we are back to the outcome (T, R) under classical game theory. Güth (1995), however, shows that the indirect evolutionary approach can have important implications when players have some, but imperfect, information about the preferences of others that allows periodic observation of preferences.

Robson and Samuelson (2011) argue that the generation of cues that makes preferences recognizable to others, ought to be endogenous to an evolutionary process. But a weakness of the indirect evolutionary approach is that these cues are exogenous. They write: “In our view, the indirect evolutionary approach will remain incomplete until the evolution of preferences, the evolution of signals about preferences, and the evolution of reactions to these signals, are all analyzed within the model.” This requirement might be too challenging, not just for this theory, but for most of game theory.

THE EVOLUTION OF MORALITY

Alger and Weibull (2013) consider the *evolution of morality* within the indirect evolutionary approach. Aside from the intrinsic worth of these results, the framework allows us to define an ESS in the presence of asymmetric information (Definition 14.10 below), which allows readers to extend their knowledge of the basic concepts.

Suppose that we have a strategy set X that is shared by players from two different populations. Individuals are randomly matched across the populations to engage in bilateral interactions, but matching can be *assortative*. So, for instance, players who have identical types may have more than an even chance of being matched (perhaps on account of some underlying green-beard effect). The material payoff that is relevant for reproductive fitness when an individual plays a strategy $x \in X$ against an opponent who plays the strategy $y \in X$ is $\pi(x, y)$, so $\pi : X \times X \rightarrow \mathbb{R}$.

In each population, there are two types of preferences, θ and τ , in the proportions $1 - \epsilon$ and ϵ ; these types are drawn from the set of types Θ . Thus, a population state is summarized by $s = (\theta, \tau, \epsilon)$. For any state, $P(\tau | \theta, \epsilon)$ gives the conditional probability that an individual of type θ is matched with a type τ individual; $P(\theta | \tau, \epsilon)$ is defined analogously. Each of the types, θ, τ , pick strategies from the set X . The *types* of players are hardwired, but they simultaneously choose their *strategies* in order to maximize their expected utility.

For any $x, y \in X$, such that type θ plays x and type τ plays y , the expected payoff or fitness of the types θ and τ is given, respectively, by

$$\Pi_{\theta}(x, y, \epsilon) = P(\theta | \theta, \epsilon) \pi(x, x) + P(\tau | \theta, \epsilon) \pi(x, y), \quad (14.90)$$

$$\Pi_{\tau}(x, y, \epsilon) = P(\theta | \tau, \epsilon) \pi(y, x) + P(\tau | \tau, \epsilon) \pi(y, y). \quad (14.91)$$

In (14.90), with probability $P(\theta | \theta, \epsilon)$, a type θ player who plays x meets a type θ player from the other population who plays x , for a payoff of $\pi(x, x)$. With the complementary probability $P(\tau | \theta, \epsilon)$, the opponent is a type τ player from the other population who plays x , giving a payoff of $\pi(x, y)$. $\Pi_{\tau}(x, y, \epsilon)$ in (14.91) can be constructed analogously.

Strategy choices are made simultaneously and types are private information, so the relevant equilibrium concept in a static asymmetric information game is a Bayesian Nash equilibrium (BNE). Suppose that payoffs, π , and utilities, u , of players coincide; denote by u_{θ} and u_{τ} , respectively, the utilities of types θ and τ . Payoffs and utilities may or may not be identical; see Definition 14.12, below.

Definition 14.9 (*Bayesian Nash equilibrium*): Choose a particular state $s = (\theta, \tau, \epsilon)$. A pair of strategies $x^*, y^* \in X$ is a Bayesian Nash equilibrium if

$$\begin{cases} x^* \in \arg \max_{x \in X} P(\theta | \theta, \epsilon) u_\theta(x, x^*) + P(\tau | \theta, \epsilon) u_\theta(x, y^*), \\ y^* \in \arg \max_{y \in X} P(\theta | \tau, \epsilon) u_\tau(y, x^*) + P(\tau | \tau, \epsilon) u_\tau(y, y^*). \end{cases}$$

Restricting attention to the set of BNE, we can define an ESS under asymmetric information; this is a useful tool for the readers to add to their toolkit.

Definition 14.10 (*Evolutionary stability*): Suppose that $x^*, y^* \in X$ is a Bayesian Nash equilibrium. Type $\theta \in \Theta$ is evolutionarily stable against type $\tau \in \Theta$ if there exists an invasion barrier $\bar{\epsilon}$ such that, for all $0 < \epsilon < \bar{\epsilon}$, the following condition holds in all Bayesian Nash equilibria for all states $s = (\theta, \tau, \epsilon)$

$$\Pi_\theta(x^*, y^*, \epsilon) > \Pi_\tau(x^*, y^*, \epsilon).$$

Type θ is evolutionarily stable if it is evolutionarily stable against all $\tau \in \Theta$ such that $\tau \neq \theta$.

Random matching takes place through an assortative matching function, ϕ , as in Bergstrom (2003); ϕ is the difference in the conditional probability that type θ is matched with another type θ and the conditional probability that type τ is matched with type θ .

Definition 14.11 (*Assortment function and index of assortativity*): Define the assortment function $\phi : (0, 1) \rightarrow [-1, 1]$ such that

$$\phi(\epsilon) = P(\theta | \theta, \epsilon) - P(\theta | \tau, \epsilon).$$

We assume that ϕ is a continuous function and

$$\lim_{\epsilon \rightarrow 0} \phi(\epsilon) = \sigma, \quad (14.92)$$

where $\sigma \in [0, 1]$ is known as the “index of assortativity” of the matching process. Letting $\phi(0) = \sigma$ allows us to extend the domain of ϕ from $(0, 1)$ to $[0, 1]$. The special case of uniform random matching between players, irrespective of type, corresponds to $\phi(\epsilon) = 0$ for all $\epsilon \in (0, 1)$. In this case, the probability that a type θ or type τ player is matched with another type θ player is identical.

We now define the precise sense in which individuals possess moral virtues.

Definition 14.12 (*Homo-moralis*): An individual is said to be Homo-moralis if his utility function is given by the following convex combination of two different tendencies.

$$u(x, y; \kappa) = (1 - \kappa)\pi(x, y) + \kappa\pi(x, x), \quad (14.93)$$

where $\kappa \in (0, 1)$ is the degree of morality of the individual, $\pi(x, y)$ is the payoff from the strategy x , against an opponent who plays y , and $\pi(x, x)$ is the payoff against a hypothetical opponent who also plays x .

From Definition 14.12, a Homo-moralis individual cares about selfishness (first term in (14.93)) and morality (second term in (14.93)). Morality takes the *Kantian form*, i.e., choose a strategy such that if the same strategy is chosen by all players, it leads to the highest possible payoff, $\pi(x, x)$. Morality as defined above, is different from altruism, inequity aversion, or efficiency concerns. Two limiting cases of Definition 14.12 are stated next.

Definition 14.13 Consider the definition of utility function $u(x, y; \kappa)$ in Definition 14.12. An individual is *Homo oeconomicus* in the limiting case $u(x, y; 0)$ and *Homo kantiensis* in the limiting case $u(x, y; 1)$. Finally, using (14.92), the individual is termed as *Homo hamiltonensis* in the special case $u(x, y; \sigma)$.

From Definition 14.13, *Homo oeconomicus* is a textbook economic individual who is *amoral*. *Homo kantiensis*, at the other extreme, chooses his strategy $x \in X$ to maximize $\pi(x, x)$. *Homo hamiltonensis* turns out to be an evolutionary important type, as we shall see below. When the degree of relatedness among the individuals is given by σ , $u(x, y; \sigma)$ is the fitness of a mutant individual who plays x when the main population plays y .

Let the type of an individual $\theta \in \Theta$ be identical to the degree of morality κ of the individual, $\kappa \in [0, 1]$, so the utility function in (14.93) can be written as $u(x, y; \theta)$. For any type $\theta \in \Theta$, denote the best reply correspondence by β_θ , where

$$\beta_\theta(y) = \arg \max_{x \in X} u(x, y; \theta) \text{ for all } y \in X.$$

Denote the fixed points under β_θ by X_θ , where

$$X_\theta = \{x \in X : x \in \beta_\theta(x)\}.$$

For the case of *Homo hamiltonensis*, the best reply correspondence is given by β_σ and the set of fixed points under β_σ is given by X_σ .

For any state $s = (\theta, \tau, \epsilon)$, let $B^{NE}(s) \subseteq X \times X$ denote the set of Bayesian Nash equilibria (Definition 14.9). We now define our final construct. For any state $s = (\theta, \tau, \epsilon)$, let ϵ be vanishingly small. Then type θ is known as the resident population and type τ , the mutant population. For any resident type $\theta \in \Theta$ we wish to identify those types $\tau \in \Theta$ that are behaviorally indistinguishable from θ . This is achieved by defining the following set

$$\Theta_\theta = \{\tau \in \Theta : \exists x \in X_\theta \text{ such that } (x, x) \in B^{NE}(\theta, \tau, 0)\}.$$

An example of behaviorally indistinguishable types are those in which the utility function of one type is an affine transformation of another. We are now in a position to state the main result.

Proposition 14.17 (Alger and Weibull, 2013):

- (a) If $\beta_\theta(x)$ is a singleton for all $x \in X_\theta$, then *Homo hamiltonensis* is an evolutionary stable strategy (see Definition 14.10) against all types $\tau \notin \Theta_\theta$.
- (b) Suppose that we do not have assortative matching, i.e., $\sigma = 0$. Consider *Homo oeconomicus* $\kappa = 0 \in \Theta$. If $\beta_0(x)$ is a singleton for all $x \in X_0$, then *Homo oeconomicus* is an evolutionary stable strategy (see Definition 14.10) against all types $\tau \notin \Theta_0$.

The implications of Proposition 14.17 are profound. Assortativity appears to be quite important in humans for reasons of kinship, networks, and social groups. Indeed, if assortativity is present to any degree, then evolution ensures that *Homo oeconomicus* is not evolutionarily stable. Under these conditions, a form of *Homo-moralis*, the *Homo hamiltonensis* is evolutionarily stable if β_0 is unique. In related results, under assortative matching, Bergstrom (2009) established the

evolutionarily stability of a semi-Kantian type, which corresponds to $\kappa = 0.5$ in the model above. Proposition 14.17(b) is consistent with the results of Ok and Vega-Redondo (2001) and Dekel et al. (2007) who find that in the absence of any assortativity, *homo oeconomicus* is an evolutionary stable strategy.

Since these results are independent of payoffs, they lead to the remarkable conclusion that economic incentives cannot crowd out human morality, at least in an evolutionary sense. We have reviewed these issues in Part 2. For instance, the experimental results of Fischbacher and Föllmi-Heusi (2013) show that roughly 40% of the subjects are *completely honest* even when lying could have improved their payoffs. A significant fraction of the subjects tell *partial lies*, even when they could have lied more strongly (see also Gneezy et al., 2013). About 22% behave in a manner consistent with the *selfish black liars* assumed in neoclassical economics. The bottom line of this analysis is that economics needs to consider a blend of selfishness and morality, but a consideration of ‘selfishness only’ is simply not tenable. Readers who have read Part 2 of the book carefully might already have reached similar conclusions.

For a separate analysis of the evolutionary development of sibling altruism and the role of the harshness of the environment, in the presence of assortative matching, see Alger and Weibull (2010). They find that in the presence of uncertainty in the economic environment, altruism may be a double edged sword. On the one hand, altruism provides an insurance function between siblings, say, if the siblings experience very different economic outcomes. On the other hand, this insurance may lead to a reduction in the effort exerted by the siblings on account of moral hazard issues. As one might expect in this case, the beneficial effects of altruism, and its evolutionary stability, depend on the degree of uncertainty in the environment.

14.9.7 Other issues in evolutionary explanations of cooperation

Strong reciprocators are predisposed to cooperate, but they are willing to punish members of the group who do not adhere to the group norms, at a personal cost (see Part 2 of the book). In Part 2, in experimental public goods games, the nature of punishment is typically unilateral, as opposed to group punishment. This may reduce the average payoff of the group. Further, in experiments conducted in non-Western countries, we may also observe antisocial punishment of individuals who cooperate; prosocial punishments are correlated with the rule of law and other social institutions. Anthropological and ethnographic evidence, however, shows that individuals typically engage only in group punishment, e.g., social ostracization, and punishment raises the average group payoffs. Using these stylized facts, Boyd et al., (2010) show that strong reciprocity can evolve as an evolutionary equilibrium.

In this chapter, we have considered random bilateral interactions between players of the same, or different populations. However, in actual practice, people may interact disproportionately with others of their own type, i.e., matchings may be assortative. Thus, cooperators may survive in prisoner dilemma type games by interacting mostly with other cooperative people in their social clusters. This *network reciprocity* is a form of spatial reciprocity; for an explanation of cooperative behavior based on network reciprocity, see Durrett and Levin (1994) and Hauert and Doebeli (2004). Nowak (2006) shows that despite the complexities of this framework, a simple form of Hamilton’s rule applies to determine the success of the cooperative outcome under network reciprocity. This rule is $\frac{b}{c} > k$ where k is the average number of neighbors that a player has in the network.

	C	D
A	$\alpha + \beta, 0$	$0, 0$
B	$\alpha, 0$	$\alpha, 0$

Figure 14.17 A two-population, 2×2 strategic form game.

14.10 Social learning and replicator-like dynamics

We now consider some simple models of social learning in which players are not hardwired to play particular strategies. Yet, these models give rise to replicator-like dynamics in models of evolutionary game theory. We then consider the broader concept of *payoff monotonicity*, of which the replicator dynamics is a special case.

14.10.1 Social imitation and emulation dynamics

Consider the normal form game in Figure 14.17 that is played between two populations of players, population 1 (row players) and population 2 (column players); $\alpha, \beta > 0$.

Players in population 2 always get a payoff of 0. Using Fudenberg and Levine's (1998) terminology, we call this a dummy population. In population 1, the respective fractions of players playing A and B are $x_t(A) \geq 0$ and $x_t(B) = 1 - x_t(A)$. In population 2, the proportion of players who play strategies C and D are denoted, respectively, by the time-invariant fractions $y(C)$ and $y(D) = 1 - y(C)$.

Suppose that at the end of each time period, and in population 1 only, new players, the "children," replace their "parents." A fraction $1 - \lambda$ of the children inherit the strategy of the parent. However, the remaining fraction, $\lambda > 0$, get a revision opportunity. They randomly sample one more player, the "stranger," from their population and are able to observe his/her payoff from the latest period. They then use the following decision rule.

Decision Rule: Of the two observed strategies, the parent's strategy and the stranger's strategy, choose the one that has a higher payoff. If the payoffs from both strategies are identical, then choose the parent's strategy. Dynamics based on this decision rule are called *emulation dynamics* (Ellison and Fudenberg, 1993; Fudenberg and Levine, 1998).

The emulation dynamics favor those strategies that have a high median payoff rather than a high average payoff, as we shall see below. Given $x_t(A)$, we would like to determine this fraction in the following period, $x_{t+1}(A)$.

Consider the children of parents at time t who played the strategy A. There are two kinds of children.

1. A fraction $x_t(A)y(C)$ are children of parents whose payoff at time t was $\alpha + \beta$. A fraction $(1 - \lambda)x_t(A)y(C)$ does not get an opportunity to revise their strategy. The remaining fraction $\lambda x_t(A)y(C)$ gets a revision opportunity. Given the decision rule, they randomly sample the payoff of a stranger from their population and compare it to the payoff of their parent. Their parent always has the highest payoff, so they continue to choose their parent's strategy, A.
2. A fraction $x_t(A)y(D)$ are children of parents whose payoff at time t was 0. A fraction $(1 - \lambda)x_t(A)y(D)$ do not get a revision opportunity. The remaining fraction $\lambda x_t(A)y(D)$

that gets the revision opportunity switches the strategy to B with probability $x_t(B)$. Thus the total fraction that switches to B is $\lambda x_t(A)y(D)x_t(B)$ and the fraction that does not switch is $\lambda x_t^2(A)y(D)$.

Now consider the children of parents in population 1 who play B at time t , who comprise a fraction $x_t(B)$ of population 1. Of these, a fraction $\lambda x_t(B)$ get a chance to revise their strategy. When they randomly sample a stranger from their population, they find that a fraction $x_t(A)y(C)$ has a higher payoff ($\alpha + \beta > \alpha$). So a total fraction $\lambda x_t(B)x_t(A)y(C)$ of the children of B parents switch to the strategy A .

Putting this information together, and noting that $(1 - \lambda)x_t(A)y(C) + (1 - \lambda)x_t(A)y(D) = (1 - \lambda)x_t(A)$, the fraction of population 1 that plays strategy A at time $t + 1$ is

$$x_{t+1}(A) = (1 - \lambda)x_t(A) + \lambda x_t(A)y(C) + \lambda x_t^2(A)y(D) + \lambda x_t(B)x_t(A)y(C).$$

Alternatively, we can write the RHS in terms of those do not get a revision opportunity and those who do get the opportunity,

$$x_{t+1}(A) = (1 - \lambda)x_t(A) + \lambda [x_t^2(A) + 2x_t(A)x_t(B)y(C)]. \quad (14.94)$$

Using $x_t(B) = 1 - x_t(A)$ and rewriting (14.94) we get

$$x_{t+1}(A) - x_t(A) = cx_t(A)(1 - x_t(A)), \text{ where } c = \lambda(2y(C) - 1). \quad (14.95)$$

To find the continuous time analog of (14.95), let the time difference between successive periods be Δ , then (14.95) can be written as

$$x_{t+\Delta}(A) - x_t(A) = \Delta cx_t(A)(1 - x_t(A)). \quad (14.96)$$

Divide both sides of (14.96) by Δ and take the limit as $\Delta \rightarrow 0$ we get

$$\dot{x}_t(A) = f(x_t(A)) = cx_t(A)(1 - x_t(A)). \quad (14.97)$$

Setting $\dot{x}_t(A) = 0$, the steady states are given by $\bar{x}(A) = 0$ and $\bar{x}(A) = 1$. Since $f'(x) = c(1 - 2x)$, at the two steady states, we have $f'(0) = c$, $f'(1) = -c$. Stability requires that $f'(x) < 0$. Using (14.95)

$$c \begin{matrix} \geq \\ \leq \end{matrix} 0 \text{ as } y(C) \begin{matrix} \geq \\ \leq \end{matrix} \frac{1}{2}.$$

Hence, $\bar{x}(A) = 0$ is a stable steady state if $y(C) < \frac{1}{2}$, while $\bar{x}(A) = 1$ is a stable steady state if $y(C) > \frac{1}{2}$. Particularly disturbing is the stable steady state $\bar{x}(A) = 0$ when $y(C) < \frac{1}{2}$. In this case, social imitation leads everyone in population 1 to play the pure strategy B . What is wrong with this? The best response for a player in population 1 is to play A rather than B if $(\alpha + \beta)y(C) > \alpha$, or

$$y(C) > \hat{y} = \left(1 + \frac{\beta}{\alpha}\right)^{-1}.$$

As an example, consider $\frac{\beta}{\alpha} = 5$, so $\hat{y} = \frac{1}{6}$. When $\frac{1}{6} < y(C) < \frac{1}{2}$, the best response of a player from population 1 is to play A , while social imitation leads the whole population to play B instead. Ordering the payoffs of population 1 players in increasing order, we get $0, \alpha, \alpha, \alpha + \beta$, thus, the median payoff is α . Since $\bar{x}(A) = 0$ for $y(C) < \frac{1}{2}$, the payoff of population 1 players is the median payoff α , hence, such dynamics are sometimes called *median enhancing emulation dynamics*.

The model of emulation dynamics sets the stage for our microfoundations of the replicator dynamics in the next section.³¹

14.10.2 Some microfoundations of replicator dynamics: imitation learning

Suppose that there are two populations of players, $K = \{1, 2\}$ who are paired randomly. The respective pure strategy sets in the two populations are, $S_1 = (s_{11}, s_{12}, \dots, s_{1n_1})$ and $S_2 = (s_{21}, s_{22}, \dots, s_{2n_2})$. The payoff matrix of a representative player in population 1 is denoted by Π_1 and the p th row in Π_1 is given by

$$\Pi_{p1} = (a_{p1}, a_{p2}, \dots, a_{pn_2}),$$

where a_{pr} is the payoff to player 1 when he plays the pure strategy s_{1p} and player 2 plays the strategy s_{2r} . Similarly, we denote row p in the payoff matrix, Π_2 , of player 2 by

$$\Pi_{p2} = (b_{p1}, b_{p2}, \dots, b_{pn_1}),$$

where b_{pr} is the payoff to player 2 when he plays the pure strategy s_{2p} and player 1 plays the strategy s_{1r} .

A proportion x_{kj} of players in population $k \in K$ plays the pure strategy $s_{kj} \in S_k$. Let $x_k = (x_{kj})_{n_k \times 1}$ denote the column vector of these proportions in population k and denote the transpose by x_k^T . Thus, the expected payoff of player 1 from the pure strategy s_{1p} is given by $\Pi_{p1}x_2$. Likewise the expected payoff of player 2 from the pure strategy s_{2p} is given by $\Pi_{p2}x_1$.

Suppose that a fraction λ_{kj} of population $k \in K$ that is playing the pure strategy s_{kj} randomly gets an opportunity to revise its strategy. Denote by ρ_{kj}^i the proportion of these individuals who adopt the pure strategy s_{ki} , having abandoned the pure strategy s_{kj} . We now make assumptions about revision opportunities.

A1: The fraction λ_{kj} of individuals who play the pure strategy s_{kj} and get a revision opportunity, is a decreasing function of the expected payoff from the strategy, s_{kj} . Thus,

$$\begin{cases} \lambda_{1j} = \phi_1(\Pi_{j1}x_2); j = 1, 2, \dots, n_1 \\ \lambda_{2j} = \phi_2(\Pi_{j2}x_1); j = 1, 2, \dots, n_2 \end{cases}, \quad (14.98)$$

where for $k \in K$, $\phi_k : \mathbb{R} \rightarrow [0, 1]$, and $\phi'_k < 0$. There are two reasons why the revision opportunity depends negatively on the expected payoff from a strategy. First, there could be some fixed costs of switching strategies, so it is worth considering a switch only for strategies that are doing particularly poorly. Second, players might have aspiration levels,

³¹ Similar models are presented in Bjornerstedt (1995), Fudenberg and Levine (1998), and Vega-Redondo (2003).

and low payoff strategies may fall below aspirations, hence, motivating a greater need to switch. For pedagogical simplicity, let us specialize ϕ_1, ϕ_2 to be identical and linear, so

$$\phi_1 = \phi_2 = \phi \text{ and } \phi(y) = a - by. \quad (14.99)$$

A2: Whenever any player gets an opportunity to revise his strategy, he randomly picks a player from his own population and adopts his strategy (*imitation*). Given the frequencies with which various pure strategies are played in one's own population, the probability that a player who plays s_{kj} meets a player who plays s_{ki} is given by x_{ki} , hence, for all k, j, i

$$\rho_{kj}^i = x_{ki}. \quad (14.100)$$

Define the state of the system, x , as the fraction of players from each population who play each of the available pure strategies at any given time (we suppress time subscripts for now)

$$x = (x_1, x_2).$$

In any population $k \in K$, among the players who have a revision opportunity, the total inflow of players into strategy s_{ki} is given by

$$I_{ki} = \sum_{j=1}^{n_k} (\lambda_{kj} x_{kj}) \rho_{kj}^i; \forall k, j \neq i, \quad (14.101)$$

where $\lambda_{kj} x_{kj}$ is the fraction of individuals who play strategy s_{kj} and get an opportunity to revise their strategy. Of this, a fraction ρ_{kj}^i will switch to strategy s_{ki} . Summing across all strategies gives the total inflow into the strategy s_{ki} . The fraction $\lambda_{ki} x_{ki}$ of players in population $k \in K$ was using strategy s_{ki} and has been given an opportunity to revise their strategy. Given A1, A2, all of these players observe and imitate the strategy of another randomly picked player from their population who does not play s_{ki} . Hence, the total outflow O_{ki} from the set of players using strategy s_{ki} is given by

$$O_{ki} = \lambda_{ki} x_{ki}. \quad (14.102)$$

The evolution of the frequencies in each population that plays different pure strategies is given by the following equations of motion.

$$\dot{x}_t(s_{ki}) = I_{ki} - O_{ki}; \forall k, i = 1, \dots, n_i. \quad (14.103)$$

Substituting (14.98)–(14.102) in (14.103), and denoting by x_{-k} , the vector of frequencies for the other population:

$$\begin{aligned} \dot{x}_t(s_{ki}) &= \sum_{j=1}^{n_k} ((a - b \Pi_{jk} x_{-k}) x_{kj}) x_{ki} - (a - b \Pi_{ik} x_{-k}) x_{ki} \\ &= \left[\sum_{j=1}^{n_k} ((a - b \Pi_{jk} x_{-k}) x_{kj}) - (a - b \Pi_{ik} x_{-k}) \right] x_{ki} \\ &= \left[a \sum_{j=1}^{n_k} x_{kj} - b \sum_{j=1}^{n_k} x_{kj} \Pi_{jk} x_{-k} - a + b \Pi_{ik} x_{-k} \right] x_{ki}. \end{aligned} \quad (14.104)$$

Since the population frequencies sum upto 1, $\sum_{j=1}^{n_k} x_{kj} = 1$. By definition, $\Pi_{jk}x_{-k} = \pi(s_{kj}, x_{-k})$ is the expected payoff to a population k player from playing the pure strategy s_{kj} . Hence, we can rewrite (14.104) as:

$$\dot{x}_t(s_{ki}) = b [\pi(s_{ki}, x_{-k}) - \pi(x_k, x_{-k})] x_{ki}. \quad (14.105)$$

For $b = 1$, the resulting population dynamics, given by (14.105), are identical to the replicator dynamics for the case of two or more populations; see (14.35). However, this result is not necessarily robust to perturbations in the form of social learning.

Other microfoundations of replicator dynamics are also available. Bjornerstedt and Weibull (1993) consider a similar model to the one we describe above, except that players receive noisy statistical information about the payoff of others. In the special case when the noise has a uniform distribution, and the function ϕ has the linear form in (14.99), one obtains the replicator dynamics. Borgers and Sarin (1997) also recover replicator-like dynamics in a stimulus response model.

14.10.3 Social dynamics based on payoff monotonicity

Assumption A1 above, required that as the payoff from a pure strategy increases, players are less likely to wish to revise it. This appears to be a desirable feature of most models of learning. One way to formalize this idea is to use the concept of *payoff monotonicity*. Relying on payoff monotonicity alone, of which the replicator dynamics is a special case, we derive some interesting properties of an *evolutionary system* (defined below).

Definition 14.14 (*Evolutionary system*): Suppose that $K = \{1, 2, \dots, n\}$ is the set of competing populations. Players in population $k \in K$ have the pure strategy set $S_k = (s_{k1}, s_{k2}, \dots, s_{kn_k})$. A proportion x_{kj} of players in population $k \in K$ plays the pure strategy $s_{kj} \in S_k$. Let $x_k = (x_{kj})_{n_k \times 1}$ denote the column vector of these proportions in population k and x_k^T , the transpose of x_k . The state of the system is given by $x = (x_1, \dots, x_n)^T$. x_{-k} denotes the frequencies of play of all pure strategies in all populations other than the k th population. A social evolutionary system is a system of differential equations,

$$\dot{x}_t(s_{ki}) = g_{ki}(x), \forall k \in K \text{ and } \forall s_{ki} \in S_k. \quad (14.106)$$

We next define payoff monotonicity, which captures the idea that strategies that have a relatively higher payoff grow at a relatively faster rate.

Definition 14.15 (*Payoff monotonicity*): An evolutionary system is said to be *payoff monotonic* if for all x, k, j, i

$$\pi(s_{ki}, x_{-k}) \geq \pi(s_{kj}, x_{-k}) \Rightarrow \frac{\dot{x}_t(s_{ki})}{x_t(s_{ki})} \geq \frac{\dot{x}_t(s_{kj})}{x_t(s_{kj})}. \quad (14.107)$$

Definition 14.16 (*Regular evolutionary system*): The evolutionary system in Definition 14.14 is said to be *regular* if g_{ki} is a C^1 function for all k, i .³²

³² This condition can be weakened to Lipschitz continuity.

Regularity ensures that a unique solution exists to the problem. The analogues of Propositions 14.7, 14.8 for replicator dynamics can be shown to exist for *payoff monotone and regular dynamics* (PMR dynamics). The proofs are analogous to those in Propositions 14.7, 14.8, so we omit them.

Proposition 14.18 *Suppose that there exists some symmetric Nash equilibrium, $\sigma = (\sigma_1, \dots, \sigma_n)$, of the symmetric game $G(2, S, \pi)$. The interpretation is that σ_i is the probability with which the pure strategy s_i is played by both the populations. Consider the symmetric population state*

$$x_t(s) = \sigma \Leftrightarrow (x_t(s_1), x_t(s_2), \dots, x_t(s_n)) = (\sigma_1, \dots, \sigma_n),$$

where $x_t(s_i)$ is the population frequency with which strategy s_i is played. Then the population state $x_t(s)$ is a dynamic equilibrium (or steady state) of the PMR dynamics.

Proposition 14.19 *Consider the game $G(2, S, \pi)$. Suppose that there is some asymptotically stable state of the PMR dynamics given by*

$$\bar{x} = (\bar{x}(s_1), \bar{x}(s_2), \dots, \bar{x}(s_n)).$$

Now define a mixed strategy

$$\bar{\sigma} = \bar{x} : \bar{\sigma}_i = \bar{x}(s_i); i = 1, 2, \dots, n.$$

Then $(\bar{\sigma}, \bar{\sigma})$ is a symmetric Nash equilibrium of the game.

At least with respect to the two important results in Propositions 14.18, 14.19, replacing replicator dynamics by PMR dynamics is without loss in generality. However, these results rely on the existence of an asymptotically stable steady state, which might not exist in some cases. In this case, one may wish to ensure a weaker notion of rationality embodied in, say, the avoidance of dominated strategies.

Definition 14.17 *(Dominance in pure strategies): Consider a bilateral game with two populations, i.e., $K = \{1, 2\}$ and let $k = 1, 2$ index the two populations. A pure strategy $s_{ki} \in S_k$ is dominated in pure strategies if there exists some other pure strategy $s_{kj} \in S_k$ such that*

$$\pi_k(s_{kj}, x_{-k}) \geq \pi_k(s_{ki}, x_{-k}) \text{ for all } x_{-k}, \quad (14.108)$$

and at least one inequality is strict; x_{-k} is the vector of frequencies with which the other population plays its pure strategies.

We now give the main result for PMR dynamics. If some strategy is dominated in pure strategies, then PMR dynamics ensure that the dominated strategy disappears eventually.

Proposition 14.20 *(Samuelson and Zhang, 1992): Consider a bilateral game with $K = \{1, 2\}$. The pure strategy set of population $k \in K$ is given by $S_k = (s_{k1}, s_{k2}, \dots, s_{kn_k})$. The vector of frequencies with which these pure strategies are played is denoted by x_k . Suppose that the PMR dynamics induce a path $x(t) = (x_1(t), x_2(t))$ starting from some fully mixed initial point $x(0) = (x_1(0), x_2(0))$. If s_{ki} is dominated in pure strategies, then it follows that $\lim_{t \rightarrow \infty} x_t(s_{ki}) = 0$.*

Proof: Since s_{ki} is dominated in pure strategies, there exists some other pure strategy $s_{kj} \in S_k$ such that

$$\pi_k(s_{kj}, x_{-k}) - \pi_k(s_{ki}, x_{-k}) \geq \varepsilon \text{ for all } x_{-k} \text{ and some } \varepsilon > 0.$$

Then from continuity, and payoff monotonicity in (14.107) we get that for all $t \geq 0$,

$$\frac{\dot{x}_t(s_{kj})}{x_t(s_{kj})} - \frac{\dot{x}_t(s_{ki})}{x_t(s_{ki})} \geq \varepsilon_1 \text{ for some } \varepsilon_1 > 0. \quad (14.109)$$

Let

$$y_t = \frac{x_t(s_{kj})}{x_t(s_{ki})}, \quad (14.110)$$

then, using (14.109) we get

$$\frac{\dot{y}_t}{y_t} = \frac{\dot{x}_t(s_{kj})}{x_t(s_{kj})} - \frac{\dot{x}_t(s_{ki})}{x_t(s_{ki})} \geq \varepsilon_1.$$

It follows that $y_t \geq y_0 e^{\varepsilon_1 t}$. Since $\varepsilon_1 > 0$, $\lim_{t \rightarrow \infty} y_t = \infty$, so, from (14.110), it must be the case that $x_t(s_{ki}) \rightarrow 0$. ■

The result in Proposition 14.20 can be strengthened (see Proposition 3.4 in Fudenberg and Levine, 1998). It can be shown that if some pure strategy s_{ki} is eliminated by iterative deletion of dominated strategies, then along any path taken by the PMR dynamics, the share of the population that plays strategy s_{ki} goes to zero. However, as shown in Bjornerstedt (1995), pure strategies that are dominated by a mixed strategy need not be eliminated by payoff monotonic dynamics. Two approaches have been taken to rule out pure strategies that are dominated by a mixed strategy. Samuelson and Zhang (1992) specify an *aggregate monotonicity condition*, which requires that if a mixed strategy has a higher current payoff relative to another mixed strategy, then it also has a higher growth rate. Hofbauer and Weibull (1996) provide a weaker condition called *convex monotonicity*. This condition requires that if a mixed strategy has a higher current payoff relative to a pure strategy, then it also has a higher growth rate.

CHAPTER 15

Models of Learning

15.1 Introduction

Models of evolutionary game theory cannot explain the rapid rate of learning in experiments. By *learning*, we refer to any change in observed behavior as, *ceteris-paribus*, players accumulate experience. Models of learning may provide the foundations of equilibrium concepts, such as Nash equilibrium, that are not based on explicit learning dynamics. Or, we may find that existing equilibrium concepts do not organize very well the data from learning games. In this case, the emerging evidence can provide important clues for the development of new theories.

Section 15.2 considers a classical model of learning, the *reinforcement model of learning*, also known as *stimulus learning* or *rote learning* (Bush and Mostellar, 1951, 1955; Luce, 1959; Suppes and Atkinson, 1960). In this model, players have minimal rationality and limited information about the game being played. Players need not know the payoff matrix, or the strategies/types of other players. All that players need to know is the set of actions that they undertook in the past and the associated payoffs. Strategies that have given rise to higher payoffs in the past are more likely to be played in the present (*law of effect*), and there is no presumption that players should play best replies. This is formalized by associating with each strategy, a number called its *attraction* or *propensity*. Propensities are then reinforced with information about payoffs arising from the play of various strategies. The available strategies are then played probabilistically, in proportion to their relative propensities.

There are many variants of reinforcement learning. Section 15.2.1 considers the one-parameter reinforcement model, which requires a specification only of the initial propensities (Arthur, 1991, 1993). Section 15.2.2 considers the three-parameter model of reinforcement learning (Roth and Erev, 1995; Erev and Roth, 1998). This extends the one-parameter model to allow for (i) *experimentation*, by updating the propensities of strategies that are not played, but are similar to those that are played, and (ii) *recency*, by letting past experience attenuate and assigning greater weight to more recent experience. Section 15.2.3 briefly discusses other variants of reinforcement learning. These include the incorporation of payoff shocks (Sarin and Vahid, 1999, 2001), and learning from the payoff experience of others (Duffy and Feltovich, 1999).

Section 15.2.4 explores the optimality of reinforcement learning. Suppose that an individual plays a game against nature; nature chooses its strategies based on a stationary distribution. Then, under a form of reinforcement learning, called *cumulative proportional reinforcement learning*, the limiting play by the individual is optimal in the sense that the chosen strategy is a best response (Laslier et al., 2001). Indeed, such limiting choices can be shown to belong to asymptotically stable states of replicator dynamics (Hopkins and Posch, 2005).

Section 15.2.5 considers the empirical evidence for reinforcement learning. The main result is that while reinforcement learning is good at predicting the direction of change, the predicted speed of learning is too sluggish (Roth and Erev, 1995; McAllister, 1991). If the speed of adjustment is increased, then the model fails in other respects, e.g., players start to play dominated actions (Van Huyck et al., 2007). Furthermore, contrary to the predictions of reinforcement learning, information about other players, their payoffs, or even their strategies, affects choices (Mookerjee and Sopher, 1994; Rapoport and Erev, 1998). Hence, unless the underlying structure of the game is poorly understood by players, and they have relatively low information, reinforcement learning leaves out important determinants of human learning.

Section 15.3 considers models of *belief-based learning*. In these models, players compute the likely mixed strategy of the opponent, given observations on the past frequencies of play by the opponent. The strategies of the opponent are assigned some initial weights that are updated, depending on their observed strategies. The probability assigned to an opponent's pure strategy is the relative weight of that strategy. Conditional on the expected mixed strategy of the opponent, players then play a myopic best response. A prominent special case is Cournot dynamics (Cournot, [1838] 1927). The rationality and informational requirements in belief-based models are greater relative to models of reinforcement learning.

It can be shown that for 2×2 games (two players, and two pure strategies per player), the frequency of play under belief-based models has a limiting distribution that corresponds to a Nash equilibrium. These results were originally demonstrated for simple 2×2 games (Brown, 1951; Robinson, 1951), and then for more general 2×2 games (Miyasawa, 1961; Monderer and Shapley, 1996). Before the game even begins, players can compute the limiting frequency of play and when they actually have a chance to choose their strategies, they simply play their best response to the limiting frequency. The resulting best response of each player constitutes a Nash equilibrium. Hence, this kind of learning is also termed as *fictitious play*. However, these results hold only for the 2×2 case, and they may fail even for 3×3 games (Shapley, 1964).

One drawback of fictitious play is that the choice of strategies is too rigid and discontinuous. This leads to several well-known problems. For instance, players do not learn from systematic patterns of play. This problem is addressed by making the best response functions *smooth* (Fudenberg and Kreps, 1993). A popular solution is to introduce additive noise in payoffs and assume that the noise follows an extreme value distribution. In this case, the best response follows the smooth logistic function.

Section 15.3.1 reviews some of the empirical evidence on fictitious learning. *Weighted fictitious play* nests ordinary fictitious play and Cournot play as two polar extremes. Using such a model, some studies find that the majority of observations are closer to the Cournot extreme (Cheung and Friedman, 1997). Others find that play is closer to the predictions of ordinary fictitious play and there is divergence between stated beliefs and the beliefs fitted to a model of fictitious play (Nyarko and Schotter, 2002). Interestingly, the actual play is consistent with players who play a best response to stated beliefs, not fitted beliefs, even though the latter converge quickly, while the former are very volatile.

There is a departure from myopic best reply dynamics when one varies the information that is available to players (Huck et al., 1999). In this case, results are sensitive to whether players have information about the market (cost/demand conditions), or information about others (quantity/profits of others). A variety of other learning mechanisms, such as *imitate the best*, *learning direction theory*, and *trial and error learning* explain well the results in several treatments. These results have been extended (Bosch-Domènech and Vriend, 2003). In summary, fictitious play, by itself, is too simple to explain all relevant aspects of human learning.

Section 15.3.2 looks at the empirical performance of fictitious learning, relative to reinforcement learning. Results tend to depend on the type of the game, the particular version of each learning model used, and the treatments allowed for. For these reasons, general conclusions are hard to draw out. Playing dominated strategies is consistent with reinforcement learning, but not with fictitious play because, in the latter, players always play best replies. There is evidence that in repeated rounds of experiments, players sometimes play dominated strategies early on, but are less likely to do so in later rounds (Mookerjee and Sopher, 1997). Reinforcement learning appears to fit better in games with relatively low dimensional strategy spaces that have a mixed strategy equilibrium. However, belief-based models appear to make relatively better predictions in coordination games.

The *optimization premium* (forgone payoff from playing a non-best response) substantially alters the dynamics of play, even when there is no change in the Nash equilibrium of the game (Battalio et al., 2001). Furthermore, the optimization premium plays an important role in understanding the long-run choice between payoff-dominant and risk dominant equilibria. These results reinforce the message that it is not sufficient to focus on best response dynamics.

Camerer and Ho (1999) proposed the five-parameter *experience weighted attraction* (EWA) model. In this model, in any round, players update propensities of strategies that are not played in that round by taking account of the forgone payoffs. They also allow for depreciation of past experience on account of limited memory, or forgetfulness. We explain the EWA model in Section 15.4. Reinforcement learning, belief-based learning models and EWA models are sometimes also known as *adaptive learning models* because, at any instant in time, players adapt their current strategies to the past choice of strategies and/or to the past observations of payoffs.¹

On the face of it, reinforcement learning and fictitious play would appear to be two distinct models. In reinforcement learning, players completely ignore the strategies/payoffs of others, concentrating only on their own history of payoffs and strategies. Since strategies that have given higher payoffs in the past are more likely to be played now, such players play a *better response*, but not a *best response*. In fictitious play, players focus on the historical record of strategies played by the opponents in order to form inferences about their likely current choice of strategies and then play a best response to these expectations. However, we show in Sections 15.4.1 and 15.4.2, that these two models are simply special cases of the EWA model for different parameter configurations. For other parameter values, EWA specializes to a range of other known models such as the Cournot model, cumulative reinforcement model, weighted fictitious play, and the average reinforcement model.

Empirical evidence on the EWA model is considered in Section 15.4.3 (Camerer and Ho, 1999; Camerer et al., 2001; Camerer and Ho, 2014). One possible objection to the EWA model is that its performance is unsurprisingly better than other models, simply because it has a larger number of free parameters. There are two ways of addressing this concern. (i) One can compare the relative performance of the alternative models by using the parameters estimated from in-sample observations to fit out-of-sample observations. There is no guarantee that a model with more free parameters performs better out-of-sample, although it probably would perform better in-sample. (ii) One can use statistical tests such as the Akaike and Bayesian information criteria that trade off better explanatory power against a larger number of parameters. On both criteria, EWA does better relative to reinforcement learning and fictitious play.

¹ It is very plausible to suppose that players have a sudden eureka moment and figure out the structure of the underlying game and/or the strategies of the opponent. Recent experimental evidence supports this form of learning and it will clearly play a more important part in learning models of the future; see, for instance, Dufwenberg et al. (2010) and Gneezy et al. (2010) for experimental evidence.

A more parsimonious alternative to the EWA model is the one-parameter *functional* EWA model or fEWA or *EWA-lite* (Ho et al., 2002); we consider this in Section 15.4.4. The parameters in this model are *self-tuning* in the sense that they respond to changes in the environment. For instance, when the environment is changing rapidly, then less weight is given to past experience. The single parameter of the model is the parameter in the logit function that is used to convert attractions to probabilities. When this model is fitted to the data, it performs better, relative to reinforcement learning and fictitious play. In a comprehensive horse race between the alternative models, tested on seven different games, EWA and fEWA outperform other models such as reinforcement learning, weighted fictitious play, and quantal response equilibrium (Ho et al. 2002; Camerer and Ho, 2014).

Adaptive players in the EWA model ignore information on the payoffs of others, even when doing so would allow them to outguess others and improve their payoffs. *Sophisticated players* are likely to do the opposite. When sophisticated players are matched with adaptive players, they may attempt to teach them, when it improves their future payoffs, even at the cost of short-run payoffs. Formal models of this phenomenon are sometimes known as models of *anticipatory learning*. Furthermore, in games like the p-beauty contest, relatively experienced players are known to converge relatively faster to the Nash equilibrium outcome. The *EWA model with sophistication* takes account of these considerations (Camerer et al., 2002; Camerer et al., 2003). We consider this model in Section 15.4.5. It is able to explain the difference in behavior among experienced and inexperienced subjects in 10 period p-beauty contest games. Experienced players are also less likely to be overconfident about the relative proportion of adaptive players in the population (Camerer et al., 2002).

In Section 15.5, we consider two alternative points of view about the nature of learning models and their generalizability that one may associate with the names of Colin Camerer and Amnon Rapoport. Colin's approach, in common with the general approach in economic modeling, is to strive to build general models that apply to all players/contexts/frames and use statistical criteria, such as AIC and BIC, that trade off greater generality of models against increased explanatory power, particularly in out-of-sample predictions. Amnon rejects this approach and argues for learning models that are tailored for different contexts and games, as well as take greater account of individual heterogeneity and the forward-looking behavior of individuals.

In Section 15.6, we consider the class of *learning direction theories*. The basic model of *learning direction theory* (Selten and Stöcker, 1986) is considered in Section 15.6.1. The essential idea can be explained by the analogy of an archer whose objective is to hit a target (his best response), but he discovers that he has hit a spot to the left of the target. In the next period, he adjusts by aiming between the previous hit and the target. The expectation is that by these successive adjustments, the archer moves qualitatively in the direction of the target.

The aim of learning direction theory is qualitative, i.e., to predict the direction of change, rather than the quantitative aim of specifying the probabilities with which players play their strategies. The theory predicts well the high degree of observed cooperation in repeated prisoner's dilemma game experiments (Selten and Stöcker, 1986). It also predicts well the direction of bid changes in Samuelson and Bazerman's takeover game (Selten et al., 2005).

In Section 15.6.2, we consider extensions of learning direction theory to enable quantitative predictions about the probabilities with which players play their strategies. These extensions have been applied mostly to 2×2 games, although the principle should apply more generally. Two of these extensions are *impulse matching learning* (Selten and Chmura, 2008; Chmura et al., 2012) and *impulse balance equilibrium* (Selten et al., 2005 and Selten and Chmura, 2008). In these theories, payoffs that are higher or lower than some aspiration level lead, respectively, to gains and losses.

One may then transform any normal form game into a game where the payoffs reflect loss-aversion if they are below the aspiration level. In each period, the strategy that could have received a higher payoff if played, receives an *impulse* from the other strategy. These impulses are then cumulated and transformed into probabilities with which various pure strategies are played. We also consider two variants of theories in this class: *action-sample learning* and *payoff-sampling learning* (Selten and Chmura, 2008). As the names indicate, these theories require sampling of past actions or payoffs to determine the current probabilities of play.

These extensions to learning direction theory allow for a direct comparison of this class of models with other models of learning. An earlier empirical attempt (Selten and Chmura, 2008) that found the class of learning direction theories to perform better than the mainstream alternatives was shown to have computation errors (Brunner et al., 2011). This was followed by a comprehensive horse-race between the main theories (Chmura et al., 2012), which ranked them in the following order from best to worse: impulse matching learning, action-sample learning, self-tuning EWA, and reinforcement learning. Importantly, it was also shown that the superior ranking of the class of learning direction theories is not an artifact of the inclusion of loss aversion.

The experimental results from *order statistic games* that we examined in detail in Part 4 of the book, pose an intriguing puzzle (Van Huyck et al., 1990, 1991). These experiments revealed that when the order statistic is the median, the initial choices are dispersed, and there is a strong history-dependence of outcomes. The final outcome, when subjects make a choice among the first seven positive integers is typically about 4 or 5. By contrast, in minimum action games, play quickly converges to the lowest integer choice available to players. Many learning models would struggle to explain the results from order statistic games. Section 15.6.3 considers the explanation offered by Crawford (1991, 1995), which shares some features that make it close, in spirit, to the class of learning direction models.

In the learning models considered so far, players try to learn aspects of the environment and update their propensities about their own, or the opponent's strategies. However, they continue to adopt the same learning rule throughout, e.g., reinforcement learning, belief-based learning, or EWA. Section 15.7 relaxes this assumption, and allows for the possibility that, given a set of rules, players also learn which rules might perform better than others (Stahl, 1999, 2000, 2001; Stahl and Haruvy, 2002). Hence, in each stage, players choose probabilistically among rules, choosing those rules that have done better, as in reinforcement learning. Each rule then prescribes a probabilistic play of pure strategies.

Even if the population is homogeneous, players draw from among the set of rules by using independently and identically distributed densities. This will typically generate heterogeneity in the choice of rules among the population. We consider choice from a plausible set of rules. An important family of such rules are *evidence based rules* in which players weight the evidence for and against the rule.

The empirical evidence supports rule-based learning relative to no rule learning. The predictions of the model, when we allow for choice from a set of evidence based rules, outperform other kinds of learning in which subjects simply play Nash with or without noise (Stahl, 2000). Further, there is strong evidence for a form of simple cooperation rule in which players play the payoff dominant action, and have an initial propensity to cooperate. These experiments do not find evidence for aspiration based learning, or the tit-for-tat rule (Stahl and Haruvy, 2002).

So far, the learning environment has typically involved single-game situations. However, it is possible that individuals might be asked to learn in situations in which nature probabilistically draws a game from a set of games. Furthermore, in different treatments, the set of games may be small or large, and information to players on payoffs and the past strategies of the opponents

may also vary. We may term these factors as reflecting the underlying *complexity* of playing in a particular treatment. How likely is it that subjects learn to play the Nash equilibrium, as the degree of complexity is varied across treatments? These issues are considered in Section 15.8 (Grimm and Mengel, 2012).

The basic hypothesis is that, as complexity increases, players with scarce cognitive resources respond by categorizing games in a convenient manner. Such categorization may be endogenous (Samuelson, 2001; Fryer and Jackson, 2008; Mengel, 2012), but we consider exogenous categorization only. Once players have categorized the games in some cognitively convenient manner, they may choose a range of simplifying heuristics that economize on scarce cognitive resources. For instance, for any category in a partition of a set of games, they may act as if they faced the average payoffs, or have the average beliefs for the category. This leads to the definition of a range of useful concepts such as *payoff bundling* and *belief bundling*.

The empirical results show that greater complexity (larger set of games and lower information) reduces the probability of playing a Nash equilibrium. Although the degree of complexity in even the most complex situation in the experiments is typically low, and the number of rounds at 100 is very large, the results are sobering for classical Nash equilibrium play. Framing effects and learning spillovers across games are observed; play in strategically equivalent games is similar; transparency of information on payoffs rather than past play of the opponent is more conducive to the Nash equilibrium being played; subjects with higher cognitive ability are more likely to play a Nash equilibrium, and action, or best-response bundling explains the data better than belief bundling.

Section 15.9 briefly illustrates two polar positions on learning.² The first position takes the view that under fairly strong restrictions, Bayesian learning, coupled with a technical condition, the absolute continuity condition, ensures that the limiting play converges to a Nash equilibrium (Kalai and Lehrer, 1993). Others may feel that this framework imposes excessive rationality requirements that are unlikely to be supported by the empirical evidence. Indeed, the evidence in this chapter points in this direction. The second position denies the possibility of such learning (Foster and Young, 2001); this can be shown, for instance, in cases where there is enough ex-ante uncertainty about the payoffs of others.

We have considered a range of models above that incorporate the idea of regret from forgone payoffs, for instance, EWA and impulse matching. The idea is that the regret arising from a strategy is the average payoff forgone from playing other strategies. One may then play pure strategies with probabilities that reflect their relative regret; this is known as *regret matching*. It can be shown that in limiting time, regret matching leads to behavioral strategies that have no associated regret, and the empirical frequency of play converges to the set of coarse correlated equilibria (Hart and Mas-Colell, 2000). This set of equilibria is typically much larger than the set of Nash equilibria, although in very special cases, the two concepts make identical predictions.

15.2 Reinforcement learning

Reinforcement learning, sometimes known as *stimulus response learning* or *rote learning*, is based on the following idea. Strategies that have resulted in a greater payoff in the past are more likely to be played in the future. This form of learning imposes very limited rationality and information

² For a more detailed account of these issues, see Young (2004).

requirements on the parts of players. For instance, players do not require information about the identities and strategies of other players, or even knowledge of the nature of strategic interaction between themselves and other players. Thus, players do not experience regret in these models because they do not have knowledge of the underlying game.

Models of reinforcement learning first appeared in the mathematical psychology literature (Bush and Mostellar, 1951; Suppes and Atkinson 1960). Arthur (1991, 1993) gives an early application of these models to problems in economics. Roth and Erev (1995) and Erev and Roth (1998) modified the basic model to better explain the evidence arising from a range of experimental games. Reinforcement learning appears to fit well human and non-human subjects in very simple games, particularly when players do not know the structure of the game. However, in games where the structure is better known to the players, reinforcement learning does not fit very well. Furthermore, the predicted speed of learning under reinforcement learning is quite sluggish relative to the speed of learning that is observed in experiments.

The basic features of the model that were emphasized in the original papers by Bush and Mostellar (1955), Luce (1959), and Suppes and Atkinson (1960) are as follows.

1. *Law of effect*: If a strategy results in a higher payoff (*stimulus*), then it is more likely to be played in the future (*response*), an idea that goes back to Thorndike (1911). Hence, the basic models in this literature are sometimes termed as *stimulus–response learning*.
2. *Choice is probabilistic*: At each point in time, players play their pure strategies in a probabilistic manner. More successful strategies that have resulted in a higher payoff in the past are more likely to be played in the present (*law of effect*).
3. *Power law of practice*: Players' learning curves are steep initially, but then gradually flatten out. However, the evidence on this is not conclusive.

15.2.1 The one-parameter reinforcement learning model

The one-parameter reinforcement learning model was used extensively until the mid 1990s (Bush and Mostellar, 1955; Arthur, 1991, 1993). Suppose that the set of players is given by $N = \{1, 2, \dots, n\}$. The pure strategy set of player i is given by $S_i = (s_{i1}, s_{i2}, \dots, s_{in_i})$. Players only play pure strategies. Time $t = 1, 2, \dots$ is discrete. Each player has certain *propensities*, sometimes also termed as *attractions*, to play each of the available pure strategies. For player i , the propensity to play strategy $s_{ij} \in S_i$ at time t is given by $q_t(s_{ij})$. We assume that the initial propensity to play any pure strategy is identical, i.e., for any $s_{ij}, s_{ik} \in S_i$ we have $q_1(s_{ij}) = q_1(s_{ik})$.

Players are not required to know the structure of the game they are playing, or even if they are playing a game against an opponent. Hence, unlike classical game theory, payoffs are not defined as mappings from the Cartesian products of strategy sets of players to the set of real numbers. For player i , the payoff function is simply $\pi_i : S_i \rightarrow \mathbb{R}$.

For any player i , we define a *reinforcement function*, $R : \mathbb{R} \rightarrow \mathbb{R}$,

$$R(\pi_i(s_{ij})) = \pi_i(s_{ij}) - \bar{\pi}, s_{ij} \in S_i, \quad (15.1)$$

where $\bar{\pi}$ is the minimum payoff.

Propensities are updated over time as follows. For $t = 1, 2, \dots$

$$q_{t+1}(s_{ij}) = \begin{cases} q_t(s_{ij}) + R(\pi_i(s_{ij})) & \text{if } s_{ij} \text{ is played at time } t \\ q_t(s_{ij}) & \text{otherwise} \end{cases}. \quad (15.2)$$

The choice of $\bar{\pi}$ as the minimum payoff ensures that propensities never become negative. The final step is the translation of propensities into probabilities of playing the various pure strategies. Let $p_t(s_{ij})$ be the probability that, at time t , player i plays strategy $s_{ij} \in S_i$. Then the probabilistic choice rule is given by

$$p_t(s_{ij}) = \frac{q_t(s_{ij})}{Q_{it}}, \text{ where } Q_{it} = \sum_{k=1}^{n_i} q_t(s_{ik}). \quad (15.3)$$

Thus, the probability with which a strategy is played equals its propensity, relative to the aggregate propensity, Q_{it} . Even if $Q_{it} \neq 1$, we have a well-defined probability measure because $\sum_{j=1}^{n_i} p_t(s_{ij}) = 1$. In applied work, sometimes the following one-parameter exponential rule for computing probabilities is used (see Appendix B)

$$p_t(s_{ij}) = \frac{e^{\lambda q_t(s_{ij})}}{\sum_{k=1}^{n_i} e^{\lambda q_t(s_{ik})}}, \quad (15.4)$$

where the single parameter λ reflects the sensitivity of probabilities to reinforcements. For instance, when $n_i = 2$ we can write (15.4) as

$$p_t(s_{i1}) = \frac{1}{1 + e^{\lambda[q_t(s_{i2}) - q_t(s_{i1})]}}. \quad (15.5)$$

Thus, if $q_t(s_{i2}) - q_t(s_{i1}) > 0$, then a higher λ reduces $p_t(s_{i1})$. If $q_t(s_{i2}) - q_t(s_{i1}) < 0$, then a higher λ increases $p_t(s_{i1})$. Thus, a higher λ magnifies the probabilities of strategies whose propensities are also higher.

From (15.2), (15.3), we have $\frac{d}{dt}(Q_{it}) > 0$. Also, from (15.2), if strategy s_{ij} is played at time t , we can write

$$Q_{it+1} = Q_{it} + R(\pi_i(s_{ij})), t = 1, 2 \dots \quad (15.6)$$

One interpretation of the probabilistic choice rule in (15.4) is that players continually experiment. The actual choice of player i in any period, which is to play one of the pure strategies in S_i , is made by using some randomizing device that reflects the choice probabilities in (15.3). We need to fix the initial value of aggregate propensity, Q_{i1} , for each player, in order to determine the initial probabilities of play in (15.3). We assume that it is identical for all players, i.e.,

$$Q_{i1} = \bar{Q}, \forall i \in N. \quad (15.7)$$

$\bar{Q}$ is the only parameter of the model. For this reason, we call this version the *one-parameter reinforcement learning model*.

We can check that the one-parameter model satisfies all the features, 1–3, of the reinforcement learning model. From (15.2), positive reinforcement is given at time t to the action that is played in that time period. From (15.1), the magnitude of reinforcement is directly proportional to the payoff associated with a strategy (law of effect). From (15.3), this directly affects the probability of playing that strategy (choice is probabilistic). To demonstrate the third feature, the power law of practice, suppose that s_{ij} is the action played at time t . Then substituting (15.2) in (15.3), we get that

$$\frac{dp_t(s_{ij})}{d(R(\pi_i(s_{ij})))} = \frac{1}{Q_{it}} > 0; \frac{d}{dt} \left(\frac{dp_t(s_{ij})}{R(\pi_i(s_{ij}))} \right) = \frac{-1}{Q_{it}^2} \frac{d}{dt}(Q_{it}) < 0. \quad (15.8)$$

Thus, the impact of a given reinforcement on the probability of playing s_{ij} is positive, but diminishes with time. So, the learning curve gets flatter with the passage of time, as required by the power law of practice.

15.2.2 The three-parameter model of reinforcement learning

Motivated by psychological evidence,³ Roth and Erev (1995) and Erev and Roth (1998) propose a *three-parameter reinforcement learning model*. They essentially add two psychological features to the one-parameter model.

1. *Experimentation*: The probabilistic nature of choice in the one-parameter model already ensures that individuals are not locked-into choices they make at any time. Experimentation ensures that players experiment with choices “similar” to those that were successful in the past.
2. *Recency*: Recent events and experiences influence behavior to a greater extent than those in the more distant past.

These two features, *experimentation* and *recency*, are incorporated by modifying the updating of propensities in (15.2), as follows. Suppose that given any strategy, s_{ij} , we can order the remaining strategies in the set S_i in a linear order such that the most adjacent strategies to s_{ij} are most “similar” to s_{ij} .⁴ This ensures that s_{ij-1} and s_{ij+1} are more similar to s_{ij} than s_{ij-2} and s_{ij+2} , and so on. Suppose that strategy s_{ij} is played at time t . Then player i updates his propensities in the following manner.

$$q_{t+1}(s_{ik}) = \begin{cases} (1 - \phi) q_t(s_{ik}) + (1 - \varepsilon) R(\pi_i(s_{ij})) & k = j \\ (1 - \phi) q_t(s_{ik}) + \frac{\varepsilon}{2} R(\pi_i(s_{ij})) & k = j - 1, j + 1 \\ (1 - \phi) q_t(s_{ik}) & \text{for } k \neq j \text{ and } k \neq j - 1, j + 1 \end{cases}, \quad (15.9)$$

where $R(\pi_i(s_{ij}))$ is defined in (15.1), and $\varepsilon > 0$ is a small number. $\phi > 0$ is a “forgetting parameter” which gradually reduces the influence of more distant experiences in the past. The first row in (15.9) shows that the strategy s_{ij} , which is most recently played, is reinforced in a manner similar to (15.2). However, reinforcement is reduced by the experimentation parameter ε . The second row shows that “similar” strategies in the immediate neighborhood of s_{ij} are also reinforced, although not to the same extent as s_{ij} . Finally, all other strategies are not reinforced at all (third row). Hence, the model has three parameters, $\phi, \varepsilon, \bar{Q}$ (where $\bar{Q}$ is defined in (15.7)).

When strategies cannot be ordered by similarity, then the three-parameter model imposes the following reinforcement of propensities following a play of s_{ij} at time t (recall that player i has n_i pure strategies).

$$q_{t+1}(s_{ik}) = \begin{cases} (1 - \phi) q_t(s_{ik}) + (1 - \varepsilon) R(\pi_i(s_{ij})) & k = j \\ (1 - \phi) q_t(s_{ik}) + \frac{\varepsilon}{(n_i - 1)} R(\pi_i(s_{ij})) & k \neq j \end{cases}. \quad (15.10)$$

³ See, for instance, Guttman and Kalish (1956) and Brown et al. (1958).

⁴ Similarity is not defined in the paper; rather, it appears to be a matter of judgment about the various strategies, exercised by the player. For instance, a bunch of strategies may be similar because they could be recognized as cooperative strategies, or relying on a some common notion of justice or fairness.

15.2.3 Some other variants of reinforcement learning

Sarin and Vahid (1999, 2001) propose the injection of greater rationality within the reinforcement learning framework. When choosing their strategies, players face some uncertainty about the future states of the world. Let Ψ be the set of possible states of the world, which may be thought of as some compact subset of $\mathbb{R}$. Payoffs are random and given by the payoff function $\pi : S_i \times \Psi \rightarrow \mathbb{R}$. Suppose that strategy s_{ij} is played at time t and $s_{ik} \in S_i$. Then, in Sarin and Vahid (1999), for $\omega \in \Psi$ the player assigns an assessment that the payoff from the strategy s_{ik} at time $t + 1$ is given by

$$u_{t+1}(s_{ik}) = \begin{cases} (1 - \lambda) u_t(s_{ik}) + \lambda \pi(s_{ik}, \omega) & k = j \\ u q_t(s_{ik}) & k \neq j \end{cases} \quad (15.11)$$

where $\lambda \in (0, 1)$, the time-invariant learning parameter is the only parameter in the model in addition to the initial assessments. The higher is λ , the greater is the responsiveness of assessments to observed payoffs. Sarin and Vahid (1999) dispense with the construction of probabilities. They directly assume that the player optimizes and chooses an action $s_i^* \in S_i$ that maximizes the assessment at time t . Sarin and Vahid (2001) find that when the results are averaged across subjects in games that have a unique mixed strategy Nash equilibrium, their model fits better overall as compared to the three-parameter Roth and Erev (1998) model.

Duffy and Feltovich (1999) consider the possibility that players learn not just from their own experiences, but also from the experiences of others. Suppose that there are only two players, A and B , with identical pure strategy sets. They respectively play strategies s_{Aj} and s_{Bl} in period t . Player A updates his propensities in the following manner that modifies the one-parameter reinforcement learning model to take account of the experience of player B .

$$q_{t+1}(s_{Ak}) = \begin{cases} q_t(s_{Ak}) + R(\pi_A(s_{Aj})) & \text{if } k = j \\ q_t(s_{Ak}) + \beta R(\pi_B(s_{Bl})) & \text{if } k = l \\ q_t(s_{Ak}) & \text{otherwise} \end{cases}, \beta \geq 0. \quad (15.12)$$

Players update strategies based on their own experience (first row) and the experience of the other player (second row). Using $\beta = 0.5$, Duffy and Feltovich (1999) find that their model gives a good fit to the data for the best shot and the ultimatum games used in Roth and Erev (1995).

The reinforcement model is also used in models of *aspiration adaptation*. Essentially, players have aspiration levels, as envisaged in Herbert Simon's notion of satisficing. Denote the aspiration level of player i at time t by a_{it} , where the initial level of aspiration a_{i0} is given for each player. Subsequent to the play of a strategy $s_{ik} \in S_i$, the aspiration levels are updated gradually in the following manner,

$$a_{it+1} = \gamma a_{it} + (1 - \gamma) \pi(s_{ik}), \gamma \in (0, 1).$$

Hence, at time $t + 1$, aspirations are a convex combination of current aspirations and the payoff received at time t from playing s_{ik} . In *aspiration-based reinforcement learning*, the following condition is imposed on propensities.

$$\text{If } \pi(s_{ik}) \geq a_{it}, \text{ then } q_{t+1}(s_{ik}) \geq q_t(s_{ik}) \text{ and } p_{t+1}(s_{ik}) \geq p_t(s_{ik}). \quad (15.13)$$

Thus, (15.13) requires that if at time t , the payoff from some strategy is greater than the current aspiration level, then the propensity and the probability of playing that strategy increases. A variety of functional forms can be used to implement (15.13).⁵ The propensities are then converted into probabilities of playing various strategies as in, say, (15.3). For evidence in support of aspiration-based learning models, see Bereby-Meyer and Erev (1998) and Huck et al. (2007).

15.2.4 Optimality of reinforcement learning in a stationary environment

Consider a single player who plays a game against nature in a stationary environment. The pure strategy set of the player is given by the finite set $A = (a_1, a_2, \dots, a_{n_A})$. Mixed strategies over this set are denoted by ΔA . Nature has the finite strategy set $B = (b_1, b_2, \dots, b_{n_B})$. Let mixed strategies over B be given by ΔB . The payoff function of the player is given by $\pi : A \times B \rightarrow \mathbb{R}$.

Consider the following updating rule for propensities

$$q_{t+1}(a_i) = \begin{cases} q_t(a_i) + \pi(a_i, b_j) & \text{if } a_i \text{ is played at time } t \\ q_t(a_i) & \text{otherwise} \end{cases}, \quad (15.14)$$

where $a_i \in A$ is the action chosen by the player and $b_j \in B$ is the action chosen by nature in period t . This updating rule is known as the *cumulative payoff matching rule* (CPM), or sometimes also as the *cumulative proportional reinforcement rule*. At time t , denote the tuple of strategies played by the player and by nature, by (a^t, b^t) . Let $t = 1$ be the initial period. Then, the history of play at the end of time period t , is given by

$$h^t = [(a^1, b^1), (a^2, b^2), \dots, (a^t, b^t)]. \quad (15.15)$$

Under reinforcement learning, players use (15.14), (15.15) as inputs into the determination of probabilities with which they play their various pure strategies, given in (15.3).

Suppose that in every period, nature plays according to some time-invariant probability distribution $p^b = (p^b(b_1), \dots, p^b(b_{n_B})) \in \Delta B$. The best response of the player to p^b , $BR(p^b)$, can be found as follows

$$BR(p^b) \in \arg \max_{a_i \in A} \sum_{j=1}^{n_B} p^b(b_j) \pi(a_i, b_j). \quad (15.16)$$

Definition 15.18 (*Optimal behavioral strategy*) A behavioral strategy of the player at time t , $s_A^t : p^b \rightarrow A$ is optimal against the stationary distribution p^b if the probability with which a_i is played under s_A^t goes to zero (almost surely) for all $a_i \notin BR(p^b)$.

The following result has been either conjectured, or shown to be true in several contributions.⁶

Proposition 15.21 (Laslier et al., 2001): For finite action sets, A, B , if players use reinforcement learning and update their propensities using cumulative payoff matching, then, in the limit, such a behavioral strategy is optimal (belongs to the set of strict best responses) against a stationary distribution of play by nature.

⁵ See for instance, Karandikar et al. (1998), Borghers and Sarin (2000), and Bendor et al. (2001).

⁶ See for instance, Arthur (1993), Rustichini (1999), Laslier et al. (2001), and Hopkins and Posch (2005).

In addition, Hopkins and Posch (2005) showed that the limiting points in Proposition 15.21 do not converge to fixed points that are not asymptotically stable under replicator dynamics. Players who use reinforcement learning behave as automata. They mechanically convert inputs, from history and their propensities, into outputs, i.e., probability distributions over their actions. However, when they use cumulative payoff matching, then in a stationary environment, reinforcement learning converges to a best response.

15.2.5 *Empirical evidence on reinforcement learning*

Reinforcement learning often correctly predicts the direction of learning; this is not surprising because it reinforces propensities of more successful strategies, relatively more. However, the predicted speed of learning is too slow relative to experimental results from a variety of games, such as the ultimatum game and the public goods game; see, Roth and Erev (1995).

McAllister (1991) used the reinforcement learning model to test the predictions of the *weak link game* of Van Huyck et al. (1990) (see Part 4 of the book). They find that speed of learning predicted by the reinforcement model is too slow relative to actual learning. If the reinforcement model is modified to include reinforcements for forgone payoffs (see this modification in the discussion of the EWA model below), then the predicted speed of learning is higher and more in conformity with the evidence.

Figure 15.1 shows the results of fitting the data from 5 simple classes of 2×2 games in Erev and Roth (1998); a total of 12 games are fitted to the data. The games are shown along the left margin. The strategies of the row player are A1 and B1, and the strategies of the column player are A2 and B2. The alphabets for each game (e.g., S & A2, M & L) are acronyms for the names of authors who originated the corresponding games. Four of the five classes of games had a unique mixed strategy Nash equilibrium that is shown in the first column in Figure 15.1. For instance, the mixed strategy equilibrium for the game in the first row (S & A2) is for the row player to play A1 with probability $1/3$ and for the column player to play A2 with probability $5/6$. However, subjects did not see the game. They pressed keys and the probability of a correct answer is determined by the payoff matrix in a manner that we do not go into here; this accounts for the variation in the mixed strategy predictions across the rounds.

Columns 2 and 3 in Figure 15.1 show the predictions of the one parameter (Roth and Erev, 1995) and the three parameter (Erev and Roth, 1998) models of reinforcement learning. The final column shows the predictions of a model of fictitious play (see Section 15.3 below). The accuracy of predictions is judged by calculating the mean square deviations (MSD) of the predictions from the actual data. The predictions are derived by populating the models with imaginary agents, and endowing them with the model parameters. Furthermore, the parameters are chosen using a grid search method to minimize MSD. The MSD were calculated for each block of data, and then averaged across the blocks. Among any two models, the one with the smaller value of MSD is to be preferred.

The results are shown in Table 15.1. Columns 2–13 correspond to one of 12 games and each row corresponds to a particular model of learning. The numbers in each results column are the values of $\text{MSD} \times 100$ for the corresponding game and the last column reports $\text{MSD} \times 100$ aggregated across all 12 games for each learning model. Other than random choice and equilibrium choice, the remaining learning models are presented in terms of the best fit across all games and the best fit in each game. Across all games, the equilibrium predictions (corresponding to mixed strategy Nash) have an MSD of 3.87, which is worse than random choice (1.87). When compared across individual games, random does better than equilibrium in five games and ties in one game.




Figure 15.1 Results of fitting the data from 5 simple classes of 2×2 games.

Source: Reprinted from *Artificial Intelligence*, 171, 423–8. Erev, I. and Roth, A. "Multi-agent learning and the descriptive value of simple models." Copyright © 2007, with permission from Elsevier.

Overall, the basic one-parameter model of reinforcement learning does better than equilibrium (MSD of 1 as compared to 3.57). However, it does not explain the late movement in many games towards equilibrium. The one-parameter model shows a greater sluggishness in response to the play of the opponent relative to the actual data. The three-parameter model that best fits the data across all games gives the following values of the parameters: $Q(1) = 9$, $\phi = 0.1$ and $\varepsilon = 0.2$. This model outperforms all the other models used. The authors argue that this is not due to the addition of extra parameters alone, but rather that the three-parameter model captures something fundamental about the nature of human learning. The three-parameter model, for instance, outperforms the one-parameter model even when the one-parameter model uses game-specific 12 parameters.

Choi et al. (2009) provide evidence from 401(k) savings plans in the US to show that investors appear to follow *naïve reinforcement learning*, hence confirming experimental findings. Investors who observe a higher than average return, or a relatively low variance return increase their savings rate in this investment vehicle relatively more. This finding cannot be explained by investor fixed effects, or time-varying investor-level heterogeneity.

Table 15.1 A comparison of the results of alternative learning models.

Game model	S&A8	S&A2	S&A3u	S&A3k	S&A3n	M&L	On	R&BI5	R&BI0	Oc9	Oc4	Ocl	Mean over all games
Random	1.08	2.04	2.53	1.46	2.11	2.46	2.19	1.07	1.38	3.88	1.78	0.45	1.87
Equilibrium	6.92	7.18	7.27	7.56	6.14	2.11	0.14	0.45	1.03	2.22	1.37	0.45	3.57
Basic reinforcement:													
best fit (1 parameter)	0.16	0.30	0.31	0.11	0.57	2.27	1.81	0.98	0.73	2.71	1.54	0.48	1.00
by game (12 parameters)	0.07	0.24	0.14	0.10	0.41	1.89	0.33	0.50	0.16	2.34	1.54	0.41	0.68
prediction	0.16	0.30	0.31	0.11	0.57	0.27	1.81	0.98	0.86	2.71	1.64	0.48	1.02
RE:													
best fit (3 parameters)	0.38	0.18	0.12	0.07	0.31	1.24	0.72	0.65	0.33	1.54	1.09	0.48	0.59
by game (36 parameters)	0.05	0.10	0.04	0.05	0.25	0.21	0.32	0.35	0.11	1.34	0.99	0.37	0.35
prediction	0.67	0.26	0.19	0.09	0.39	1.24	0.87	0.83	0.48	1.54	1.17	0.51	0.69
FP-like:													
best fit (4 parameters)	0.34	0.20	0.16	0.09	0.37	1.26	1.05	0.71	0.44	2.04	1.48	0.42	0.71
by game(48 parameters)	0.05	0.09	0.04	0.03	0.29	0.45	0.04	0.14	0.17	1.70	1.19	0.28	0.37
prediction	0.77	0.44	0.24	0.09	0.37	1.54	1.24	0.71	0.44	2.10	1.65	0.45	0.84

Source: Erev and Roth (1998).

Van Huyck et al. (2007) consider an *order statistic game* with five players; see Part 4 for the details. The strategy of each player is a choice of number $x_i \in [0, 1]$. Denote the strategy choice of players other than i by the vector x_{-i} , and the strategy profile by $x = (x_i, x_{-i})$. Let the median strategy choice be x_m . Then, the payoff of the i th player is given by

$$\pi_i(x_i, x_{-i}) = 0.5 - |x_i - \lambda x_m(1 - x_m)|, \lambda > 0. \quad (15.17)$$

By inspection, the optimal choice of the i th player is $x_i^* = \lambda x_m(1 - x_m)$. In one of the games considered by the authors, $\lambda = 2.44$. For this game, the symmetric Nash equilibrium is given by the solution to $x^N = 2.44x^N(1 - x^N)$ for all i . There are two symmetric Nash equilibria, $x^N = 0$ and $x^N = 0.59$. In one treatment, each experimental subject was informed that there were four other players and the game would be repeated over 75 periods. In particular, the payoff functions of the other players were not revealed. Play settles down to the interior Nash equilibrium in a zig-zag pattern after about 20 periods.

Van Huyck et al. (2007) then ask if this is in accord with the predictions of reinforcement learning. However, the learning dynamics induced by reinforcement learning are too slow. If the parameters in the model are altered to speed up learning, then it gives rise to other undesirable side-effects such as the play of dominated strategies.

The information requirements in models of reinforcement learning are limited. Players only need information about their own past payoffs. However, Mookerjee and Sopher (1994) found that the experimental results depend on extra information, in addition to the data on players' own past payoffs. A similar result is found in Rapoport and Erev (1998) and Battalio et al. (2001); the latter found that the provision of extra information can speed up the learning process substantially.

15.3 Belief-based models of learning

How do players form beliefs about others and about the environment around them, and how do these beliefs evolve over time? In the literature, the general class of models that address this question is known as *belief-based models of learning*. In models of reinforcement learning, players are not required to be informed of the underlying game and they only play a *better response*, but not a *best response*. By contrast, in belief-based models, there is a greater degree of rationality.

Players in belief-based models know the game that they are playing, and conditional on their beliefs about the opponent, they play a *best response*. However, they have limited rationality in the sense that they do not realize that their opponents are also engaged in a similar exercise. Beliefs are typically formed by observing the past strategies of the opponent. For instance, in the Cournot model, players play a best response to the previous period's output of the opponent, believing that the opponent will choose the same output this period. However, the general class of belief-based models allows for a sophisticated weighting of past actions of the opponent in order to form the relevant beliefs.

An important subclass of belief-based models is the model of *fictitious play*. Most results in the literature on fictitious play are available only for two-player simultaneous move games, so we restrict our analysis to such games. In models of fictitious play, each of the two players observes the past frequency of play of different actions by the opponent. This data is then used to compute the probabilities with which the opponent is expected to play various pure strategies; opponents are assumed to draw their strategies from a stationary distribution. Once such probabilities are

computed, each player behaves in a myopic manner and maximizes the current payoff. Thus, players take into account the entire history of play (unlimited memory), but assume a stationary environment.

An interesting question is whether the computed frequencies of play of the opponent under fictitious play converge to some limiting distribution. Also, does the play converge to a Nash equilibrium, conditional on this limiting distribution? The answer to both questions is in the affirmative for two-player games. Indeed, the original motivation to study fictitious play (Brown, 1951; Robinson, 1951) was to find an alternative method of computing a Nash equilibrium.

If the strategy profile in the game converges to a Nash equilibrium (as we shall see below), then one could envisage that players engage in a mental exercise prior to the game. As part of this mental exercise, each player goes through the steps of computing a limiting distribution of play of the opponent's actions. Subsequently, when each player is given a chance to play the game, they can straightaway play the Nash equilibrium, using the limiting distribution as the initial expectation over the opponent's actions. Hence, the name, *fictitious play*.

Denote the two-player simultaneous move game by $G(2, S_1, S_2, \pi_1, \pi_2)$, where π_1, π_2 are von Neumann–Morgenstern utility functions. The set of players is $N = \{1, 2\}$, the pure strategy sets of players 1 and 2 are, respectively, given by $S_1 = (s_{11}, s_{12}, \dots, s_{1n_1})$, $S_2 = (s_{21}, s_{22}, \dots, s_{2n_2})$ and the payoffs by $\pi_i : S_1 \times S_2 \rightarrow \mathbb{R}$. Time is discrete, $t = 0, 1, 2, \dots$. Each player $i \in N$ associates, at time t , a weight $w_t^i(s_{jk})$ to the pure strategy s_{jk} of the opponent $j \in N$, where $j \neq i$ and $s_{jk} \in S_j$. Denote the vector of these weights by

$$w_t^i = (w_t^i(s_{j1}), \dots, w_t^i(s_{jn_j})).$$

There is an initial set of non-negative weights w_0^i , $i \in N$ that is taken as given and $\sum_{k=1}^{n_j} w_t^i(s_{jk}) > 0$, $j \neq i$. Over time, the weights are adjusted as follows.

$$w_{t+1}^i(s_{jk}) = \begin{cases} w_t^i(s_{jk}) + 1 & \text{if } s_{jk} \text{ is played at time } t \\ w_t^i(s_{jk}) & \text{otherwise} \end{cases}; \quad i, j \in N, j \neq i \text{ and } s_{jk} \in S_j. \quad (15.18)$$

Hence, if a pure strategy is played by the opponent, then its weight is augmented by 1 unit in the next period, but the weight on all other pure strategies is unchanged.

Each player uses the weights, constructed as in (15.18), to assign probabilities to the opponent's pure strategies. Define

$$\bar{w}_t^i = \sum_{k=1}^{n_j} w_t^i(s_{jk}). \quad (15.19)$$

Denote by $x_t^i(s_{jk})$, the probability with which player i expects player j ($j \neq i$) to play the pure strategy $s_{jk} \in S_j$. Then, under fictitious play

$$x_t^i(s_{jk}) = \frac{w_t^i(s_{jk})}{\bar{w}_t^i} \geq 0. \quad (15.20)$$

From (15.20), we have $\sum_{k=1}^{n_j} x_t^i(s_{jk}) = 1$. Denote the vector of probabilities with which player i expects player j to play all pure strategies in the set S_j by

$$x_t^i = (x_t^i(s_{j1}), x_t^i(s_{j2}), \dots, x_t^i(s_{jn_j})); \quad i, j \in N, j \neq i \text{ and } s_{jk} \in S_j. \quad (15.21)$$

Finally, fictitious play requires that each player chooses a pure strategy $s_{ik}^* \in S_i$, such that it is a best response to beliefs, x_t^i . Thus,

$$s_{ik}^* \in BR(x_t^i) \equiv \{s_i \in S_i : \pi(s_i, x_t^i) \geq \pi(\hat{s}_i, x_t^i), \forall \hat{s}_i \in S_i\}, i = 1, 2. \quad (15.22)$$

Under fictitious play, players do not try to influence the opponent's action, nor do they take into account the strategic effect of their actions on others. One feature of belief-based models is that the best reply functions of players may be non-differentiable. An example is the battle-of-sexes game. In order to smooth the best response functions, one can extend belief-based models to noisy play in many ways. In Boylan and El-Gamal (1993), with probability ϵ , players randomly choose one of their actions and with probability $1 - \epsilon$, they play their best response, as described above.

Let us now consider a version of fictitious play known as *smooth fictitious play* by introducing a small additive shock to the payoff of players, so the payoff of player i becomes

$$\pi_i(s_i, \sigma_{-i}) + \eta_i, \quad (15.23)$$

where η_i is a random shock. Suppose that the cumulative distribution of the shock is given by

$$\ln \Pr(\eta_i < y) = -e^{-\lambda_i y}, \quad (15.24)$$

where λ_i is an individual specific parameter. Then it can be shown that the best response of the player is given by the *logistic* or *quantal response function* (see Appendix B). Suppose that the probability that player i plays strategy s_{ik} when the other players are expected to play a mixed strategy σ_{-i} is given by $p_t^i(s_{ik} | \sigma_{-i})$. Then, under (15.23), (15.24), the best response of player i is given by the quantal response function

$$p_t^i(s_{ik} | \sigma_{-i}) = \frac{e^{-\lambda_i \pi_i(s_{ik}, \sigma_{-i})}}{\sum_{j=1}^{n_i} e^{-\lambda_i \pi_i(s_{ij}, \sigma_{-i})}}, \quad (15.25)$$

which is completely smooth.

In Section 15.2.4, we showed that in special circumstances, reinforcement learning is consistent with fully rational behavior. It can also be shown that fictitious play is consistent with fully rational optimizing behavior under the following two conditions (Fudenberg and Levine, 1998). (1) The initial beliefs of a player about an opponent's actions are given by the Dirichlet distribution (which is a generalization of the beta distribution). (2) Each player believes that in each subsequent period, the opponent's play is a sequence of iid multinomial random variables.

Example 15.7 In 2×2 games, if the probabilities of the play of others are given in (15.20) and a strict Nash equilibrium is played, then Fudenberg and Kreps (1990) showed that under fictitious play, this Nash equilibrium is played for all time periods. To illustrate this result, consider the symmetric 2×2 coordination stage-game, shown in Figure 15.2, that is played in every period.

This game has two strict Nash equilibria: (s_1, s_1) and (s_2, s_2) . Denote the row and the column players by players 1 and 2 respectively. Suppose that player 2 plays action s_1 with probability q and s_2 with probability $1 - q$. Then player 1's best response is to play action s_2 if $9q \leq 7q + 8(1 - q)$ or $q \leq 0.8$.

	s_1	s_2
s_1	9, 9	0, 7
s_2	7, 0	8, 8

Figure 15.2 A symmetric 2×2 coordination game.

Suppose that initial weights are identical and given by the ordered pair $w_0^1 = w_0^2 = (6, 4)$. Using (15.20), the probabilities assigned by each player to the actions (s_1, s_2) played by the other player are symmetrically given by the ordered pair $x_0^1 = x_0^2 = (0.6, 0.4)$. Since $0.6 < 0.8$, the best response of each player is to play s_2 . Using (15.18), the weights in period 1 are $w_1^1 = w_1^2 = (6, 5)$. From (15.20) the opponent is expected to play each of the pure strategies with probabilities $x_1^1 = x_1^2 = (0.545, 0.455)$. Since $0.545 < 0.8$, the best response by each player is again to play s_2 . The calculations for the next two periods, each eliciting a best response of s_2 from each player are: $w_2^1 = w_2^2 = (6, 6)$, $x_2^1 = x_2^2 = (0.5, 0.5)$ and $w_3^1 = w_3^2 = (6, 7)$, $x_3^1 = x_3^2 = (0.46, 0.54)$. This is a strictly decreasing sequence, ensuring that s_2 is played in every period.

Do the probabilities given in (15.21) converge? Robinson (1951) showed that they converge for the case of zero-sum games. Miyasawa (1961) and Monderer and Shapley (1996) show that convergence takes place for generic 2×2 games. We state this result without proof.

Proposition 15.22 Consider the game $G(2, S_1, S_2, \pi_1, \pi_2)$. Suppose that $n_1 = n_2 = 2$, initial weights are w_0^i , $i = 1, 2$, and the sequence $(x_t^1, x_t^2)_0^\infty$ is induced by (15.18)–(15.22). Then $\lim_{t \rightarrow \infty} x_t^i = x^{i*} = (a, b) \in [0, 1] \times [0, 1]$, $i = 1, 2$.

For Example 15.7, the vector of probabilities converges to the tuple $x^{1*} = x^{2*} = (0, 1)$. The best response of each player is to play s_2 , so the limiting outcome is the Nash equilibrium (s_2, s_2) . The next proposition shows that this is not a coincidence.

Proposition 15.23 Consider the game $G(2, S_1, S_2, \pi_1, \pi_2)$. Suppose that $n_1 = n_2 = 2$ and initial weights are given by w_0^i , $i = 1, 2$. Let the sequence $(x_t^1, x_t^2)_0^\infty$ be induced by (15.18)–(15.22). Suppose that $\lim_{t \rightarrow \infty} x_t^i = x^{i*} = (a, b) \in [0, 1] \times [0, 1]$. Now define a mixed strategy for each player as follows:

$$\sigma_i^* = x^{i*}; i = 1, 2.$$

Then (σ_1^*, σ_2^*) is a Nash equilibrium of $G(2; S_1, S_2; \pi_1, \pi_2)$.

Proof: Suppose by way of contradiction that (σ_1^*, σ_2^*) is not a Nash equilibrium. Let there be some $s_1^* \in S_1$ in the support of σ_1^* and another $\hat{s} \in S_1$ that is not in the support of σ_1^* . Then we have

$$\pi(s_1^*, \sigma_2^*) < \pi(\hat{s}, \sigma_2^*). \quad (15.26)$$

Since the sequence $\{x_t^1, x_t^2\}$ converges to (σ_1^*, σ_2^*) and π is continuous, along the path of play, there is a neighborhood of σ_1^* such that (15.26) holds for all σ_1 in that neighborhood. Alternatively, there exists some integer value τ such that for all $t \geq \tau$, the condition (15.26) holds. Thus, s_1^* is

never played by player 1 subsequent to time period τ . Hence, using (15.18), its decision weight is never updated, so

$$w_t^2(s_1^*) = w_\tau^2(s_1^*), \forall t \geq \tau. \quad (15.27)$$

From (15.19) we get that

$$\bar{w}_t^2 = \bar{w}_\tau^2 + (t - \tau), \forall t \geq \tau. \quad (15.28)$$

Using, (15.20) we expect that player 2 assigns a decreasing sequence of probabilities to player 1 playing s_1^* ever again. Indeed in the limit, such a probability may vanish as we now see. Since $x_t^2(s_1^*) = \frac{w_t^2(s_1^*)}{\bar{w}_t^2}$ and using (15.27), (15.28) we have that

$$\sigma_1^*(s_1^*) = \lim_{t \rightarrow \infty} x_t^2(s_1^*) = \lim_{t \rightarrow \infty} \frac{w_\tau^2(s_1^*)}{\bar{w}_\tau^2 + (t - \tau)} = 0.$$

But if $\sigma_1^*(s_1^*) = 0$, then it contradicts the condition that s_1^* belongs to the support of σ_1^* . An analogous proof applies if, in this symmetric game, we had started off with player 2 instead. ■

Propositions 15.22, 15.23 apply to special classes of 2×2 games. Shapley (1964) showed that for 3×3 games, the sequence $\{x_t^1, x_t^2\}$ induced by (15.18)–(15.22) might not converge. Consider, for instance the symmetric two-player game shown in Figure 15.3, originally due to Shapley (1964).

Suppose that the initial weights are such that the equilibrium induced by (15.18)–(15.22) is (s_1, s_2) . Any set of initial weights in which the row player assigns the highest initial weight to the column player's choice of s_2 and the column player assigns the highest initial weight to the row player's choice of s_2 will do. One can then check that the play induced by (15.18)–(15.22) in the following periods is similar to the Cournot dynamics, i.e., each player plays a best response to the previous period's play by the opponent. Thus, play in the subsequent periods is (s_1, s_3) , (s_2, s_3) , (s_2, s_1) , (s_3, s_1) , (s_3, s_2) , (s_1, s_2) , $\dots$. Along the path that we have sketched out, play begins at (s_1, s_2) and returns to (s_1, s_2) . Thus, the entire process keeps cycling repeatedly and the empirical frequencies of play never converge.

Consider another example due to Fudenberg and Kreps (1993). Suppose that we have the symmetric two-player game shown in Figure 15.4. There are two pure strategy Nash equilibria in this game, (s_1, s_2) and (s_2, s_1) . There is also a mixed strategy Nash equilibrium in which each player randomizes between his two pure strategies. Now consider fictitious play. Suppose that

	s_1	s_2	s_3
s_1	0, 0	1, 0	0, 1
s_2	0, 1	0, 0	1, 0
s_3	1, 0	0, 1	0, 0

Figure 15.3 Non-convergence of probabilities under fictitious play.

Source: Fudenberg, Drew, and David K. Levine, *The Theory of Learning in Games*, Figure 2.2, page 34, © 1998 Massachusetts Institute of Technology, by permission of The MIT Press.

	s_1	s_2
s_1	0, 0	1, 1
s_2	1, 1	0, 0

Figure 15.4 A symmetric two-player coordination game.

Source: Fudenberg, Drew, and David K. Levine, *The Theory of Learning in Games*, Figure 2.3, page 36, © 1998 Massachusetts Institute of Technology, by permission of The MIT Press.

the initial weights for each player on the strategies of the opponent, (s_1, s_2) , are identical and given by $w_0^1 = w_0^2 = (1, \sqrt{2})$. Thus, the probabilities of play of the opponent are given by $x_0^1 = x_0^2 = (0.41, 0.59)$. Each player's best response is to choose s_1 . The weights in period 1 are $w_1^1 = w_1^2 = (2, \sqrt{2})$, and the probability vectors are $x_1^1 = x_1^2 = (0.58, 0.42)$. So the best response of each player is s_2 . The weights in period 2 are $w_2^1 = w_2^2 = (2, 2.41)$, the probability vectors are $x_2^1 = x_2^2 = (0.45, 0.55)$, hence, the best response of each player shifts back to s_1 . It can be shown, using an induction argument, that the sequence of play oscillates as follows: (s_1, s_1) , (s_2, s_2) , (s_1, s_1) , (s_2, s_2) , $\dots$

Since the game in Figure 15.4 is a 2×2 game, Proposition 15.22 applies. Hence, there is a unique limiting distribution of the probabilities with which the opponents are expected to play each of their pure strategies. In this case, the unique distribution (from Proposition 15.23) coincides with the mixed strategy Nash equilibrium of the game, i.e., $x^{1*} = x^{2*} = (0.5, 0.5)$. While the marginal distribution of the player's strategies converges to $(0.5, 0.5)$, the joint distribution of play oscillates repeatedly between the two worse outcomes.

A striking aspect of this equilibrium is that each player gets a payoff of 0 in each period, when the expected payoff under the mixed strategy Nash equilibrium is 0.5. Since play incessantly cycles, if one of the players departs from fictitious play, while the other does not, then he can improve his payoff to 1. Hence, behavior in fictitious play is too rigid and players do not seem to learn from systematic patterns of play.

Another variant of fictitious play called *weighted fictitious play* was proposed by Cheung and Friedman (1997). Suppose that there are two players. Pick any player i and refer to the other player $j \neq i$ as his opponent. The strategy played at time t by player i is denoted by $s_i(t) \in S_i$, where $S_i = \{s_{i1}, s_{i2}, \dots, s_{in_i}\}$.⁷ The opponent plays the strategy $s_j(t) \in S_j$ at time t , where $S_j = \{s_{j1}, s_{j2}, \dots, s_{jn_j}\}$. Define an indicator function I over pairs of the opponent's strategies as follows

$$I(s_{jk}, s_j(t)) = \begin{cases} 1 & \text{if } s_{jk} = s_j(t) \\ 0 & \text{if } s_{jk} \neq s_j(t) \end{cases}, s_{jk}, s_j(t) \in S_j.$$

Under *weighted fictitious play*, the beliefs of player i about the strategy of the opponent (the analog of (15.20)) are given by

$$x_{t+1}^i(s_{jk}) = \frac{I(s_{jk}, s_j(t)) + \sum_{\tau=1}^{t-1} \phi_i^\tau I(s_{jk}, s_j(t-\tau))}{1 + \sum_{\tau=1}^{t-1} \phi_i^\tau} \in [0, 1]; k = 1, 2, \dots, n_j, \quad (15.29)$$

⁷ Recall that under fictitious play, all pure strategies have some probability of being played. However, players actually play only one pure strategy.

where ϕ , the weighting parameter, allows this framework to nest a range of models. Once players form beliefs as in (15.29), they play a best response; so player i chooses $s_i \in S_i$ to maximize expected payoff $\pi_i(s_i, \sigma_{-i})$, where σ_{-i} is the vector of mixed strategies of the opponent, computed in accordance with (15.29). The *Cournot adjustment dynamics* correspond to the case $\phi = 0$, so $x_{t+1}^i(s_{jk}) = I(s_{jk}, s_j(t))$, i.e., players believe that the last period choice of the opponent will be played again and they play a best response to it. When $\phi = 1$, weighted fictitious play corresponds to ordinary fictitious play.

15.3.1 Empirical evidence on some belief-based models

Cheung and Friedman (1997) test for the predictions of weighted fictitious play in four games: a hawk–dove game, a coordination game (which is a stag-hunt game⁸), a buyer–seller game, and the battle of the sexes game. They estimate the parameter ϕ separately for the individuals in their experiment. The median values of ϕ lie between 0.25 to 0.50, which is closer to the Cournot extreme rather than the other extreme of ordinary fictitious play. For instance, in the battle of the sexes game, 51 individuals behaved in a Cournot manner, 18 in the manner of ordinary fictitious play (i.e., $\phi = 1$), 6 were adaptive ($0 < \phi < 1$), and 4 were fickle ($\phi < 0$). The type distribution is similar across the four games.

A more direct approach to testing for fictitious play is to test if the measured beliefs in actual games correspond to fictitious play. Nyarko and Schotter (2002) elicited beliefs in repeated play of the 2×2 game in Figure 15.5, which has a unique mixed strategy equilibrium.

Subjects in experiments played 60 rounds of the game in four sessions. Actual frequencies of play were intermediate between the unique mixed strategy equilibrium and uniform randomization. Since beliefs were directly measured, they could be fitted to a model of fictitious play. The results are as follows. (i) The parameter ϕ in (15.29) is close to 1, which corresponds to ordinary fictitious play. (ii) There is divergence between stated and fitted beliefs. (iii) Fitted beliefs converge rapidly, while stated beliefs vary substantially from one period to another and do not converge; see Figure 15.6. (iv) The actions of players seem to be consistent with the play of best response, given stated beliefs, but not, given the fitted beliefs; this holds in 75% of the cases. One implication is that using fictitious play as a model of stated beliefs is inappropriate; perhaps players are using information and a level of sophistication that exceeds what is required for fictitious play.

Cheung and Friedman (1998) compare the relative performance of replicator dynamics and weighted fictitious play using two games: a hawk–dove game, and a buyer–seller game. The

	Green	Red
Green	6, 2	3, 5
Red	3, 5	5, 3

Figure 15.5 A 2×2 stage game.

Source: Nyarko and Schotter (2002) with permission of The Econometric Society.

⁸ In a stag-hunt game, there are two pure strategy Nash equilibria, one is payoff-dominant, while the other is risk-dominant. Examples are considered below.




Figure 15.6 A comparison of stated and fictitious play beliefs.

Source: Nyarko and Schotter (2002) with permission of The Econometric Society.

Table 15.2 Summary of the treatments in Huck et al. (1999).

Treatment	Cost/demand information given?	Information on q_{ti} , p_t revealed at the end of each period?
<i>BEST</i>	Yes	No
<i>FULL</i>	Yes	Yes
<i>NOIN</i>	No	No
<i>IMIT</i>	No	Yes
<i>IMIT+</i>	Partial	Yes

learning model performs slightly better than the replicator model for the one population hawk–dove game, but substantially better in the two-population buyer–seller model.

Huck et al. (1999) test a range of belief-based theories. They consider a model of four symmetric firms with linear demand and unit marginal costs. At time t , the demand curve is given by

$$p_t = \max \{100 - Q_t, 0\},$$

where p_t is the market price, q_{ti} is the quantity chosen by the i th firm and $Q_t = \sum_{i=1}^4 q_{ti}$ is the total market quantity at time t . The four firms played this game for 40 time periods. Inertia is built into the model by allowing each firm to revise its strategy, every period, with probability equal to $2/3$; the probabilities of revision are independent across firms. The cost function of firm i is $C(q_{ti}) = q_{ti}$ and the profit function is $\pi_{ti} = q_{ti}(p_t - 1)$. The five treatments are summarized in Table 15.2.

The “partial information” in the treatment *IMIT+* provided to the subjects is the following statement: “as more is sold on the market, the lower is the price per unit that one obtains.” A simple

calculation shows that the symmetric Cournot–Nash outcome is given by $q_{ti}^N = 19.8$, $Q_t^N = 79.2$, $p_t^N = 20.6$. The Walrasian outcome (price equals marginal cost for each firm) is given by $q_{ti}^W = 24.75$, $Q_t^W = 99$, $p_t^W = 1$. The symmetric collusive outcome, which maximizes joint profits of all four firms, is given by $q_{ti}^C = 12.375$, $Q_t^C = 49.5$, $p_t^C = 50.5$.

The following main theories are tested.

- A. Under *myopic best reply dynamics*, each individual plays a best response in quantity, given the other players' choice of quantity in the previous period. This requires information on the cost and demand conditions. Both these pieces of information are given in the BEST and FULL treatment cases (see Table 15.2). The myopic best response dynamics are predicted to converge in finite time to the Cournot–Nash equilibrium.
- B. Under *imitation dynamics*, when a firm gets an opportunity to revise its strategy, then with probability $1 - \varepsilon$ (for some small $\varepsilon > 0$) it imitates the best choice by any of the firms' last period. This learning rule is also sometimes called *imitate the best*.⁹ But with probability ε , it chooses one of its feasible strategies (i.e., quantity) randomly. It was shown by Vega-Redondo (1997) that such imitation dynamics converge to the Walrasian outcome as $\varepsilon \rightarrow 0$.¹⁰ The intuition is simple. Since all firms face the same price, and have identical costs, thus the firm which produces the higher quantity makes the highest profits.

Suppose that in period $t - 1$, price exceeds marginal cost (i.e., $p_t > 1$). Consider a firm that in period t , gets an opportunity to choose its quantity, q_{ti} . Then, it will imitate the firm that produced the highest quantity in period $t - 1$. So its quantity chosen will increase over the previous period. This process continues until in some future period, price is equated to marginal cost. The converse process takes place, if at time $t - 1$, price falls below marginal cost, i.e., $p_t < 1$. In this case, the highest profits are made by the firm that produced the least. Hence, imitation dynamics dictate a reduction in quantity. This process continues until parity between price and marginal cost is restored. The information requirement for imitation dynamics is the quantity levels and profit levels of all firms in the last period. In particular, cost and demand conditions need not be known. The treatments FULL and IMIT contain this information (see Table 15.2).

- C. *Other models: Learning direction theory* associated most closely with Reinhard Selten, predicts that, over time, subjects should move in the direction of the best reply; see Section 15.6. Since players are always moving in the direction of their best reply, this dynamics naturally converges to the Cournot–Nash equilibrium.

In *trial and error learning*, subjects do not repeat their mistakes. So, for instance, if by reducing quantity, a firm finds that its profits have reduced, then it does not reduce quantity again. Huck et al. (2004) show that trial and error learning leads to the collusive outcome, which maximizes joint profits.

The results are as follows.

1. In all five treatments, on average, the total output produced exceeded the Cournot level. In treatment IMIT+, the actual outcome is very close to the competitive outcome, $Q_t^W = 99$. The collusive outcome was not observed in any treatment. The collusive price of $p_t^C = 50.5$ was exceeded in only 12 out of 1200 observations.

⁹ Another possible learning rule is *imitate the average* of the other firms rather than *imitate the best*. This can be motivated on grounds of caution.

¹⁰ Apesteguia et al. (2010) show that this result is not theoretically robust to asymmetric costs. However, in their experiments with asymmetric costs, they find that the outcomes still converge to a Walrasian equilibrium. They propose an alternative learning model to explain their results.

2. Greater *information about the market*, i.e., the cost/demand conditions, results in a *less* competitive outcome. To test for the role of information on market conditions, the treatments *BEST* and *NOIN* are compared as are the treatments *IMIT* and *IMIT+*. In each case, the mean quantity was lower in the treatment with the higher information about the market (i.e., information about cost/demand conditions). By contrast, more *information about the behavior of others*, i.e., quantities, profits of other firms, results in a *more* competitive outcome. This is seen by a comparison of the conditions *BEST* versus *FULL* and *NOIN* vs *IMIT*; in each of these comparisons, more information about others results in a higher mean quantity produced.
3. Firms produce higher output levels in the treatment *FULL* relative to the treatment *BEST*; in the latter treatment, the quantity converges fairly quickly to the Cournot–Nash solution. Insofar as only treatment *FULL* reveals the quantity/profit choices of others, imitation of others appears to play some role.
4. When subjects have the information necessary to play best replies, then most play a best reply. Similarly, when information on the quantity/profits in the previous period is available, then at least a few firms become pure imitators.
5. Learning direction theory does better in the treatment *BEST* relative to the treatment *FULL*. Furthermore, it predicts correctly in more than 70% of the cases and performs better when an upward, rather than a downward adjustment is required. However, trial and error learning does better than learning direction theory and “imitate the average” learning.

The role of imitation learning is also tested by Bosch-Domènech and Vriend (2003) in their duopoly and tripoly experiments. They consider a generic demand curve $p_t(Q_t) = a + bQ_t$, where $Q_t = \sum_i q_{ti}$ and the exact values of a, b vary over the treatments. Players compete over 22 rounds of play. The cost curve is given by $C(q_{ti}) = kq_{ti} + F$, where F is the fixed cost and $k > 0$. The profit of the i th firm is $\pi_{ti} = q_{ti}p_t(Q_t) - C(q_{ti})$.

Three different treatments are considered: Easy, hard, and hardest. The underlying market conditions and the information provided is identical in each of three treatments. However, the treatments differ in two dimensions. First, the time allowed to make a decision decreases as we move from the easy to the hardest treatment. In the easy treatment, there is no time pressure, while in the hardest treatment, subjects must make their output choices in one minute. Second, the framing of information is most transparent in the easy condition, most opaque in the hardest treatment, and intermediate in the hard treatment. Hence, one would expect that subjects are more likely to play their best response in the easy treatment. On the other hand, in the other two treatments, cognitive bounds are more likely to lead players away from best response to something like imitation learning. The latter should lead to the Walrasian outcome as $\varepsilon \rightarrow 0$ (see discussion of the previous model).

As one moves from the easy to the hardest treatments, we get the following results. (i) There is greater variability in output choices. (ii) Average output increases, but falls short of the Walrasian output levels. (iii) There is no trend movement in output towards the Walrasian level. (iv) There is no increase in the level of imitation. Thus, there is only weak support for the theoretical predictions.

15.3.2 Reinforcement learning versus fictitious play: empirical evidence

Mookerjee and Sopher (1997) test various models of learning, including reinforcement learning and fictitious play. They conduct their experiments using two zero-sum games, a 4×4 game,

and a 6×6 game. Each strategy for player i , $s_{ij} \in S_i$, is given a *score*. The score given to strategy s_{ij} is

$$\phi_{ij} = \sum_{k=1}^{n_i} \alpha_{jk} A_{t-1}(s_{ik}),$$

where $\alpha_{jk} \in \mathbb{R}$, and $A_{t-1}(s_{ik})$ is the previous period's *attraction/propensity*, or *weight* to play strategy s_{ik} . Using scores, and a logit rule, players compute the probability of playing any strategy s_{ij} as

$$p_{ij} = \frac{e^{\phi_{ij}}}{\sum_{k=1}^{n_i} e^{\phi_{ik}}}. \quad (15.30)$$

The logit rule used in (15.30) is very general; it allows the attractions of one strategy to affect the probabilities of playing other strategies. In this sense, Mookerjee and Sopher (1997) also generalize the insights obtained from the quantal response equilibrium of McKelvey and Palfrey (1995). Indeed, the author's preferred name for their learning mechanism is *quantal response learning*.

The results are as follows. Players learn from experience. The bulk of the evidence points to reinforcement learning. Fictitious play does not fit their data as well. The estimated parameters of the learning model fluctuate from game to game and vary between the early and later stages of play. The cross-attraction effects are generally insignificant. The Nash equilibrium for constant-sum games implies minimax play (i.e., minimize the maximum loss). However, substantial deviations from minimax play are found. The initial play influences the path of play in subsequent periods in important ways (history dependence). Dominated strategies are also sometimes played, but early rather than later during the 40 rounds of play.

Playing dominated strategies is consistent with reinforcement learning, but not deterministic fictitious play because in the latter, players play a best response given their beliefs about the other player. Dominated strategies can be accommodated within fictitious play only when players make mistakes, or there are some random shocks to the payoff function. This is not the case in the games studied here. However, the early play of dominated strategies is more consistent with reinforcement learning because players make mistakes early on and then learn not to play their dominated actions.

Battalio et al. (2001) consider three *stag-hunt* games. In a stag-hunt game, there are two pure strategy Nash equilibria, one is payoff-dominant, while the other is risk-dominant. The focus of these experiments is to see which of the two Nash equilibria is played. The three games are labeled as 2R, R, and 0.6R, and are shown in Figure 15.7. The payoffs shown in each cell are the actual number of cents received by each player. Subjects in experiments have complete information

Game 2R				Game R				Game 0.6R			
		X	Y			X	Y			X	Y
X	45, 45	0, 35		X	45, 45	0, 40		X	45, 45	0, 42	
Y	35, 0	40, 40		Y	40, 0	20, 20		Y	42, 0	12, 12	

Figure 15.7 Three stag-hunt games with varying levels of the optimization premium.

Source: Battalio et al. (2001) with permission of The Econometric Society.

about their own payoffs and the payoff of others. There is no preplay communication. In each of the three games, the payoff Nash equilibrium is (X, X) and the risk-dominant Nash equilibrium is (Y, Y) .

For all three games, suppose that the probability with which a player plays X is given by q . The essential difference between each game is the incentive/penalty attached to not playing one's best response. To see this, denote in game $j = 2R, R, 0.6R$, the payoff to a player from playing strategy $s = X, Y$ when the opponent is expected to play X with probability q by $\pi_j(s, q)$; this payoff is symmetric. The benefit from playing a best response, relative to playing a non-best response is termed as the *optimization premium*. Let $q^* = 4/5$. For any of the games, the best response is X if $q > q^*$, and Y if $q < q^*$. If $q = q^*$, then the player is indifferent between X and Y . Let

$$\Delta\pi_j = \pi_j(X, q) - \pi_j(Y, q); j = 2R, R, 0.6R.$$

For game $j = 2R, R, 0.6R$, and for beliefs about the opponent's play of X given by q , we denote the optimization premium as $r_j(q)$. A simple calculation gives

$$\begin{aligned} r_{2R}(q) &= \Delta\pi_{2R} = 45q - [35q + 40(1 - q)] = 50q - 40 = 50(q - q^*) = \delta_{2R}(q - q^*), \\ r_R(q) &= \Delta\pi_R = 45q - [40q + 20(1 - q)] = 25q - 20 = 25(q - q^*) = \delta_R(q - q^*), \\ r_{0.6R}(q) &= \Delta\pi_{0.6R} = 45q - [42q + 12(1 - q)] = 15q - 12 = 15(q - q^*) = \delta_{0.6R}(q - q^*), \end{aligned}$$

where δ_j is called the *optimization premium parameter*. Suppose that $p_j(q; \lambda)$ is the probability that “ X is chosen” for game j , where λ is a parameter of the probability density function, and q is the probability with which the opponent is expected to play X . The logistic response function that is used is given by

$$p_j(q; \lambda) = \frac{e^{\lambda\delta_j(q-q^*)}}{1 + e^{\lambda\delta_j(q-q^*)}}. \quad (15.31)$$

If $\lambda = 0$, then $p_j(q; \lambda) = 0.5$, so each player plays both strategies with equal probability. On the other hand, as λ increases, one approaches best response behavior.

Definition 15.19 (McKelvey and Palfrey, 1995): For a given value of λ , a *logit* (or a *quantal response*) *equilibrium* is a fixed point of the two players' logit response functions, i.e., $q^* = p_j(q^*; \lambda)$.

The Nash equilibrium is invariant to the optimization premium. However, holding λ fixed, $\frac{\partial p_j}{\partial \delta_j} > 0$, hence we get the first hypothesis under a quantal response equilibrium:

H1: Actual behavior of subjects is more responsive to beliefs if the optimization premium parameter is higher.

The second hypothesis follows by using the single population continuous-time logistic-response dynamic. This dynamic is given by $\dot{q} = p_j(q; \lambda) - q$, where q is now interpreted as the frequency of X choices in the population; see for instance Fudenberg and Levine (1998).

H2: The convergence to an equilibrium is faster if the optimization premium is higher.

To get a feel for the final hypothesis, consider the following intuition. Suppose that learning is noisy. For instance, with some probability players choose their best response, but with a complementary probability, they make a tremble and randomize equally between their strategies.

In this case, for a smaller optimization premium, the likelihood of entering the basin of attraction of the payoff dominant equilibrium increases. This can be shown by using the aspiration adaptation model of Binmore and Samuelson (1997), but the explanation relies on tools that we develop in Chapter 16.

Whenever the payoffs are below the aspiration level, players are more likely to experiment by trying other strategies, so learning is relatively noisier. In the games shown in Figure 15.7, the payoffs in the risk dominant equilibrium (Y, Y) fall as the optimization premium falls, but the payoffs in the payoff dominant equilibrium are unchanged. Hence, as the optimization premium falls, the greater experimentation of players with alternative strategies in the risk dominant equilibrium makes it relatively less costly to escape its basin of attraction. The cost of escaping from the basin of attraction of the payoff dominant equilibrium is unchanged. Hence, as the optimization premium falls, the payoff dominant equilibrium is more likely to be *stochastically stable*; See Appendix C and Chapter 16. This leads to a prediction that is not made by best-response, logistic-response, or replicator dynamics.

H3: There is a greater likelihood of the behavior converging to the payoff dominant equilibrium if the optimization premium is lower.

All three hypotheses are confirmed by the experimental data. Individuals are more sensitive to the past play of the opponent, the greater the optimization premium. Play converges faster to an equilibrium, if the optimization premium is larger. Thus, the fastest convergence is in game 2R, followed by games R and 0.6R. The payoff dominant equilibrium is more likely to emerge, the smaller the optimization premium and this effect is more pronounced in later periods of play. The main message of this paper is that it is not sufficient to focus on the best response dynamics—other factors such as the optimization premium matter too.

Individual behavior is studied, using a modified version of (15.31), which is written as

$$p_{it}(q_{it}; \lambda) = \frac{e^{\alpha_i + \lambda \delta_j(q_{it} - q^*)}}{1 + e^{\alpha_i + \lambda \delta_j(q_{it} - q^*)}}; i = 1, 2, j = 2R, R, 0.6R \quad (15.32)$$

where $i = 1, 2$ is an index for players, p_{it} is the probability that player i plays X at time t , q_{it} is the probability with which the opponent is believed to be playing X at time t , and α_i is a player-specific drift parameter. They also consider the case where $\delta_j(q_{it} - q^*)$ is the average of all play up to that point in time. This corresponds to the *average-payoff reinforcement model*. The two main findings are that the belief-based model fits the data better than the average-payoff reinforcement model, and a combination of the belief-based and the reinforcement model does even better.

Feltovich (1999, 2000) argues that the existing experiments use games that are too simple. He introduces complexity through the channel of asymmetric information. Nature chooses one of the two games in Figure 15.8, Left or Right. The row player (player 1) is the informed player and knows nature's choice of the game. The column player (player 2) is the uninformed player and his prior belief that the game is Left is given by p . Denote this game by $G(p)$. In each game, player 1 has a weakly dominated strategy (B in Left and A in Right). Once nature chooses one of these games, it is played over two stages. The total payoffs for each player equal the sum of the payoffs over the two stages. The payoffs in the first stage are not revealed to player 2, but the actions of both players are publicly revealed at the end of the first stage.

If player 1 uses the strategy “play A if chosen game is Left and play B if the chosen game is Right,” then these actions fully reveal the type of game to player 2. These are the stage dominant actions of player 1 that are completely revealing of his type. In this case, player 2's optimal second

Left matrix			Right matrix		
	A	B		A	B
A	1, 0	0, 1	A	0, 1	0, 1
B	0, 1	0, 1	B	0, 1	1, 0

Figure 15.8 Introducing complexity through the channel of asymmetric information.

Source: Feltovich (2000) with permission of The Econometric Society.

period strategy is to play B if player 1 plays A in the first period and play A if player 1 plays B in the first period. In each of these cases, player 1 gets a payoff of zero in the second period. Hence, player 1 might do better by playing a mixed strategy in the first period that does not reveal the game to player 2. We focus only on the first stage play of player 1, contingent on the choice by nature of either the Left or the Right payoff matrix.

For the game $G(0.5)$, the following is an equilibrium (Feltovich, 2000). Player 1 randomizes equally between his two first period strategies irrespective of the payoff matrix chosen by nature. Suppose that Player 2 observes a choice of A by player 1, then he plays his best response in a completely revealing equilibrium with probability $\beta \in [0.5, 1]$. If player 2 observes a choice of B, then he plays his best response in a completely revealing equilibrium with probability $1.5 - \beta$. This describes the Nash equilibrium of the game. This equilibrium is not unique because it depends on the parameter β . In a second game, $G(0.34)$, the equilibrium is unique. Here, in the Left matrix, player 1 plays his stage dominant action A with probability 0.97 and in the Right matrix, he randomizes equally between the two actions. The authors then pit this equilibrium model against the reinforcement and the belief-based learning models.

The results are as follows. In both games, play starts and ends (over several rounds) substantially away from the Nash equilibrium outcome. Averaged across 40 rounds of learning, the results are shown in Table 15.3. There is little support for the Nash predictions in the behavior of either of the two players. These results, and many other reviewed in this chapter and elsewhere in the book, should caution us against assuming that given sufficiently many periods of learning, the outcome is a Nash equilibrium.

Examining the time pattern of play in Figure 15.9, both learning models, reinforcement and belief-based learning, are able to explain the qualitative features of actual play. Whilst the

Table 15.3 Results averaged across 40 rounds of play. * denotes significantly different from the predicted value at 5% and ** denotes significantly different at 0.1% level.

	Player 1 probabilities of playing the stage dominant action		Player 2 probabilities of playing the best response to completely revealing action by player 1	
	Predicted	Actual	Predicted	Actual
G(0.5)	0.500	0.808**	[0.5-1]	0.869
G(0.34)	0.660	0.818**	0.670	0.749*

Source: Feltovich (2000).




Figure 15.9 Aggregate behavior across all experiments over 40 rounds of play. The horizontal axis is the proportion of times player 1 played his stage dominant action and the vertical axis is the proportion of times player 2 played a best response to the play of a stage dominant action by player 1. The point labeled NE shows the Nash equilibrium.

Source: Feltovich (2000) with permission of The Econometric Society.

belief-based model is better at predicting aggregate behavior (shown in Figure 15.9), the reinforcement model does better at explaining individual behavior.¹¹

Overall, based on the evidence reviewed above, it is difficult to declare a clear winner in the comparison between these two alternative models of learning. Confounding reasons include: (i) The papers differ very much in the sorts of games that are analyzed. (ii) Differences in the way that the experiments are implemented. (iii) Differences in the versions of the two learning models used. Reinforcement learning appears to fit better than belief-based models in games with mixed strategy equilibrium that have relatively low dimensional strategy spaces (Erev and Roth, 1998; Mookerjee and Sopher, 1994). However, belief-based models make relatively better predictions in coordination games (Battalio et al., 2001; Ho and Wiegelt, 1996).

15.4 The experience weighted attraction model of learning

Models of reinforcement learning and fictitious play are sometimes seen as two entirely separate models of learning. Under reinforcement learning, players are not required to know the structure of the game. They do not use the historical information on the strategies of the opponent, nor are they required to be aware of any such information. Their objective is to play those strategies that, in their experience, have done well in the past. More successful strategies are then used more often. Under fictitious play, players know the structure of the game. They are not interested in their own historically successful strategies. Rather, they are interested in the history of play of their opponents, and use it to predict their current choices; and they play a best response to these

¹¹ The reinforcement learning model is also found to fit better as compared to the belief-based model in Blume et al. (2001) (for a sender-receiver game) and by Nagel and Tang (1998) (for the centipede game).

predictions. Camerer and Ho (1999) show that both these models are actually special cases of an overarching model which they call *experience weighted attraction* model of learning (EWA).¹²

Suppose that each player $i = 1, 2, \dots, n$ has a set $S_i = \{s_{i1}, s_{i2}, \dots, s_{in_i}\}$ of *pure strategies*; denote a typical element of S_i by s_i . In any period, the profile of strategies that is played is $s = \{s_1, s_2, \dots, s_n\} \in S = S_1 \times S_2 \times \dots \times S_n$. Denote the pure strategy profile of all opponents of player i by s_{-i} . The players have *payoff functions* $\pi_i : S \rightarrow \mathbb{R}$. Time is discrete, $t = 0, 1, 2, \dots$

The EWA model has two main components.

1. At any time t , the experience weight N_t encapsulates the experience of the players. It is also known as the *observation-equivalents of past experience*.
2. The second component is similar to *propensities* under reinforcement learning. It is denoted by $A_t^i(s_{ij})$, $\forall i, j$ and is referred to as the *attraction* that player i has at time t for his j th pure strategy, s_{ij} .

The initial values of $N_t, A_t^i(s_{ij})$, denoted respectively by $N_0, A_0^i(s_{ij})$, reflect pre-game experience of the players, say, from having played similar games in the past. The updating rules for $N_t, A_t^i(s_{ij})$ are as follows.

1. The updating rule for experience, N_t , is given by

$$N_t = (1 - \kappa)\phi N_{t-1} + 1, t \geq 1, \quad (15.33)$$

where $(1 - \kappa)\phi \in [0, 1]$. The parameter ϕ allows for the depreciation of past experience, on account of (i) forgetting, or limited memory, or (ii) a rapidly changing environment. In most empirical work, N_0 is set equal to 1.

2. Although *attractions* under EWA are similar to *propensities* under reinforcement learning, there is an important difference. Under reinforcement learning, only the strategies that are actually played get reinforced (*law of effect*). Players do not, however, take account of the forgone payoff that could have resulted from playing other strategies.

EWA alters this assumption by proposing, instead, a *law of simulated effect*. This does not dispute the law of effect, but, in addition, also updates the attractions of unchosen strategies by taking account of forgone payoffs from having not played these strategies.

Suppose that player i plays $s_i(t) \in S_i$ at time t , while the opponents play $s_{-i}(t) \in S_{-i}$. Then EWA specifies attractions as follows.

$$A_t^i(s_{ij}) = \frac{1}{N_t} [\phi N_{t-1} A_{t-1}^i(s_{ij}) + [\delta + (1 - \delta) I(s_{ij}, s_i(t))] \pi_i(s_{ij}, s_{-i}(t))], \forall s_{ij} \in S_i. \quad (15.34)$$

The constant $\phi \in [0, 1]$ enables previous levels of experience weighted attractions, $N_{t-1} A_{t-1}^i(s_{ij})$, to decay. The second term in (15.34) allows for the laws of actual and simulated effects to work out. $\delta \in [0, 1]$ is a parameter that gives the weight placed on forgone payoffs; one may interpret it as a form of *regret* over forgone payoffs. $I(s_{ij}, s_i(t))$ is an indicator function that is defined as

$$I(s_{ij}, s_i(t)) = \begin{cases} 1 & \text{if } s_{ij} = s_i(t) \\ 0 & \text{if } s_{ij} \neq s_i(t) \end{cases},$$

¹² Hopkins (2002) showed that the expected dynamics under stochastic fictitious play and reinforcement learning with experimentation can be shown to be a perturbed form of replicator dynamics. It is possible that this result may also hold for the EWA model.

where $s_{ij} \in S_i$ is any pure strategy of player i , and $s_i(t) \in S_i$ is the strategy that is actually played at time t . $\pi_i(s_{ij}, s_{-i}(t))$ is the payoff at time t of player i from playing s_{ij} against the “actual” strategies of the opponents, $s_{-i}(t) \in S_{-i}$. When $s_{ij} = s_i(t)$, the law of effect operates and when $s_{ij} \neq s_i(t)$, the law of simulated effect operates. Thus,

$$A_t^i(s_{ij}) = \begin{cases} \frac{1}{N_t} [\phi N_{t-1} A_{t-1}^i(s_{ij}) + \pi_i(s_{ij}, s_{-i}(t))] & \text{if } s_{ij} = s_i(t) \\ \frac{1}{N_t} [\phi N_{t-1} A_{t-1}^i(s_{ij}) + \delta \pi_i(s_{ij}, s_{-i}(t))] & \text{if } s_{ij} \neq s_i(t) \end{cases} \quad (15.35)$$

The parameter κ in (15.33) determines the growth rate of attractions; it is also referred to as the *lock-in parameter*. When $\kappa = 0$, attractions are weighted averages of past payoffs and lagged attractions. When $\kappa = 1$, attractions cumulate and can exceed the stage game payoffs.

The final step in the model is similar to the conversion of propensities into probabilities in the reinforcement learning model. In EWA, we convert attractions into probabilities with which players play their strategies. These probabilities of play can take several forms, logit, exponential, power, probit, or normal. Camerer and Ho (1999) prefer the logit form. Let $p_t^i(s_{ik})$ be the probability with which player i plays pure strategy s_{ik} at time t , then, the logit form dictates that

$$p_{t+1}^i(s_{ik}) = e^{\lambda A_t^i(s_{ik})} / \sum_{j=1}^{n_i} e^{\lambda A_t^i(s_{ij})}. \quad (15.36)$$

The parameter λ measures the sensitivity of probabilities to attractions, as in the reinforcement model. If payoffs are negative, i.e., $\pi_i(s_{ij}, s_{-i}(t)) < 0$, then the corresponding attractions can also be negative. However, this is not a problem if one uses the logit form, because it is invariant to adding a constant to all attractions (this can be directly checked from (15.36)). One could also use the *power form*, which is given by

$$p_{t+1}^i(s_{ik}) = (A_t^i(s_{ik}))^\lambda / \sum_{j=1}^{n_i} (A_t^i(s_{ij}))^\lambda. \quad (15.37)$$

It is easy to check that the power form is invariant to multiplying all attractions by a constant. It is not established, yet, which of these forms, logit, or power, fits better. Erev and Roth (1998) and Chen and Tang (1998) are unable to conclude that any of these forms is superior to the other, while Camerer and Ho (1997) find that the logit form fits relatively better.

We now show that reinforcement learning and belief-based learning are special cases of this model.

15.4.1 Special case: reinforcement learning

Section 15.2 considered a variety of reinforcement learning models. Consider the one-parameter model in (15.2), but allow for past levels of reinforcements to decay at the rate $\phi \in [0, 1]$, and specialize the reinforcement function to $R(\pi_i) = \pi_i$. Hence, (15.2) can be written as:

$$q_{t+1}(s_{ij}) = \begin{cases} \phi q_t(s_{ij}) + \pi_i(s_{ij}, s_{-i}(t)) & \text{if } s_{ij} \text{ is played at time } t \\ \phi q_t(s_{ij}) & \text{otherwise} \end{cases} \quad (15.38)$$

In the EWA model, suppose that $\delta = 0$, $N_0 = 1$, and $\kappa = 1$, so $(1 - \kappa)\phi = 0$. Using (15.33), this implies that $N_t = 1$ for all t . (15.35) implies that

$$A_{t+1}(s_{ij}) = \begin{cases} \phi A_t(s_{ij}) + \pi_i(s_{ij}, s_{-i}(t)) & \text{if } s_{ij} \text{ is played at time } t \\ \phi A_t(s_{ij}) & \text{otherwise} \end{cases}. \quad (15.39)$$

Interpreting attractions as propensities, (15.38), (15.39) are identical.

When $\delta = 0, \kappa = 0$, and $N_0 = \frac{1}{1-\phi}$, we get that

$$A_{t+1}(s_{ij}) = \begin{cases} \phi A_t(s_{ij}) + (1 - \phi) \pi_i(s_{ij}, s_{-i}(t)) & \text{if } s_{ij} \text{ is played at time } t \\ \phi A_t(s_{ij}) & \text{otherwise} \end{cases} \quad (15.40)$$

which takes the form of *average choice reinforcement*, as in Mookerjee and Sopher (1997) and Sarin and Vahid (2001).

15.4.2 Special case: belief-based learning

Section 15.3 considered a simple model of fictitious play. Here we consider a more general version of the model of fictitious play. Denote the initial beliefs of player i about the frequencies of choices of other players by $N_0(s_{-i}) > 0, s_{-i} \in S_{-i}$; these are the initial weights w_0^i in Section 15.3. For all $s_{-i} \in S_{-i}$, define $N_t^i = \sum_{s_{-i} \in S_{-i}} N_t(s_{-i})$, $t = 0, 1, 2, \dots$. Then, the prior belief of player i that the other players will play the strategy profile $s_{-i}^k \in S_{-i}$ is given by

$$B_0^i(s_{-i}^k) = \frac{N_0(s_{-i}^k)}{N_0^i}. \quad (15.41)$$

Notice that we model beliefs over strategy profiles of opponents here, rather than beliefs over each individual pure strategy. Beliefs about the strategy profiles of others, for the next period, are then constructed recursively, as follows

$$B_t^i(s_{-i}^k) = \frac{\rho N_{t-1}(s_{-i}^k) + I(s_{-i}^k, s_{-i}(t))}{\sum_{j=1}^M \rho N_{t-1}(s_{-i}^j) + I(s_{-i}^j, s_{-i}(t))}, \quad (15.42)$$

where M is the cardinality of the set S_{-i} , $\rho = (1 - \kappa)\phi$ is a decay parameter. $I(s_{-i}^k, s_{-i}(t))$ is an indicator function such that

$$I(s_{-i}^k, s_{-i}(t)) = \begin{cases} 1 & \text{if } s_{-i}^k = s_{-i}(t) \\ 0 & \text{if } s_{-i}^k \neq s_{-i}(t) \end{cases}, \quad (15.43)$$

where $s_{-i}(t)$ is the actual strategy profile played by players others than i at time t . Using (15.41), we can rewrite (15.42) in terms of the beliefs in the previous period, as

$$B_t^i(s_{-i}^k) = \frac{\rho N_{t-1}^i B_{t-1}^i(s_{-i}^k) + I(s_{-i}^k, s_{-i}(t))}{\rho N_{t-1}^i + 1}. \quad (15.44)$$

By definition, the expected payoff of player i from playing some strategy s_{ij} is defined by

$$\Pi_t^i(s_{ij}) = \sum_{k=1}^M B_t^i(s_{-i}^k) \pi(s_{ij}, s_{-i}^k). \quad (15.45)$$




Figure 15.10 The EWA cube.

Source: Reprinted from *Handbook of Game Theory with Economic Applications*, Young, P. and Zamir, S. (eds.), "Behavioral game theory experiments and modeling," Camerer, C. and Ho, T.-H. Volume 4, pages 517–73. Copyright © 2015, with permission from Elsevier.

Substituting (15.44) into (15.45) we can express the period t expected payoffs in terms of the period $t - 1$ expected payoffs as follows:

$$\Pi_t^i(s_{ij}) = \frac{\rho N_{t-1}^i \Pi_{t-1}^i(s_{ij}) + \pi(s_{ij}, s_{-i}^k)}{\rho N_{t-1}^i + 1}. \quad (15.46)$$

Consider the attractions given in (15.34). Let $\delta = 1, \kappa = 0$, and let initial attractions equal the expected payoffs $(A_0^i(s_{ij}) = \Pi_0^i(s_{ij}) = \sum_{k=1}^M B_0^i(s_{-i}^k) \pi(s_{ij}, s_{-i}^k))$. Then the attractions in (15.34) are identical to (15.46), so the weighted belief model is a special case of EWA. Not all learning models can, however, be subsumed within the EWA model.¹³

The various special cases of the EWA model are shown in (κ, ϕ, δ) space in the cube in Figure 15.10. Cournot best-response dynamics is shown by the triple $(\kappa, \phi, \delta) = (0, 0, 1)$, standard fictitious play by $(\kappa, \phi, \delta) = (0, 1, 1)$, and weighted fictitious play by the line $(\kappa, \phi, \delta) = (0, \phi, 1)$. Various forms of reinforcement learning correspond to $\delta = 0$ and different values of κ .

¹³ These include, for instance, certain kinds of belief learning models (Crawford, 1997; Broseta et al., 2000) and rule learning models (Stahl, 2000, 2001).

15.4.3 Empirical evidence for the EWA

The EWA model has been tested in a range of games such as the centipede game, signaling games, and patent race games.¹⁴ From these studies, Camerer et al. (2001) conclude the following. When there are no mixed strategy equilibria in a game, then the value of δ lies between 0.5 and 1. In games with mixed strategy equilibria, δ is close to 0. ϕ lies close to a value of 0.9.¹⁵ However, Wilcox (2006) shows that if parameter heterogeneity across players is not considered, then one obtains biased estimates that favor reinforcement learning over belief-based learning.

One potential criticism of the EWA model is that it has a greater number of free parameters than most models, so it is expected to fit better. In response, Camerer et al. (2001) first fit the model to the data and then they use the estimated parameters of the model in order to check for the success of out-of-sample predictions. A model with a greater number of parameters need not predict better out-of-sample. The EWA does better (often modestly) than the reinforcement learning model in 11 out-of-sample comparisons. And in 9 out of 11 out-of-sample cases, it does better (often dramatically) than belief-based learning. Using statistical tests such as the Akaike and the Bayesian criteria, which penalize more complex models, EWA does better than belief-based learning in all in-sample comparisons and better than reinforcement learning in 16 out of 20 comparisons.

In Figure 15.10, each of the dotted lines with a hollow circle at the top represents a three-parameter estimate of ϕ, δ, κ in the cube in 20 different datasets, each from a different game. The results from various studies are summarized by Camerer et al. (2001) as follows.

1. The points generally lie away from the edges or corners that represent an extreme case. The only cluster of points is toward the region where $\phi \cong 1, \delta = \kappa = 0$ (this is the average reinforcement model). This data typically arises from games with a mixed strategy equilibrium.
2. The estimated triples are dispersed throughout the cube, which suggests that no single special case of EWA describes well the data from diverse games.
3. One particular vertex of the cube $\delta = \kappa = 1$ corresponds to no known learning theory. However, estimates from three order-statistic games and the p-beauty contest¹⁶ lie in this region. This justifies the extra flexibility incorporated in the EWA model.

15.4.4 Functional EWA (fEWA)

In response to the criticism of “too many parameters” in the EWA model, Ho et al. (2007) propose a one-parameter family of learning models that they call *functional EWA* or fEWA. In fEWA models, for any player i , the three parameters ϕ, δ, κ in EWA are replaced by deterministic functions of experience, $\phi_i(t), \delta_i(t)$. These functions are player and time dependent, hence, this model is also sometimes known as self-tuning EWA. The parameter κ is set to 0 because it does not appear to have a significant effect on the fit of the model (Ho et al., 2008).

This leaves only two more parameters to be determined relative to EWA. These are the initial attractions and the parameter λ in (15.36) that enables the conversion of attractions to probabilities. The initial attractions are estimated from first period data. Hence, λ is the single parameter

¹⁴ See, for instance, Camerer and Ho (1999), Camerer et al. (1999), Andersen and Camerer (2001), Morgan and Sefton (1999), Rapoport and Amaldoss (2000), and Stahl (2001).

¹⁵ See Camerer et al. (2001) for a summary of these values over many different empirical studies.

¹⁶ These games are discussed in Part 4.

in fEWA models that is estimated using maximum likelihood methods. For this reason, relative to the EWA model, this model is also sometimes known as EWA-lite.

In the EWA model, the parameter ϕ is termed as the forgetting parameter. By contrast, fEWA unpacks this parameter and replaces it by the function $\phi_i(t)$ that is designed to detect the speed with which the environment is changing. $\phi_i(t)$ is called the *change-detecting function*. The function ϕ_i is self-tuning in the sense that if player i detects at time t that the environment is changing rapidly, then $\phi_i(t)$ falls, so that the weight placed on past experience is reduced.

It is important to distinguish real change and change that simply results from noise. In fEWA this is accomplished through a *surprise index* that detects the difference between the historical frequency of strategies of other players relative to their current choices. Denote by $h_i(t)$, the *cumulative history vector*, such that an element of this vector is given by

$$h_i^k(t) = \frac{1}{t} \sum_{\tau=1}^t I(s_{-i}^k, s_{-i}(\tau)), \quad (15.47)$$

where $I(s_{-i}^k, s_{-i}(\tau))$ is defined in (15.43). Thus, (15.47) gives the fraction of times that player i observes strategy s_{-i}^k to be played in the last t time periods. In contrast, the immediate history vector, $g_i(t)$, records the current choices of opponents. All elements of $g_i(t)$ are either zero or one. If $s_{-i}^k = s_{-i}(t)$, then the k th element in the vector $g_i(t)$ takes a value 1, and 0 otherwise, so $g_i^k(t) = I(s_{-i}^k, s_{-i}(t))$.

The *surprise index* for player i at time t is defined as the sum of squared deviations of the current strategy from the historical observations, as follows.

$$S_i(t) = \sum_{\tau=1}^t (h_i^k(\tau) - g_i^k(\tau))^2 \in [0, 2). \quad (15.48)$$

Suppose that the strategy of the opponents is identical in the last t periods, say, $s_{-i}^k = s_{-i}(\tau)$, $\tau = 1, \dots, t$, then $S_i(t) = 0$. This is a stationary environment, so there is no surprise. At the other extreme, suppose that for the first $t-1$ periods, the opponents played the strategy s_{-i}^k , but in period t they switch to another strategy $s_{-i}^r \neq s_{-i}^k$, then (15.48) can be written as

$$S_i(t) = \left(\frac{t-1}{t} \right)^2 + (-1)^2. \quad (15.49)$$

For the case depicted in (15.49) we have $\lim_{t \rightarrow \infty} S_i(t) = 2$, which gives the upper limit of $S_i(t)$ in (15.48). This is the most turbulent environment faced by the individual. The *change-detecting function* is then given by

$$\phi_i(t) = 1 - \frac{1}{2} S_i(t) \in (0, 1]. \quad (15.50)$$

The bounds in (15.50) follow from the bounds on $S_i(t)$ in (15.48). For the case depicted in (15.49) we have $\phi_i(t) = \frac{2t-1}{2t^2}$. For values of $t = 2, 3, \dots$ the value of $\phi_i(t)$ can be seen to be declining. Thus, the value of ϕ_i is lower and past experience is discounted more, the longer the environment is stationary followed by an abrupt change in strategy.

In the EWA model, δ allowed the player to take account of forgone payoffs (see (15.34)). However, δ may also be linked to the attention that subjects pay to the game. More inattentive players

will experience more ex-post regret from discovering that they have forgone higher payoffs. In order to take account of these concerns, define the *attention function*

$$\delta_{ij}(t) = \begin{cases} 1 & \text{if } \pi_i(s_i^j, s_{-i}(t)) \geq \pi_i(t) \\ 0 & \text{otherwise} \end{cases}, \quad (15.51)$$

where $\pi_i(t)$ is the actual payoff of player i at time t and $\pi_i(s_i^j, s_{-i}(t))$ is the payoff of player i at time t if he chooses s_i^j , when the strategy profile of others is $s_{-i}(t)$. Thus, all strategies that give a weakly higher payoff relative to the realized payoff are reinforced, and all remaining strategies receive zero reinforcement. If players play a best response, then $\delta_{ij}(t) = 0$ for all strategies that were not chosen. At the other extreme, if subjects choose their worse strategy, then $\delta_{ij}(t) = 1$ for all their strategies. The attention function in (15.51) describes a position in between these two extremes that captures limited attention. This is consistent with learning direction theory in which players gradually move in the direction of their ex-post best responses; see Section 15.6.

In the early stages of a game, the choices of players are quite variable, and because they may have limited experience and limited attention, they may make lots of mistakes. Thus, initially $\phi_i(t)$ is low and $\delta_{ij}(t)$ is high, so past experience is relatively highly discounted and players pay relatively low attention. In many experimental games, the behavior stabilizes over time, so $\phi_i(t)$ rises and $\delta_{ij}(t)$ falls. Thus, as the game proceeds, the weight on experience goes up and the forgone strategies are reinforced less.

Once these functions feed into the determination of attractions, the attractions are converted to probabilities of play, using (15.36). Ho et al. (2007) examine the empirical properties of fEWA using seven games, most of which were discussed in detail in Part 4 of the book. These games are listed in the first column of Table 15.4 that is an updated table of results from Camerer and Ho (2014). The labels of these games and the source is as follows. (i) The Cont'l divide game is the continental divide game (Van Huyck et al., 1997). (2) The Median action game (Van Huyck et al., 1991). (3) p-bc is the p-beauty contest game (Ho et al., 1998). (4) Price matching game, or the traveller's dilemma (Capra et al., 1999). (5) Mixed games are games with a unique mixed strategy Nash equilibrium (Mookerjee and Sopher, 1997). (6) Patent race games (Rapoport and Amaldoss, 2000). (7) Pot games are a form of market entry games (Amaldoss et al., 2005).

The subject pool is divided into two parts. First, data from 70% of the subjects is used to estimate the value of λ (see (15.36)), using maximum likelihood methods in an fEWA model that is described above. Second, the estimated model is used to form predictions for the behavior of the remaining 30% subjects, and the out-of-sample predictions are then compared to actual data. Two measures are used to evaluate the success of out-of-sample predictions. These are: the hit rate (fraction of the highest probability choices using (15.36) that are actually chosen by the 30% of subjects) and the log likelihood (denoted by LL). Ho et al. (2007) also report results for the Bayesian information criterion (BIC) for the estimates, based on 70% of the data. This criterion trades off the larger number of parameters in a model against better explanatory power.

In Table 15.4, for any game, the highest hit rates and the best fits are shown in bold. When two hit rates or two values of LL are not statistically different then they are both shown in bold. From left to right in Table 15.4, the theories tested are the following. (1) Thinking+fEWA is fEWA in which the initial attractions are computed from a cognitive hierarchy model (1 parameter). (2) fEWA (1 parameter). (3) EWA (5 parameters). (4) Weighted fictitious play, which is a belief-based model (2 parameters). (5) Reinforcement learning with payoff variability as in Roth and Erev (1995) and Erev and Roth (1998). (6) QRE, which is short for quantal response equilibrium

Table 15.4 Hit rates and the best fits arising from an application of various learning models to different games.

Game	Thinking +fEWA		fEWA		EWA		Weighted fict. play		Reinf. with PV		QRE	
	%Hit	LL	%Hit	LL	%Hit	LL	%Hit	LL	%Hit	LL	%Hit	LL
Cont'l divide	45	-483	47	-470	47	-460	25	-565	45	-557	5	-806
Median action	71	-112	74	-104	79	-83	82	-95	74	-105	49	-285
p-BC	8	-2119	8	-2119	6	-2042	7	-2051	6	-2504	4	-2497
Price matching	43	-507	46	-445	43	-443	36	-465	41	-561	27	-720
Mixed games	36	-1391	36	-1382	36	-1387	34	1405	33	1392	35	1400
Patent race	64	-1936	65	-1897	65	-1878	53	-2279	65	-1864	40	-2914
Pot games	70	-438	70	436	70	137	66	471	70	-429	51	-509
Pooled	50	-6986	51	-6582	49	-7100	40	-7935	46	-9128	36	-9037
KS p-BC			6	-309	3	-279	3	-279	4	-344	1	-346

Source: Camerer and Ho (2014). The sample sizes for the games in the first column are as follows: 315, 160, 580, 160, 960, 1760, 739, 4674 (pooled), 80.

(see Part 4 of the book for a discussion). Based on the hit rates, the three best theories are EWA, Thinking+fEWA and fEWA (highest in 6 out of 7). Reinforcement learning just follows behind these models (highest in 4 out of 7), followed by weighted fictitious play (highest in 3 out of 7), while QRE is the worst (highest in 1 out of 7). Based on the LL, EWA is best (highest in 4 out of 7) followed by reinforcement learning (highest in 2 out of 7). QRE is the worse in this respect too.

When the data is pooled across all games and parameters are estimated based on pooled data (except for value of λ that is estimated for individual games), then Thinking+fEWA gets the highest hit rate and fEWA the highest LL. Thus on both criteria, fEWA outperforms other models for pooled out-of-sample data. Below the row labeled pooled data in Table 15.4, the authors also report results from an analysis of the p-beauty contest game from Kocher and Sutter (2005); fEWA again performs better than the other models.

15.4.5 Sophisticated EWA learning and strategic teaching

The EWA model contradicts observed behavior in several respects.

1. It is based on players who use *adaptive learning*. Thus, they do not use knowledge of the payoffs of others, even if such knowledge is available, in order to outguess them. Experimental evidence indicates that players use information on the payoffs of others (Cachon and Camerer, 1996; Mookherjee and Sopher, 1997). When players have the option of using or ignoring the payoff information about others, they often opt to use it. This is implemented in some experiments by checking if players opt to open boxes containing payoffs of others (Camerer et al., 1993; Costa-Gomes et al., 2001).
2. Consider the *p-beauty contest*. A group of players are asked to guess a number between 0 and 100. The winner chooses a number closest to $p > 0$ times the average. When $p \in$

(0, 1], the Nash equilibrium is for each player to choose 0 (see Part 4 of the book). In p-beauty contests, it is found that experienced players converge to the Nash equilibrium faster than inexperienced players (Ho et al., 1998). Camerer et al. (2002) take this as suggestive evidence that players learned about how others were learning.

In order to account for these stylized facts, Camerer et al. (2002) propose the *EWA learning with sophistication* model. A proportion α of the players are sophisticated, and the rest, $1 - \alpha$, are adaptive, as in the fEWA model. The adaptive players use the fEWA model, while the sophisticated players form an estimate of the fraction of the adaptive players, $1 - \alpha'$ and take this information into account when forming their attractions. Sophisticated players may be overconfident and overestimate the proportion of adaptive players, $\alpha > \alpha'$, or they might be cautious, and underestimate this fraction, $\alpha < \alpha'$. Sophisticated players believe that other sophisticated players also believe that the proportion of adaptive players is $1 - \alpha'$.

Denote the type of the players by $\theta = s, a$ for, respectively, sophisticated and adaptive. The relation between probabilities and attractions for player i of type θ is given by

$$p_{t+1}^i(s_{ik}; \theta) = \frac{e^{\lambda A_t^i(s_{ik}; \theta)}}{\sum_{j=1}^{n_i} e^{\lambda A_t^i(s_{ij}; \theta)}}, \theta = s, a \text{ and } i \in N. \quad (15.52)$$

where $A_t^i(s_{ik}; \theta)$ is the attraction of strategy s_{ik} for player i of type θ at time t . Sophisticated players form attractions as follows

$$A_t^i(s_{ij}; \theta) = \sum_{k=1}^{n_{-i}} \left[\alpha' p_{t+1}^{-i}(s_{-i}^k; s) + (1 - \alpha') p_{t+1}^{-i}(s_{-i}^k; a) \right] \pi_i(s_{ij}, s_{-i}^k). \quad (15.53)$$

In (15.53), $s_{-i}^k \in S_{-i}$ is the k th strategy profile of others and there is a total of n_{-i} such strategies; $p_{t+1}^{-i}(s_{-i}^k; s)$ and $p_{t+1}^{-i}(s_{-i}^k; a)$ are, respectively, the probabilities of play of sophisticated and adaptive players, other than player i , at time $t+1$, which are determined by an expression analogous to (15.52). Given the interdependence between (15.52) and (15.53), the model is solved recursively. It can be shown that a Nash equilibrium and a quantal response equilibrium (QRE) are special cases of this model when $\alpha = \alpha' = 1$ (Camerer et al., 2002).

The model is applied to a ten period p-beauty contest game using data from Ho et al. (1998). Two values of p are considered in the experiments, $p = 0.7$ and $p = 0.9$. *Inexperienced subjects* first play a game with one value of p . They then play the second game with a different value of p ; in the second game they are referred to as *experienced subjects*. The results are shown in Table 15.5, drawing on Camerer and Ho (2014); Table 2 on p. 156 in Camerer et al. (2002) gives additional results on the QRE.

The results are as follows. When subjects are inexperienced, the adaptive EWA model generates estimates that are consistent with Cournot beliefs ($\delta = 0.9, \phi = 0$). The estimated value of ϕ is larger in the sophisticated EWA for both experienced and inexperienced subjects. Sophistication also improves the log-likelihood value for in-sample and out-of-sample observations for experienced and inexperienced subjects. Among inexperienced subjects, about 24% of the players are estimated to be sophisticated and for these players, $\alpha' = 0$, so they are overconfident. The estimated fraction of sophisticated subjects among experienced players is larger at 77% and these players are less overconfident; for these players $\alpha' = 0.41$. Hence, for experienced subjects, adaptive EWA fits much worse than sophisticated EWA. Hence, learning appears to occur in the

Table 15.5 Estimates for the parameters of the sophisticated and adaptive EWA models. Standard errors in parentheses.

	Inexperienced subjects		Experienced subjects	
	Sophisticated EWA	Adaptive EWA	Sophisticated EWA	Adaptive EWA
θ	0.44 (0.05) ²	0.00 (0.00)	0.29 (0.03)	0.22 (.02)
δ	0.78 (0.08)	0.90 (0.05)	0.67 (0.05)	0.99 (0.02)
α	0.24 (0.04)	0.00 (0.00)	0.77 (0.02)	0.00 (0.00)
α'	0.00 (0.00)	0.00 (0.00)	0.41 (0.03)	0.00 (0.00)
LL				
(in sample)	-2095.32	-2155.09	-1908.48	-2128.88
(out of sample)	-968.24	-992.47	-710.28	925.09

Source: Camerer and Ho (2014).

expected direction, however, the repeated p-beauty contest may be particularly amenable to these results, which need to be checked for other classes of games.

The matching protocol used in experiments, e.g., random matching in each round, or repeated matching with the same opponent is known to affect the outcome (Andreoni and Miller, 1993; Clark and Sefton, 1999). This is not predicted by the EWA model. When *sophisticated players* are matched with *adaptive players* (all players are adaptive in the standard EWA model) in a repeated game setting, the former usually have an incentive to teach the latter in a strategic manner. Such *strategic teaching* can take the form of deliberately playing strategies that could have a lower short-run payoff, but which alters the behavior of the adaptive player in a manner which is to the long-run benefit of the sophisticated player. This often leads to repeated game equilibria and behavior that is consistent with reputation formation.¹⁷

Camerer et al. (2002) discuss an application of a model of strategic teaching to a finitely repeated trust game. In 3×3 games played over 20 rounds, Hyndman et al. (2012) give empirical evidence to show that certain players emerge as “teachers.” These players repeatedly play the Nash outcome when it is not in their best short-term interests. However, by their actions, they are able to shepherd learning in the direction of a Nash equilibrium. In conditions where such teaching is more difficult, there is less frequent convergence to Nash behavior.

Unlike adaptive models (e.g., reinforcement learning, belief-based learning, and EWA) that are backward looking, models of strategic teaching are forward looking. For this reason, they are sometimes referred to as models of *anticipatory learning*. Other models in this class include models of rule-based learning that we consider in Section 15.7.

¹⁷ See for instance, Fudenberg and Levine (1989), Watson (1993).

15.5 Learning models: two points of view

We complete this discussion by giving two opposing points of view on the class of EWA models, which includes reinforcement learning and belief-based models as special cases. The first point of view is already implicit in our discussion above; we may call this the Colin Camerer view. The idea is to construct ever improving models in the EWA class. If this improvement comes about by adding extra parameters (a counter example is the fEWA), then we use statistical criteria such as the AIC and the Bayesian information criteria to trade off the extra explanatory power against the larger number of parameters. Furthermore, we may use estimated parameters based on in-sample estimation to test the model based on out-of-sample data.

A second point of view may be associated with the work of the eminent mathematical psychologist, Amnon Rapoport.¹⁸ This view has two components.

1. Individuals exhibit *heterogeneity in learning*, and there are important context dependent features in learning. Heterogeneity in learning among individuals reflects differences in the learning models used and differences in their cognitive processes. These effects cannot be captured by simply allowing for a distribution of parameters in a common learning model, such as the EWA model. Amnon Rapoport argues that:¹⁹ “any attempt to account for the complex adaptive learning process over iterations of the stage game is doomed to fail. It is like attempting to place multiple individuals of different height into a Procrustean bed.”
2. The class of EWA models, except models of strategic teaching, is backward looking. It ignores the forward-looking aspects of human learning and the ingenuity of human learning. For instance, the end effects in a prisoner’s dilemma game are not captured by purely backward-looking models. In other cases, individuals may sacrifice short-term gains to credibly signal an intention to shift the group behavior in a direction that is mutually beneficial to the players. Indeed, sophisticated EWA learning and strategic teaching that we consider in Section 15.4.5 already considers this possibility in a formal manner. So there are already attempts underway to address this concern.

In the Amnon Rapoport view, statistical tests such as the AIC and the Bayesian information criteria are not very useful if the class of models that is being considered is mis-specified. In this view, a mis-specified model omits heterogeneity in learning/cognitive processes. The alternative is to write context-dependent learning models. Examples include: variations of the reinforcement model, augmented for context and individual variations in learning, to apply to market entry games (Rapoport et al., 2000, 2002), two-player bargaining games (Rapoport et al., 1998), common pool resource problems (Botelho et al., 2014), route choice in congestible networks (Mak et al., 2014), and a three-player centipede game (Rapoport et al., 2003). In each case, these augmented models are successful in explaining the particular features of the experimental results.

As an example, consider the market entry games in Rapoport et al. (2000, 2002). In this model, players simultaneously decide to enter two different markets. Players use a simple cutoff rule to make their entry decision, based on the combined capacity in both markets. Conditional on entry, players probabilistically decide to enter each market; these probabilities, in turn, are based on propensities to enter each market. In the manner of reinforcement learning, the actual observations are used, in each round, to update the cutoff rules to determine entry as well as the propensities to enter the markets. In line with earlier experiments on single market entry games, the

¹⁸ I am grateful to Amnon Rapoport for discussing and clarifying his position on these issues.

¹⁹ The source for this quote is personal communication.

authors find a high degree of implicit coordination on entry that conforms well with equilibrium predictions. The proposed learning model also conforms well with actual data.

The Rapoport et al. (2003) experiments extend the standard two-player centipede games (see Part 4 of the book) to three players, more rounds of learning (from 10 rounds to 60 rounds), higher stake payoffs,²⁰ and terminal payoffs of zero. With these changes, the authors find that, in their four sessions, the equilibrium prediction (play down at the first node) holds in 46.3%, 39.3%, 30.3%, and 40.7% of the cases.

They then propose a reinforcement learning model to explain their results. In this model, players begin with some initial probabilities of playing “down” at each of the nodes in the centipede game. They then update the probabilities at the end of each round, based on the observed play of “down” in that round. The model is fairly parsimonious. It requires a specification of the initial probabilities and the single parameter of the reinforcement learning model. The predicted conditional probability of moving down at the first node increases over successive rounds of learning, as does the observed probability.

One disadvantage of this approach may be that we could have a plethora of learning models that are designed for each situation; something that most economists would find unattractive, given their approach to model building. However, against this criticism is the maintained hypothesis in this view that it is impossible to construct general models anyway. At a minimum, these results highlight the importance of considering individual differences in learning and cognition, and the necessity of incorporating forward-looking elements into existing learning models.

15.6 The class of learning direction theories

In this section, we consider learning models that are either inspired by *learning direction theory*, or are close relatives.

15.6.1 *Learning direction theory*

There has been growing recent interest in *learning direction theory*. A popular analogy used to describe the theory is the problem faced by an archer (Selten and Stöcker, 1986). The best response of an archer is to hit the absolute middle of the target. Suppose instead that the archer hits to the left of the target. In the next step, the archer realizes that he needs to aim a little bit to the right. Thus, in each period, the player appears to best respond to last period’s play, as under Cournot dynamics—a form of *ex-post rationality*. An attractive feature of this theory is the low informational requirement on account of the limited objective to predict only the direction of change.

Selten et al. (2005, p. 6) write “Learning direction theory is not meant to be a complete explanation of adaptive behavior. Therefore, it does not propose the hypothesis that ex-post rationality prevails in all cases. Sometimes other influences might move the decision in the opposite direction. However, it is assumed that ex-post rationality is the strongest influence.” The theory is qualitative. It does not specify the magnitude of adjustment from one period of learning to another—the interest lies only in the direction of change.

²⁰ In the four move centipede game, the maximum payoff was 200 times higher relative to the McKelvey and Palfrey (1992) two-player centipede experiments considered in Part 4; the corresponding figure for the six move game (see also Part 4) was 50.

In successive steps of learning, players may converge to the target, so, for instance, the archer may eventually hit the bulls eye. The objective of *impulse balance theory*, outlined in Selten et al. (2005) is to describe the long-run consequences of such directional learning. The idea in this theory is to find an *impulse balance point* at which, on average, one is equally likely to move to the left or to the right. We defer a discussion of these issues to Section 15.6.2.

Suppose that player $i \in N = \{1, 2, \dots, n\}$ chooses strategy $s_i(t)$ at time t , while others choose the strategy profile $s_{-i}(t)$. Let the pure strategy set of player i be given by S_i . The *ex-post best response* of player i is

$$b_i(t) \in \arg \max_{s_i \in S_i} \pi_i(s_i, s_{-i}(t)).$$

Then, under learning direction theory, and as indicated in the example of the archer above, the strategy chosen by player i in period $t + 1$, $s_i(t + 1)$, lies between $s_i(t)$ and $b_i(t)$. In relation to other learning models, two comments are in order. First, the notion of moving in the direction of improving payoffs is a feature of many other learning models, such as belief-based models, aspiration adaptation theory, or the EWA model. Second, and as noted by Camerer (2003), suppose that the current action is not the best response and that there are several “intermediate actions” between the current action and the best action. It is quite conceivable that some of these intermediate actions might have a very low payoff. Learning direction theory does not predict which, if any, of these intermediate low payoff actions will be played. However, belief-based, or EWA learning models are able to make such predictions, giving, in some cases, a sharper predictive alternative to learning direction theory.

The comparison of learning direction theory with some of the other models that we have considered is as follows. Unlike reinforcement learning, the central issue for players who use learning direction theory is not the payoff in the previous period. Instead, they ask if they could they have achieved a higher payoff if they had followed another action? This consideration of *forgone hypothetical payoffs* is similar to some belief-based models and EWA. However, unlike the adaptive learning models that we have considered, learning direction theory does not predict the probabilities with which players play their strategies.

In order to test the predictions of learning direction theory, Selten and Stöcker (1986) had six groups of subjects in experiments play 25 supergames. Each supergame was a ten period repeated prisoner’s dilemma game. They were interested in examining the variations in the defection strategy on account of learning, as players play successive supergames. The neoclassical prediction is that players must defect in the first round.

Cooperative behavior in these experiments is defined as the play of the cooperative strategy for at least four periods, followed by a defection from at least one player, which, in turn, is followed by mutual defection from both players until the tenth period. On this definition, most players, particularly experienced players (those who had played several supergames) choose to cooperate.

Let $k \in \{1, 2, \dots, 10\}$ be the earliest period that a player wishes to play the strategy defect. This can be revealed through an actual defection, or if a defection is not observed, through a written response from the player at the end of the supergame. For each of the six groups, k fell as the players successively played the supergames. Table 15.6 shows the average value of k for six groups of experienced subjects for supergames 13–19 and 20–25.

Classical game theory predicts $k = 1$. However, from Table 15.6, for experienced subjects, $k > 4$, hence, players play the cooperative outcome (this also holds when one looks at data that is disaggregated by supergames). There is variability across groups. For instance, on average, in supergames 13–19, there is almost perfect cooperation in groups V and VI. Even in group I where

Table 15.6 Average defection period, $k=1, 2, \dots, 10$, for experienced subjects.

Supergames	I	II	III	IV	V	VI	Average
13–19	7.1	8.6	9.9	7.3	9.6	10.0	8.8
20–25	5.9	7.8	9.6	6.4	8.8	8.7	7.8

Source: Selten and Stöcker (1986).

the average cooperation is the lowest, on average, no player defects until round 7 of the ten round supergame. There is a reduction in k , typically by about one round on average, across the two sets of average figures, across the supergames reported in Table 15.6.

The authors propose a model based on learning direction theory to organize the data. Suppose that players begin with some intended period of deviation, k . If they find that their opponent defected before them, then in the next supergame, they reduce k by one unit with probability α . In the case of simultaneous defection, players reduce k by one unit with probability β . Finally, if they find that they defected first, then in the next supergame they increase k by one unit with probability γ . The estimated values are $(\alpha, \beta, \gamma) = (0.50, 0.14, 0.23)$; these might be taken to be the estimates of learning direction theory for the finitely repeated prisoner's dilemma game.

Selten et al. (2005) show that learning direction theory can explain the data arising from the *takeover game*, due to Samuelson and Bazerman (1985), that we describe in detail in Part 4 of the book. In these experiments, subjects repeatedly place bids in auctions, over 100 periods. At the end of each period, they learn about the valuation of the firm that they were bidding for. Learning direction theory prescribes that if the bid turns out to be greater (lower) than the valuation, then it should be reduced (increased) in the next period. The results show that the majority of the changes are in the direction predicted by learning direction theory.

15.6.2 Impulse matching theories

In this section we are interested in two-player normal form games in which each player has two actions (2×2 games). The *impulse balance equilibrium* was introduced by Selten et al. (2005) and Selten and Chmura (2008). However, we first outline *impulse matching learning*, which is a closely related concept (Selten and Chmura, 2008; Chmura et al., 2012). Unlike traditional learning direction theories, that could only predict the direction of change, these newer concepts are able to predict the probabilities with which players play their strategies. Thus, they allow the predictions of learning direction theory to be directly compared with other theories that predict probabilities of play.

IMPULSE MATCHING LEARNING

The basic idea is that if in any round, a player were to choose an alternative strategy that could have given him a higher payoff, then in the next round, he receives an *impulse* to play the alternative strategy. Consider, for instance, the 2×2 game shown in Figure 15.11, which is game 3 from Selten and Chmura (2008). Suppose that in some round, the column player plays L and the row player plays D. However, the row player's best response would have been U, thus, in the next round, the row player receives an impulse to play U. One could calculate impulses based on the forgone payoff, as in theories that incorporate regret. However, Chmura et al. (2012) incorporate a form of loss aversion by first defining an *aspiration level* that is taken to be the *maximin payoff*.

	L	R		L	R
U	8, 6	0, 14		U	7.5, 6
D	7, 7	10, 4	→	D	7, 6.5
					8.5, 4

Figure 15.11 Original game (left) and transformed game (right) based on an aspiration level of 7 for the row player and 6 for the column player.

Source: Reprinted from *Games and Economic Behavior* 76(1): 44–73. Chmura, T., Goerg, S., and Selten, R. “Learning in experimental 2×2 games.” Copyright © 2012, with permission from Elsevier.

In the original game in the left panel in Figure 15.11, the maximin payoffs for the row and the column players are, respectively, 7 and 6. Payoffs that exceed the aspiration levels are termed as gains and those that fall below are termed as losses. Taking a value of the loss aversion parameter to be 2, gains above the aspiration level are cut by half. Consider, for instance, the payoff of 8 of the row player from (U, L). Relative to the aspiration level of the row player, 7, this is a gain of 1, which is cut by half. Thus, in the transformed matrix in Figure 15.11, the payoff of the row player is $7 + \frac{8-7}{2} = 7.5$. Applying this transformation to each payoff in the original matrix on the left gives the transformed matrix in the right panel in Figure 15.11. The payoffs used below shall all refer to the payoffs in the transformed matrix.

Suppose that the row and column players are labeled 1 and 2, respectively, and their respective strategy sets are $S_1 = \{s_{11}, s_{12}\}$ and $S_2 = \{s_{21}, s_{22}\}$. At any time $t = 1, 2, \dots$, suppose that the action of the opponent is s . Fixing s , let the payoff of player $j = 1, 2$ from his two actions s_{j1}, s_{j2} be $\pi_{j1} = \pi_j(s_{j1}, s)$ and $\pi_{j2} = \pi_j(s_{j2}, s)$. Then, the impulse from action s_{j2} to action s_{j1} for player $j = 1, 2$ is given by

$$r_{j1}(t) = \max \{0, \pi_{j1} - \pi_{j2}\}, j = 1, 2. \quad (15.54)$$

Similarly, the impulse from action s_{j1} to action s_{j2} is $r_{j2}(t) = \max \{0, \pi_{j2} - \pi_{j1}\}, j = 1, 2$. The impulses in any time period contribute to the *cumulative impulses*. The cumulative impulse for strategy i of player j , denoted by R_{ji} , is calculated as follows. Initial values of the cumulative impulses are $R_{ji}(1) = 0, i, j = 1, 2$. For any $t > 1$, the cumulative value of the impulse for strategy i of player j is

$$R_{ji}(t+1) = R_{ji}(t) + r_{ji}(t), j = 1, 2 \text{ and } i = 1, 2.$$

Player j plays the two pure strategies in $S_j = \{s_{j1}, s_{j2}\}$ with the following probabilities

$$p_{ji} = \begin{cases} \frac{R_{ji}}{R_{j1} + R_{j2}} & \text{if } R_{j1} > 0 \text{ and } R_{j2} > 0 \\ 0.5 & \text{otherwise} \end{cases}, j = 1, 2 \text{ and } i = 1, 2.$$

IMPULSE BALANCE EQUILIBRIUM

The idea of an *impulse balance equilibrium* is similar to impulse matching learning except for the calculation of impulses in (15.54). In this case, the analog of (15.54) is

$$r_{j1}(t) = \begin{cases} \max \{0, \pi_{j1} - \pi_{j2}\} & \text{if the chosen action is } s_{j2} \in S_j \\ 0 & \text{otherwise} \end{cases}, j = 1, 2,$$

where π_{ji} is the payoff to player j from the strategy s_{ji} , as given in the transformed matrix in Figure 15.11.²¹ Thus, there is an impulse from strategy s_{j2} to s_{j1} if the payoff would have been higher from playing s_{j1} instead of s_{j2} . A potential impulse in the other direction, $r_{j2}(t)$, can be explained in an analogous manner,

ACTION-SAMPLE LEARNING

In *action-sample learning*, players take a sample of past observations of the opponent's strategies and optimize with respect to the average (Selten and Chmura, 2008). More formally, let $A_t = \{a_1, \dots, a_m\}$ be a sample of size m of the past actions of the opponent, taken without replacement at time t . Let $\pi_{jk}(a_i)$ be the payoff of player $j = 1, 2$ from his action $s_{jk} \in S_j = \{s_{j1}, s_{j2}\}$ against the opponent's sampled action $a_i \in A_t$. Let $P_{jk}(t) = \sum_{i=1}^m \pi_{jk}(a_i)$ be the cumulative payoffs of the player when he uses the strategy s_{jk} , against the set of opponent's sampled actions, A_t . Then the probabilities of the two actions of player j , s_{jk} and s_{jl} , are given by

$$p_{jk}(t) = \begin{cases} 1 & \text{if } P_{jk}(t) > P_{jl}(t) \\ \frac{1}{2} & \text{if } P_{jk}(t) = P_{jl}(t), k, l = 1, 2, k \neq l. \\ 0 & \text{if } P_{jk}(t) < P_{jl}(t) \end{cases}$$

The initial probabilities of the two actions of player j are set at $p_{j1}(1) = p_{j2}(1) = 0.5$ until the opponent plays both of his possible actions at least once.

PAYOFF-SAMPLING LEARNING

The final learning model that we consider in this section is *payoff-sampling learning* (Selten and Chmura, 2008). In this model, players take two samples of equal size from observed past play. Any player, $j = 1, 2$, takes one sample of m observations in which he played his strategy s_{j1} in the past and another sample of m observations when he played the strategy s_{j2} . Denote by $T_{jk}(t) = \sum_{i=1}^m \pi_{jk}(a_i)$, $k = 1, 2$ the total payoff of player j from his strategy s_{jk} when he uses the sample relevant to s_{jk} . Then, the probabilities of the two actions of player j are given by

$$p_{jk}(t) = \begin{cases} 1 & \text{if } T_{jk}(t) > T_{jl}(t) \\ \frac{1}{2} & \text{if } T_{jk}(t) = T_{jl}(t), k, l = 1, 2, k \neq l. \\ 0 & \text{if } T_{jk}(t) < T_{jl}(t) \end{cases}$$

EMPIRICAL EVIDENCE ON LEARNING DIRECTION MODELS

Selten and Chmura (2008) pit various models of learning against each other by using 12 games, of which 6 were constant sum, and 6 were non-constant-sum games. There were 12 subject groups of 4 members each, for each constant-sum game and 6 subject groups for the non-constant-sum games. For each game, members of each subject group were randomly matched to play the game for 200 rounds. Based on their results, the authors ranked the various learning models from best to worse, as follows: impulse balance equilibrium, payoff-sampling equilibrium, action-sampling equilibrium, quantal response equilibrium, and Nash equilibrium. However, Brunner et al. (2011) reanalyzed the data and found computation errors. They undertook two revisions of the results. First, they augmented the other learning models also with loss aversion; impulse balance equilibrium, payoff-sampling equilibrium, and action-sampling equilibrium already incorporate loss

²¹ Ockenfels and Selten (2014) apply impulse balance equilibrium to the newsvendor problem that we considered in Part 1 of the book. Ockenfels and Selten (2005) apply this concept to explaining experimental results in private value auctions with feedback on losing bids.

aversion, as shown, for instance, in Figure 15.11. Results showed that the other learning models performed as well as impulse balance equilibrium. Second, in 2×2 matching pennies games where loss aversion plays no role, they find that quantal response equilibrium performs as well as impulse balance equilibrium. In their reply, Selten et al. (2011) contest both these claims, although they admit the computation errors.

It is against this backdrop that the paper by Chmura et al. (2012) runs a comprehensive horse-race between several competing one-parameter learning models, including the ones introduced in this section. They use the same 12 games, and a similar experimental structure to that in Selten and Chmura (2008) described above. One of the games is shown in Figure 15.11; the row player's strategies are U and D, while the column player's strategies are L and R. They run 500 simulations with computer agents to derive the predictions of the following models: impulse balance equilibrium, payoff-sampling equilibrium, action-sampling equilibrium, reinforcement learning, one-parameter self-tuning model, and Nash equilibrium.

Using their simulations and the data from the experiments, they propose the *mean quadratic distance*, Q , to rank the various models; Q is defined as follows

$$Q = \frac{1}{12} \sum_{i=1}^{12} \left(\frac{1}{500} \sum_{k=1}^{500} (\phi_{ik}^L - f_i^L)^2 + (\phi_{ik}^U - f_i^U)^2 \right),$$

where 12 is the total number of games, 500 is the number of simulations, f_i^L and f_i^U are the experimentally observed frequencies of play of L (by the column player) and U (by the row player). ϕ_{ik}^L and ϕ_{ik}^U are, respectively, the frequencies of play of L and U in game i and simulation k . Q lies between 0 and 1. A smaller value of Q shows a better fit with the data. The parameters of each model are chosen so as to minimize the corresponding value of Q for that model. For instance, for




Figure 15.12 Q values for various learning models. Stationary concepts are depicted in dark bars and learning models are depicted in light bars.

Source: Reprinted from *Games and Economic Behavior* 76(1): 44–73. Chmura, T., Goerg, S., and Selten, R. "Learning in experimental 2×2 games." Copyright © 2012, with permission from Elsevier.




Figure 15.13 Q values of various models with and without loss aversion.

Source: Reprinted from *Games and Economic Behavior* 76(1): 44–73. Chmura, T., Goerg, S., and Selten, R. "Learning in experimental 2×2 games." Copyright © 2012, with permission from Elsevier.

action-sample learning, a value of $m = 12$ (sample size) minimizes Q and for the one-parameter EWA model, a value of $\lambda = 0.2775$ (parameter used to compute probabilities, based on attractions) minimizes Q .

Figure 15.12 shows the value of Q for the various learning models considered. This allows a clear ranking of models from best to worse: impulse matching learning, action-sample learning, self-tuning EWA, reinforcement learning. The last two are barely distinguishable, however, reinforcement learning does better than EWA in constant-sum games, but the converse is true for non-constant-sum games. The Q value for EWA is 18 times that of impulse matching learning. Reinforcement learning and EWA are outperformed by the three stationary models (Nash, impulse-balance, action-sampling).

Is the superior performance of impulse matching due to loss aversion? Figure 15.13 gives a comparison of Q values when loss aversion is incorporated in the other models as well (by using the transformed matrix, as in the right panel of Figure 15.11). Impulse matching and reinforcement learning benefit by the incorporation of loss aversion, but it does not alter the ranking of models. Hence, in these results, the superior performance of impulse learning is not an artifact of loss aversion. The ranking is stable over time. The overall ranking is same for rounds 1–100, 101–200, and 1–200.

15.6.3 An explanation of the data from order statistic games

In this section we consider the work of Crawford (1991, 1995) and Crawford and Broseta (1998); their work is typically not included in this category, yet in important ways it is a close relative of learning direction theory. Recall from Part 4 of the book, the experimental results on *order*

statistic games from Van Huyck et al. (1990, 1991). In these games, a set of $N = \{1, 2, \dots, n\}$ players simultaneously choose a strategy (a positive integer) $s_i \in \{1, 2, \dots, \bar{s}\}$, where $\bar{s}$ is the maximum number that can be chosen. The payoff of player i is given by

$$\pi(s_i, M) = aM - b(M - s_i)^2 + c; \quad a > 0, b > 0. \quad (15.55)$$

In (15.55), M is an order statistic, which could be a *median*, or *minimum value*. There are multiple symmetric Nash equilibria that can be Pareto ranked. In the Pareto optimal action, all players should choose $\bar{s}$. In general, the payoff is increasing in one's own choice, but decreasing in the deviation of that choice from the order statistic.

When M is the median, two interesting results are that initial choices are quite dispersed and there is a strong effect of history—the period 10 median is always the same as the median in the first period of play. When $\bar{s} = 7$, the typical choice is 4 or 5. When M is the minimum, initial choices are again dispersed, suggesting the lack of a common prior among the players. The equilibrium typically converges to a play of 1 by each player by round 10 and convergence is often fairly quick.

An initial explanation of these results was given in Crawford (1991) using an evolutionary explanation, while taking particular care to distinguish between the results from small and large groups. We now briefly focus on Crawford (1995). The underlying model is a belief-based learning model that explains the data from the games in Van Huyck et al. Players ignore the effect of their choice on the order statistic. Suppose that the choice of a strategy by player i at time t is given by s_{it} and the period t order statistic is M_t . The order statistic may be written as

$$M_t = f(s_{1t}, s_{2t}, \dots, s_{nt}), \quad t = 1, 2, \dots \quad (15.56)$$

The function f satisfies

$$f(c + ds_{1t}, c + ds_{2t}, \dots, c + ds_{nt}) = c + df(s_{1t}, s_{2t}, \dots, s_{nt}), \quad (15.57)$$

which is satisfied by the median and the minimum. Suppose that the initial choices of strategies are given by

$$s_{i0} = \alpha_0 + \varsigma_{i0}, \quad (15.58)$$

and the strategy choices in any period, $t \geq 1$, are given by

$$s_{it} = \alpha_t + \beta_t M_{t-1} + (1 - \beta_t) s_{it-1} + \varsigma_{it}, \quad \beta_t \in (0, 1] \text{ and } \lim_{t \rightarrow \infty} \alpha_t = 0. \quad (15.59)$$

In (15.58) and (15.59), α_t is a individual-invariant, but time-specific drift parameter. The variable ς_{it} is random with mean zero and variance $\sigma_{\varsigma_t}^2$ and represents individual-specific shocks to beliefs. A high value of β_t captures a greater responsiveness to experience (in this case, the observation of the order statistic in the previous period). It can be shown that (15.59) allows for fictitious play and Cournot best response dynamics. *Deterministic dynamics* allow for an initial dispersion of beliefs by $\sigma_{\varsigma_0}^2 > 0$, but no subsequent belief differences among players, so $\alpha_t = \varsigma_{it} = 0$ and $\sigma_{\varsigma_t}^2 = 0$ for $t = 1, 2, \dots$

We now show that even in the absence of any noise and absence of the drift parameter (deterministic evolutionary dynamics) the equilibrium is history-dependent, which is one of the central findings in the median action game.

Proposition 15.24 (*Deterministic evolutionary dynamics; Crawford, 1995*) Suppose that there exists some time $\tau \geq 1$ such that $\alpha_t = \varsigma_{it} = 0$ and $\sigma_{\varsigma_t}^2 = 0$ for $t = \tau, \tau + 1, \dots$.²² Then s_{it} approaches $M_{\tau-1}$ monotonically and $M_t = M_\tau$ for $t = \tau, \tau + 1, \dots$. This result is independent of the number of players, n , and the functional form of f , provided (15.57) holds.

The intuition behind Proposition 15.24 is that players do not have idiosyncratic beliefs and using (15.59) we have

$$s_{it} - M_{t-1} = (1 - \beta_t)(s_{it-1} - M_{t-1}). \quad (15.60)$$

From (15.60), the signs of $s_{it} - M_{t-1}$ and $s_{it-1} - M_{t-1}$ are identical, but $|s_{it} - M_{t-1}| < |s_{it-1} - M_{t-1}|$. Hence, over time, $s_{it} \rightarrow M_{\tau-1}$, which occurs independently of n and f . However, under deterministic dynamics one can determine which of many possible steady states will occur by invoking the role of history, e.g., by invoking differences in the value of M_τ in Proposition 15.24.

An alternative approach is to examine stochastic dynamics in which the strategies of players are subject to random mutations. Under the assumptions that $\alpha_t = 0$, $\beta_t = \beta$, and $\sigma_{\varsigma_t}^2 = \varepsilon > 0$ in (15.59) for $t = 1, 2, \dots$ the stochastic dynamics are ergodic and initial history does not play a role; however, the most recent history plays a role. When the probability of mutations approaches zero, the long-run equilibrium of the ergodic distribution approaches one of the steady states (that arise in the absence of mutations) with probability one. The long-run steady state depends on the difficulty of moving from various steady states to the basin of attraction of other steady states; the stochastically stable steady state is the one whose basin of attraction is the hardest to escape from. This is essentially the framework that we consider in Chapter 16. The next proposition summarizes the results for alternative order statistics in relation to the median for the games in Van Huyck et al. (1990, 1991).

Proposition 15.25 (*Crawford, 1995*) Suppose that $\alpha_t = 0$, $\beta_t = \beta$, and $\sigma_{\varsigma_t}^2 = \varepsilon > 0$ for $t = 1, 2, \dots$. Consider the games in Van Huyck et al. (1990, 1991) and an order statistic that satisfies (15.56) and (15.57).

- (i) If the order statistic is below (above) the median, the long-run equilibrium assigns probability 1 to the lowest (highest) possible strategy choice.
- (ii) If the order statistic equals the median, then the long-run equilibrium assigns positive probability to all every equilibrium.

Recall that when the order statistic is the minimum, then the equilibrium is a symmetric strategy choice of the lowest possible number, 1. This is consistent with the results in Proposition 15.25. However, the result in Proposition 15.25 is history independent, while the empirical results demonstrate the strong effect of history.

In order to incorporate history dependence, let us consider adaptive dynamics in the sense that $\lim_{t \rightarrow \infty} \sigma_{\varsigma_t}^2 = 0$. In contrast, under stochastic evolutionary dynamics, the variance in beliefs is constant. However, the assumption $\lim_{t \rightarrow \infty} \sigma_{\varsigma_t}^2 = 0$ allows for heterogeneous belief shocks that converge to zero, hence, giving rise to non-stationary transition dynamics.

²² This is the sense in which we have deterministic evolutionary dynamics.

Proposition 15.26 (Crawford, 1995) Suppose that $\lim_{t \rightarrow \infty} \alpha_t = 0$, $\sigma_{s_0}^2 > 0$, $\lim_{t \rightarrow \infty} \sigma_{s_t}^2 = 0$ and $\beta_t = \beta$, then the unique solution to (15.58) and (15.59), is given by:

$$\begin{aligned} s_{it} &= \alpha_0 + \sum_{j=0}^{t-1} \beta g_j + z_{it} \\ M_t &= \alpha_0 + \sum_{j=0}^{t-1} \beta g_j + g_t \end{aligned}$$

where

$$z_{it} = \sum_{j=0}^t (1 - \beta)^{t-j} \text{ and } g_j = f(z_{1j}, z_{2j}, \dots, z_{1nj}). \quad (15.61)$$

Proposition 15.26 allows for drift ($\alpha_t \neq 0$) that disappears in the time limit ($\lim_{t \rightarrow \infty} \alpha_t = 0$). Thus, the eventual outcome is determined by cumulative drift up to that point and as belief uncertainty is eliminated ($\lim_{t \rightarrow \infty} \sigma_{s_t}^2 = 0$) on account of learning by the players, the system settles down into one of the asymptotically stable states. In the absence of noise, it stays there forever. In the large-group minimum-action games in Van Huyck et al. (1990), the strong downward drift causes the outcome to converge to the minimum possible and history plays no role. By contrast, in the median action games in Van Huyck et al. (1991), there is no drift and history determines the outcomes.

We can use Proposition 15.26 to compute the expected values of the players' choices and the order statistic.

Proposition 15.27 (Crawford, 1995) Suppose that Proposition 15.26 holds, and that the following conditions are satisfied. (i) The z_{it} defined in (15.61) are iid across players and have a common variance across players that equals $\sigma_{z_t}^2$. (ii) $\frac{z_{it}}{\sigma_{z_t}}$ has mean zero and unit variance (iii) Let $\mu_t = Ef \left(\frac{z_{1t}}{\sigma_{z_t}}, \dots, \frac{z_{nt}}{\sigma_{z_t}} \right)$.

(a) Under these conditions the expected values of s_{it} and M_t , respectively, Es_{it} and EM_t take a simple form.

$$Es_{it} = \alpha_0 + \beta \sum_{j=0}^{t-1} \sigma_{zj} \mu_j \quad (15.62)$$

$$EM_t = \alpha_0 + \beta \sum_{j=0}^{t-1} \sigma_{zj} \mu_j + \sigma_{zt} \mu_t. \quad (15.63)$$

(b) In addition to the assumptions stated above, let $\frac{z_{it}}{\sigma_{z_t}}$ be normally distributed, so $\mu_t = \mu$ for all t and let $\lim_{t \rightarrow \infty} \sum_{j=0}^t \sigma_{zj} = \theta$, then

$$\lim_{t \rightarrow \infty} Es_{it} = \lim_{t \rightarrow \infty} M_t = \alpha_0 + \beta \mu \theta. \quad (15.64)$$

Proposition 15.27(a) shows that the drift in the model (second terms in (15.62), (15.63)) depends on the behavioral parameters of the model as well as the variance in the beliefs of the players. (15.64) shows the mean level of coordination (the expected strategy of each player equals the order statistic which is a symmetric Nash equilibrium of the game). Thus, mean coordination also depends on the behavioral parameters of the model, number of players, and the dispersion of player beliefs (as represented by the terms μ, θ). Consider the games in Van Huyck et al. (1990, 1991). In the median and the random-pairing minimum games,²³ due to symmetry, $\mu = 0$, thus, (15.64) implies that $\lim_{t \rightarrow \infty} Es_{it} = \lim_{t \rightarrow \infty} M_t = \alpha_0$ and there is no predicted drift.

²³ In the random pairing minimum games, subjects only observed the current pair's minimum. However, given the payoff parameters, it is optimal for risk-neutral players to try to match the population minimum, of which the pair minimum can be viewed as a noisy estimate.

However, for the large-group minimum treatment, Crawford (1991) shows that $\mu = -1.74$, thus, there is a downward trend towards the minimum in large groups, which is consistent with the experimental data. This leads to a lack of history dependence in this treatment (unlike the median action game). This is confirmed in simulated trials when the game is run 500 times (the Van Huyck et al. experiments were run for 9 trials).

15.7 Rule-based learning

In most models of learning, such as reinforcement learning, fictitious play, EWA, individual players learn about which strategies to use. However, it might be that in many contexts, players also learn about *which rules to use*. Rules could themselves be alternative models of learning such as reinforcement learning, belief-based learning, EWA, aspiration adaptation learning, level- k reasoning, tit-for-tat, and maximin. Models in which players choose probabilistically among a set of rules in each time period, and update the probabilities of choosing alternative rules, based on the history of play, are known as models of *rule-based learning*; for instance, Stahl (1999, 2000, 2001), Stahl and Haruvy (2002).

As motivation for rule learning, Stahl and Haruvy (2002) give the following example of human behavior in a 3×3 coordination game. Consider a single population of players in which players are randomly matched to play the bilateral cooperation game shown in Figure 15.14. The strategy C is dominated by a mixed strategy in which strategies A and B are played with identical probabilities.

The experimental results over various rounds of play are shown in Figure 15.15. The vertical and horizontal axes are, respectively, the probabilities of playing B and A. Any point in the probability triangle represents a triple (p_A, p_B, p_C) , where p_i is the probability of playing $i = A, B, C$, and p_C is found residually from $p_C = 1 - p_A - p_B$. The player is indifferent between A and B when $p_B = 0.5 - \frac{1}{3}p_A$; this equation is shown as the dotted line, or *separatix* in Figure 15.15. Above the separatix, B is preferred to A, and below it A is preferred to B. Of particular interest is the dynamic path shown in the region $p_B < 0.5 - \frac{1}{3}p_A$ that begins at 1. Since it is in the region where A is preferred, *myopic rational best response* (e.g., belief learning and EWA) implies that play should continue to stay in this region. However, from rounds 4 to 5, play jumps across the separatix and eventually converges to all play B ($p_B = 1$). This is inconsistent with myopic best response models and provides one rationale for studying rule-based learning.

	A	B	C
A	20, 20	0, 0	60, 10
B	0, 0	60, 60	0, 25
C	10, 60	25, 0	25, 25

Figure 15.14 A 3×3 coordination game.

Source: Reprinted from *Journal of Mathematical Psychology* 46(5): 531–53. Stahl, D. and E. Haruvy, "Aspiration-based and reciprocity-based rules in learning dynamics for symmetric normal-form games." Copyright © 2002, with permission from Elsevier.




Figure 15.15 Choice frequencies over various rounds of play.

Source: Reprinted from *Journal of Mathematical Psychology* 46(5): 55–53. Stahl, D. and E. Haruvy, "Aspiration-based and reciprocity-based rules in learning dynamics for symmetric normal-form games." Copyright © 2002, with permission from Elsevier.

Binmore and Samuelson (1997) provide one possible explanation of similar, separatix crossing, behavior. Players may have aspiration levels and if the current payoff is below the aspiration level, they may engage in experimentation with their next round strategies. For instance, unless a large fraction of the population plays the strategy C, players who choose A as their best response (rounds 1 to 4 in Figure 15.15) may find that their payoffs are mostly low numbers, 20 or 0, possibly below their aspiration level. If players then experiment with B, provided a high number of other players play B, they would experience a payoff of 60, possibly above their aspiration level. They then stop experimenting and all play B may be the eventual outcome. However, the experimental results of Stahl and Haruvy (2002) do not support this particular channel of aspiration based learning.

15.7.1 A formal model of rule learning

Let us now consider a formal model of rule learning. Consider a two-player game $G(2, S, \pi)$ in which (i) players from the same population are randomly matched together in bilateral encounters, and (ii) in each period, every player is matched with every other player. The pure strategy set of each player is given by $S = (s_1, s_2, \dots, s_n)$. The probability distribution of pure strategies in S belongs to the set Σ , with typical element σ .

Denote by p^t , the empirical distribution of pure strategies in the population in period t . For any $s \in S$, $\pi(s, p)$ is the expected payoff of a player from playing the strategy s when the frequency distribution of pure strategies in the population is given by p . Denote by p^0 , the mixed strategy $(\frac{1}{n}, \frac{1}{n}, \dots, \frac{1}{n})$ in which all pure strategies are played with an identical probability. At the

beginning of period t , the history of the game is given by $h^t = \{p^0, p^1, \dots, p^{t-1}\}$, where we have assumed that the null history at time $t = 1$, is p^0 .

At the end of each period, players receive information about their own payoffs, and the choices made by the other players. Thus, the information available to any player at the beginning of period t is given by $\Omega^t = (G, h^t)$.

Definition 15.20 A behavioral rule, ρ_j , is a mapping from the available information to mixed strategies, i.e., $\rho_j : \Omega^t \rightarrow \Sigma$. We denote by $\rho_j(\Omega^t)$ the mixed strategy that is generated by rule ρ_j when the available information is Ω^t . Denote the set of all behavioral rules under consideration by $R = \{\rho_1, \rho_2, \dots, \rho_r\}$.

At any time t , the rules in the set R are selected using a probability measure φ such that $\varphi(\rho_j, t)$ gives the probability of choosing the rule $\rho_j \in R$ at time t . We assume that there are propensities or attractions of using rules (formally defined below) so that $w(\rho_j, t)$ is the log propensity to use rule $\rho_j \in R$ at time t . We assume that

$$\varphi(\rho_j, t) = \frac{e^{w(\rho_j, t)}}{\sum_{\rho_k \in R} e^{w(\rho_k, t)}}. \quad (15.65)$$

Thus, each rule is chosen with the probability given by (15.65). In turn, each rule prescribes probabilistic play of pure strategies that is given in Definition 15.20. Thus, at time t , the probability that pure strategy $s_k \in S$ is played, $p_k(t)$, is given by

$$p_k(t) = \sum_{\rho_j \in R} \varphi(\rho_j, t) \rho_j(s_k), \forall s_k \in S,$$

where $\rho_j(s_k)$ is the probability with which strategy s_k is played under rule ρ_j . Thus, $p^t = (p_1(t), p_2(t), \dots, p_n(t))$.

Rules that perform better in one period are more likely to be chosen in the next period. This idea, sometimes known as the *law of effect*, is implemented in the model by the following evolution of log propensities that is similar to reinforcement learning,

$$w(\rho_j, t+1) = \beta_0 w(\rho_j, t) + g(\rho_j, \Omega^{t+1}), t \geq 1, \beta_0 \in (0, 1] \text{ and } \forall \rho_j \in R, \quad (15.66)$$

where $g(\rho_j, \Omega^{t+1})$ is the reinforcement to the propensity of playing ρ_j . $g(\rho_j, \Omega^{t+1})$ is defined as the expected payoff that would have occurred if at time t the individual had played the rule ρ_j when the actual frequency of strategies played was p^t , thus,

$$g(\rho_j, \Omega^{t+1}) = \beta_1 \pi(\rho_j, p^t), \beta_1 > 0. \quad (15.67)$$

Equations (15.66) and (15.67) feed into (15.65) to determine the probabilities of drawing various rules in any time period. These probabilities are identical for all players. However, for each player, in each period, a separate and independent draw from the distribution φ determines the choice of a rule. This is the only source of heterogeneity among individuals in the population. In order to operationalize the model, we need to specify two elements. (1) The set of rules, R , and (2) the initial vector $w(\rho, 1)$ of log propensities to play all the rules.

The model allows the specification of a wide array of rules.

Example 15.8

- (i) Replicator dynamics obtains by setting the number of rules equal to the number of strategies, so $r = n$, and for each $\rho_j \in R$, and for all t

$$\rho_j(s_t) = \begin{cases} 1 & \text{if } s_t = s_j \\ 0 & \text{otherwise} \end{cases}.$$

Thus each rule prescribes one hardwired pure strategy in the set S .

- (ii) Belief-based models such as Cournot dynamics arise when at any time t , $R = \{\rho_1\}$ and ρ_1 prescribes a best response to the time $t - 1$ strategy of the opponents, i.e.,

$$\rho_1(s_k) = \begin{cases} 1 & \text{if } s_k \in \arg \max_{s \in S} \pi(s, p^{t-1}) \\ 0 & \text{otherwise} \end{cases}.$$

We now describe a class of rules, the family of *evidence based rules*. In these rules, players base their choice of a rule on the basis of the *evidence* for and against a rule. In Stahl and Haruvy (2002), evidence, in turn, is generated from the following three archetypical theories of the mind.

(1) *Null model of other players*: In this case, for the initial period $t = 1$, and based on the principle of insufficient reason, players think that each action of the opponent is equally likely. Then the evidence in favor of each action $s_k \in S$ in period $t = 1$ is given by $y_1(s_k, \Omega^1) = \pi(s_k, p^0)$, where $p^0 = (\frac{1}{n}, \frac{1}{n}, \dots, \frac{1}{n})$.

Let p^1 be the vector of empirically observed frequencies of play in period $t = 1$; p_k^1 is the observed frequency of s_k . For period $t = 2$, the evidence in favor of each action, $s_k \in S$, arises by using simple distributed lag forecasting, so $y_1(s_k, \Omega^2) = \pi(s_k, (1 - \theta)p_k^0 + \theta p_k^1)$, where $\theta \in [0, 1]$. Letting $q_k^0(\theta) = p_k^0$, we can write the evidence as $y_1(s_k, \Omega^2) = \pi(s_k, q_k^2(\theta))$, where $q_k^2(\theta) = \theta q_k^0 + (1 - \theta)p_k^1$. Using this recursive chain, in general, for any $t \geq 1$, the evidence in favor of action $s_k \in S$ is²⁴

$$y_1(s_k, \Omega^t) = \pi(s_k, q_k^t(\theta)), \quad (15.68)$$

where

$$q_k^t(\theta) = \theta q_k^{t-1} + (1 - \theta)p_k^{t-1}, \theta \in [0, 1]. \quad (15.69)$$

The vector q^t is defined by $(q_1^t, \dots, q_k^t, \dots, q_n^t)$

(2) *Level-2 model of others*: In Part 4 of the book, we discussed level- k models at length. Stahl and Wilson (1995) formulated level-2 players who believe that all other players are level-1. In any period t , level-1 players play a best response to $q^t(\theta)$, defined in (15.69). Denote this best response by $b(q^t(\theta)) \in \Sigma$, which puts equal probability on all pure strategy best responses to $q^t(\theta)$, but zero probability on all other pure strategies. This may be termed as a *second order theory of the mind*. Thus, the level-2 evidence in favor of action $s_k \in S$ at time t is the expected payoff from playing s_k , i.e.,

$$y_2(s_k, \Omega^t) = \pi(s_k, b(q^t(\theta))). \quad (15.70)$$

²⁴ Notice that using this definition we have $q^1(\theta) = \theta q^0 + (1 - \theta)p^0 = q^0(\theta)$.

(3) *Nash equilibrium theory*: In this case, the evidence in favor of action $s_k \in S$, at time t , is the expected payoff from playing s_k when others are assumed to play a Nash equilibrium, p^{NE} , of the game G , i.e.,

$$y_3(s_k, \Omega^t) = \pi(s_k, p^{NE}). \quad (15.71)$$

The three kinds of evidence, y_1 , y_2 , and y_3 in (15.68)–(15.71), are weighted using respective weights v_1 , v_2 , and v_3 to generate the average evidence in favor of action $s_k \in S$ at time t

$$\bar{y}(s_k, \Omega^t) = v_1 y_1(s_k, \Omega^t) + v_2 y_2(s_k, \Omega^t) + v_3 y_3(s_k, \Omega^t).$$

On account of ease of estimation, the multinomial logit specification is used to specify the probabilities of using various pure strategies. Define the vector $v = (v_1, v_2, v_3)$. The probability of choosing action $s_k \in S$ at time t for any player is

$$\hat{p}_k(\Omega^t, v, \theta) = \frac{e^{\bar{y}(s_k, \Omega^t)}}{\sum_{s_l \in S} e^{\bar{y}(s_l, \Omega^t)}}. \quad (15.72)$$

Equation (15.72) defines a mapping, $\hat{p} : \Omega^t \rightarrow \Sigma$, hence, using Definition 15.20, it is a behavioral rule.

Stahl (2000) discusses another possible theory of the mind, *follow the herd*. This corresponds to a level-0 player in a level- k model. Such an individual follows the herd in the sense that the past is followed with inertia, as in (15.69). The follow-the-herd theory specifies that the evidence in favor of any action $s_k \in S$ at time t

$$y_0(s_k, \Omega^t) = \ln q_k^t(\theta).$$

Thus, if the only theory of mind were follow-the-herd model, then in conjunction with (15.72) this will ensure that $\hat{p}_k$ at time t equals q^t . In general, one may introduce an extra weight v_0 and integrate y_0 into the definition of $\bar{y}$. In this case we define the vector $v = (v_0, v_1, v_2, v_3)$ and the parametric vector as (v, θ) . The model can be extended to allow for other behavioral rules, or to consider combinations of a subset of the rules.

The initial vector of propensities $w(\rho, 1)$ is specified as normally distributed in rule space with mean $(\bar{v}, \bar{\theta})$ and variance σ_w^2 to be determined from the data. In the experiments, subjects play one game for T periods followed by another game for T periods. Thus, one needs to specify the *transference of learning* from one game to the other, i.e., the vector of propensities $w(\rho, T+1)$ in the first period of the second game. This is specified as

$$w(\rho, T+1) = (1 - \tau) w(\rho, 1) + \tau w(\rho, T^+), \tau \in [0, 1], \quad (15.73)$$

where $w(\rho, T^+)$ indicates an update to propensities in period $T+1$ based on the new game. If $\tau = 1$, then the experience from the older game is irrelevant, while if $\tau = 0$, then the older experience is the only relevant information used; (15.73) allows for any convex combination of the two.

15.7.2 Evidence from rule-based models

Stahl (1999, 2000) and Stahl and Haruvy (2002) conduct a comprehensive horse race between the main models. We omit a discussion of the results in Stahl (1999) because subjects play each of the 15 games over two periods, which does not give them an adequate opportunity to learn. These results are superseded by the comprehensive exercise in Stahl (2000) who uses a computer search over billions of randomly generated 5×5 matrices to identify the most suitable games for testing.

One of these games is shown in Figure 15.16. Each player was matched with every other player and received the average payoff over all matchings. These games have the following features. The Nash equilibrium (strategy ne) is a strict equilibrium. Strategy b1 is a best response to the belief that the opponent randomizes uniformly among the various strategies. Strategy b2 is a best response to b1. The strategy wd corresponds to the “worldly type” in Stahl and Wilson (1995); in the model outlined above, this corresponds to the specification $v_1 = 0.015$, $v_2 = 0.035$, and $v_3 = 0.05$. The final strategy is either the maximax strategy (mx), or the strictly dominated strategy dm (dm is not applicable to Figure 15.16).

The results are shown in Table 15.7. Maximizing the log likelihood (LL) for each participant and then summing across the 91 participants gives an unrestricted LL value of -1896.44 . Uniformly random play ($\sigma_w^2 = \bar{v} = \beta_1 = 0$ and $\beta_0 = 1$), labeled “Noise” in Table 15.7, yields a LL of -4393.75 . The difference in the likelihoods multiplied with 2 is distributed as a χ^2 variable with 819 degrees of freedom (d.f.) and has a p-value lower than 10^{-587} , thus, the model is informative.

The choices of participants are typically not the Nash equilibrium, so consider Nash with choice errors. This involves introducing a parameter ε to take account of trembles in choice and setting $\bar{v}_1 = \bar{v}_2 = 0$, $\beta_0 = 1$, $\beta_1 = 0$, and $\sigma_w^2 = 0.25$. Then, for each participant, values of $(\bar{v}_2, \bar{\theta})$ that maximized the LL were found; summed across 91 participants, this gives -3917.14 . However, when we compare this LL value with the value under the unrestricted model, -1896.44 , Nash equilibrium predictions with noise are decisively rejected despite the higher number of parameters in the unrestricted rule learning model.

If we impose no rule learning ($\beta_0 = 1$, $\beta_1 = 0$), then the aggregate maximized LL does much better than Nash at -2303.01 , but is rejected in favor of rule-learning. At the individual level, no rule learning can be rejected for 54.9% of the participants. Cournot dynamics in which the only non-zero parameters are v_1 and θ is also rejected in favor of rule learning (and also in favor of no-

	b1	ne	b2	wd	mx
b1	19, 19	43, 28	96, 67	85, 40	85, 16
ne	28, 43	62, 62	88, 21	74, 58	24, 15
b2	67, 96	21, 88	38, 38	48, 0	38, 86
wd	40, 85	58, 74	0, 48	15, 15	92, 99
mx	16, 85	15, 24	86, 38	99, 92	79, 79

Figure 15.16 A 5×5 strategic form game from Stahl (2000). Key: b1 (best response to uniform), ne (Nash equilibrium), be (best response to b1), wd (SW “worldly” choice), mx (maximax).

Source: Reprinted from *Games and Economic Behavior* 32: 105–38. Stahl, D. O. “Rule learning in symmetric normal-form games: theory and evidence.” Copyright © 2000, with permission from Elsevier.

Table 15.7 Fits of various learning models based on a comparison of the log likelihoods.

Hypothesis	LL	d.f.	p<
Noise	-4393.75	819	10^{-587}
Nash	-3917.14	637	10^{-486}
No rule lrn	-2303.01	273	10^{-54}
Cournol	-2829.30	637	10^{-120}
Homogeneous	-3060.88	810	10^{-146}
$\bar{v}_0 = 0$	-2014.41	91	10^{-14}
$\bar{v}_1 = 0$	-2315.34	91	10^{-120}
$\bar{v}_2 = 0$	-2022.61	91	10^{-16}
$\bar{v}_3 = 0$	-1987.33	91	10^{-7}
$\bar{\theta} = 0$	-2167.25	91	10^{-64}
$\beta_1 = 0$	-2151.30	91	10^{-58}
$\tau = 0$	-2125.14	91	10^{-49}
$\beta_0 \leq 1$	-2012.66	91	10^{-13}
Unrestricted	-1896.44	0	

Source: Stahl (2000).

rule learning). Finally, significant population heterogeneity is found and each of the parameters of the rule learning model is statistically significant in the aggregate.

Stahl and Haruvy (2002) introduce several further rules. They find no evidence for *aspiration based experimentation*, *reciprocity-based cooperation*, or the *tit-for-tat* rules. However, there is strong evidence for a simple cooperation rule in which reinforcements to existing propensities depend on the payoffs. This takes the form of playing the symmetric payoff dominant action and an initial propensity to cooperate.

15.8 Multiple games, complexity, and categorization

In most of the learning models considered so far, the object of interest was learning in a single game. Grimm and Mengel (2012) explore the possibility that subjects learn while playing multiple games. In different treatments, they control for the amount of information that subjects have access to. The main idea is that when players play a large number of games under limited information, they economize on scarce cognitive resources by *categorizing* the available games. Three kinds of behaviors under categorization are explored, as we now explain.

Suppose that we have six games (as in the experiments reported in this section), indexed by $\gamma = \{1, \dots, 6\}$ and game γ occurs with frequency $f(\gamma)$. In any game γ , the payoff matrix is $\Pi(\gamma)$, the generic action of player i is $a_i(\gamma)$, the mixed strategy profile of all other players is $\sigma_{-i}(\gamma)$ and player i 's beliefs about $\sigma_{-i}(\gamma)$ are denoted by $\mu_i(\gamma)$.

Suppose that the decision maker forms a *partition* of the games and $G_j \subseteq \{1, \dots, 6\}$ is a *category* of games. Here we assume that the partition is exogenous, but there have been attempts at forming endogenous partitions; see, for instance, Samuelson (2001), Fryer and Jackson (2008), and Mengel (2012). Define the *average game matrix* for category G_j to be

$$\bar{G}_j = \sum_{\gamma \in G_j} \left(\frac{f(\gamma)}{\sum_{\gamma \in G_j} f(\gamma)} \right) \Pi(\gamma).$$

Also, define the average behavior of players other than player i in category G_j by

$$\bar{\sigma}_{-i}(G_j) = \sum_{\gamma \in G_j} \left(\frac{f(\gamma)}{\sum_{\gamma \in G_j} f(\gamma)} \right) \sigma_{-i}(\gamma).$$

Denote the beliefs of player i about $\bar{\sigma}_{-i}(G_j)$ by $\bar{\mu}_i(G_j)$.

We can now define the various modes of optimization in categorized games. Let $BR_i(\gamma, \mu_i(\gamma))$ be the best response of player i in game γ when his beliefs about the strategies of others are given by $\mu_i(\gamma)$.

1. *Best response bundling* (BRB). In this case, for any category of games, G_j , players have distinct beliefs, $\mu_i(\gamma)$, about the play of others for each game γ in G_j . However, they act as if they are facing the average game matrix $\bar{G}_j$. Thus, their chosen action, $a_i(\gamma)$, satisfies

$$a_i(\gamma) \in BR_i(\bar{G}_j, \mu_i(\gamma)), \gamma \in G_j.$$

2. *Belief bundling* (BB): In this case, behavior is similar to the case of *analogy-based expectations equilibrium* that we considered in Part 4 of the book. So players apply a common set of beliefs, $\bar{\mu}_i(G_j)$, to all games in category G_j , but they treat the payoffs from each game separately (unlike the case of BRB). Hence, their chosen action, $a_i(\gamma)$, satisfies

$$a_i(\gamma) \in BR_i(\gamma, \bar{\mu}_i(G_j)), \gamma \in G_j.$$

3. *Action bundling* (AB): In this case, players apply the average payoffs to any game in category G_j (as in BRB) and also use the average beliefs about the strategies of others in category G_j (as in BB). Hence, their chosen action, $a_i(\gamma)$, satisfies

$$a_i(\gamma) \in BR_i(\bar{G}_j, \bar{\mu}_i(G_j)), \gamma \in G_j.$$

If the category G_j is a singleton, $G_j = \{\gamma\}$, then action bundling is equivalent to $a_i(\gamma) \in BR_i(\gamma, \mu_i(\gamma))$, as in the neoclassical case.

The six games used in the experiments are shown in Figure 15.17. Games 4, 5, and 6 are derived, respectively, from games 1, 2, and 3 by adding 5 or 10 to the payoffs. All games have a strict Nash equilibrium with corresponding payoffs shown by bold numbers; this also corresponds to the maximin prediction. In games 2 and 5, the equilibrium is in strictly dominated strategies, which is quite easy to spot. In games 3 and 6, iterated deletion of strictly dominated strategies leads to the unique Nash equilibrium. Finally, in games 1 and 4, the Nash equilibrium can be achieved by deletion of weakly dominated strategies. In each treatment (see below), the games to be played in each round are picked randomly using a uniform distribution. Subjects are initially assigned to the roles of row and column players and their role is unchanged throughout. In each round, subjects are randomly rematched with others in a group of eight players over 100 rounds.

In treatments that have the letter F (few games) subjects play only games 1 and 2, while in treatments that have the letter M (many games) subjects played all six games. The treatments are

Payoff matrices, payoffs associated with the unique strict Nash equilibrium are in bold.

Game 1	<i>a</i>	<i>b</i>	<i>c</i>	Game 4	<i>a</i>	<i>b</i>	<i>c</i>
A	20, 20	15, 10	15, 10	A	25, 25	20, 15	20, 15
B	10, 15	25, 10	0, 10	B	15, 20	30, 15	5, 15
C	10, 15	15, 35	35, 0	C	15, 20	20, 40	40, 5
Game 2	<i>a</i>	<i>b</i>	<i>c</i>	Game 5	<i>a</i>	<i>b</i>	<i>c</i>
A	5, 5	15, 20	5, 10	A	15, 15	25, 30	15, 20
B	10, 15	5, 25	10, 10	B	20, 25	15, 35	20, 20
C	20, 5	25, 15	15, 10	C	30, 15	35, 25	25, 20
Game 3	<i>a</i>	<i>b</i>	<i>c</i>	Game 6	<i>a</i>	<i>b</i>	<i>c</i>
A	15, 10	20, 20	15, 15	A	20, 15	25, 25	20, 20
B	15, 20	10, 15	5, 10	B	20, 25	15, 20	10, 15
C	20, 5	15, 35	35, 0	C	25, 15	20, 40	40, 5

Figure 15.17 Six normal form games used in the experiments of Grimm and Mengel (2012).

Source: Reprinted from *Journal of Economic Theory* 147: 2220–59. Grimm, V. and Mengel, F. "An experiment on learning in multiple game environments." Copyright © 2012, with permission from Elsevier.

as follows. In the *highly accessible treatments*, FI and MI, payoff matrices are displayed prominently on the decision screen in front of the subjects, and the play of the opponent over the last five rounds is revealed. At the other extreme are the two *low accessibility treatments*, F and M, in which participants observe the payoff matrices for a limited time and do not observe the past play of opponents. Finally, there were two *intermediate accessibility treatments*, M+G and M+F. In treatment M+G, the payoff information is displayed prominently on the decision screen (as in MI), but the average action choices of the opponents are not revealed. Finally, in Treatment M+F, average past play of the opponent is revealed, but payoff information is not.

Two further treatments are run to check for learning spillovers; these are M* and M**. In treatment M*, in games 1 and 4, the row player's payoffs from (A, a), (A, b) and (B, b) are reduced by 5; the Nash equilibrium is unchanged and there are no changes in any other game. In treatment M**, actions B and C in games 2 and 5 are relabeled, so the Nash equilibrium is (B, b) rather than (C, b); there are no changes in other games.

The results are as follows.

1. *Complexity and learning*: Does greater complexity of games hinder learning? Games 2 and 5 are the simplest because there is a Nash equilibrium in strictly dominated strategies. In these games, the overall share of Nash choices is more than 98% across all treatments. Games 3 and 6 are slightly more complex because the Nash equilibrium is in iterated deletion of strictly dominated strategies. In the highly accessible treatment, MI, players learn to play the Nash equilibrium in 94% cases in game 3 and 85% in game 6. However, at the other extreme, in the low accessibility treatment, M, only 50% of row players play their part of the Nash equilibrium despite the column players playing their part of the equilibrium. The results for the most difficult games, 1 and 4, in which Nash equilibrium is achieved by eliminating weakly dominated strategies, is shown in Table 15.8.

It is clear that complexity hinders convergence to a Nash equilibrium. In the easiest treatment, F, more than 95% of the row and column players learn to play the Nash equilibrium in game 1. This is not very surprising because this is a fairly low dimensional game and the

Table 15.8 Frequencies of various choices of actions in games 1 and 4. Data from the last 50 rounds. The Nash equilibrium (A, a) is shown in bold. RP denotes row player and CP denotes column player.

Game 1 (RP)	A	B	C	(CP)	a	b	c
FI	0.90	0.08	0.02	FI	0.92	0.07	0.00
F	0.96	0.04	0.00	F	0.95	0.03	0.02
MI	0.62	0.32	0.06	MI	0.74	0.26	0.00
M	0.37	0.50	0.13	M	0.45	0.55	0.00
Game 4 (RP)	A	B	C	(CP)	a	b	c
MI	0.76	0.20	0.03	MI	0.83	0.17	0.00
M	0.40	0.52	0.07	M	0.48	0.52	0.00

Source: Grimm and Mengel (2012).

number of learning rounds (100) is quite high. However, in the most difficult treatment, M, only 37% of row players and 45% of column players learn to play a Nash equilibrium. For game 4, the corresponding percentages are 40% and 48%.²⁵ The overall complexity of even the most complex games here is relatively low, and the number of rounds is high, hence, the results are sobering for classical game theory.

2. *Framing effects*: It is plausible to argue that the treatments M, M*, and M** are equally complex. Thus, if there is any difference in the behavior of subjects across these three treatments, then it is likely to be due to framing effects and possibly due to learning spillovers (which is the authors' preferred explanation).²⁶ In these treatments, games 3 and 6 are identical and the Nash equilibrium is (A, b). Yet, while the proportion of column players who play b is nearly identical in all treatments, a significantly higher proportion of row players play A in treatment M** relative to treatment M (71% as compared to 50% in game 3; and 66% as compared to 44% in game 6). There are also significantly less C choices in treatment M** relative to treatment M.
3. *Play in strategically equivalent games*: The following pairs of games are strategically equivalent: 1 and 4; 2 and 5; 3 and 6. Between pairs of these equivalent games, and for the last 50 periods, the difference in choice frequencies of all strategies are statistically insignificant. This result holds also for the most complex treatment, M. Thus, learning ensures that strategically equivalent games are played in the same manner.
4. *Source of complexity*: Consider the two pairs of strategically equivalent games, 1 and 4, and 3 and 6. In both games, the row player should play A in a Nash equilibrium. Figure 15.18 shows the predicted frequencies of row players who play A , based on logit regressions in the six-game treatments (M, MI, M+F, and M+G).

In all cases, learning drives choices towards all-play A , but with varying degrees of success. Complexity plays a role. The Nash choices are highest in the highly accessible

²⁵ Oechssler and Schipper (2003) show that subjects are not very good at guessing the payoff functions of the opponents when they do not have this information, yet they often learn to play the Nash equilibrium.

²⁶ Bednar et al. (2012) consider players who play two different games with different opponents. They find evidence of behavioral spillovers from one game to the other. Individuals develop common behaviors across games, which could be a form of categorization.




Figure 15.18 Predicted shares of A choices by row player, based on logit regressions, in treatments M, MI, M+F, and M+G.

Source: Reprinted from *Journal of Economic Theory* 147: 2220–59. Grimm, V. and Mengel, F. “An experiment on learning in multiple game environments.” Copyright © 2012, with permission from Elsevier.

treatment, MI, lowest in the least accessible treatment, M, and intermediate in the treatments, M+F and M+G. Among the two intermediate treatments, M+F and M+G, there are more Nash choices in the latter. These results suggest that displaying payoff information prominently is more conducive for Nash learning as compared to revealing information about average past play of opponents.

Behavior in the two treatments, M+F and M, is quite similar, which suggests that information on average past play of opponents is not critical to Nash learning. The behavior in treatments M+G and M differs. This is particularly true in games 3 and 6 where there are relatively more Nash choices in the treatment M+G. Thus, providing payoff information on the screens contributes to a reduction in the underlying complexity of the game. In a comparison of treatments M+G and MI there are more Nash choices in the latter, which again suggests a role of complexity (players have information about the opponent’s past choices in MI, but not in M+G). In a comparison of games 1 and 4, the most difficult games, the differences between the treatments M+G and M+F are more muted (only in game 4 does the difference achieve statistical significance at the 10% level).

5. *Modes of categorization:* Consider two exogenous, but plausible, partitions of the six games. The first partition separates the most difficult games, 1 and 4, in a separate category to give $\{\{1, 4\}, \{2, 3, 5, 6\}\}$. The second partition is the finest; it obtains when pairs of similar games

are grouped, i.e., $\{\{1, 4\}, \{2, 5\}, \{3, 6\}\}$. Consider action, or best response bundling in which there is a distribution of the population that holds either of the two given partitions.²⁷ The authors show that this explains well 90% of the behavior in treatment M, which, as we noted above, significantly departs from the predicted Nash behavior. This mode of categorization explains better than belief-based categorization.

6. *Influence of cognitive ability*: Subjects who score more highly in the *cognitive reflection test* choose the Nash equilibrium more often. Further, 85% of these high achievers exhibit identical choices in strategically equivalent games.

15.9 Possibility and impossibility of rational learning

Under reinforcement learning, players only observe the payoffs arising from their own past experience. By contrast, under fictitious play, players observe the past frequency of play of the opponent and attempt to infer the mixed strategies that the opponent is likely to play. In both cases, players behave as *automata* and do not evaluate the efficiency of their learning rule. In this section, we consider two alternative learning rules, *regret learning* and *hypothesis learning*, which allow players to periodically evaluate the efficiency of their learning strategy.

Other models of learning that we have considered above also incorporate some of these features. For instance, regret from forgone payoffs is incorporated in EWA and impulse matching learning (Sections 15.4, 15.6). Players also evaluate the efficiency of their learning strategies in models of rule learning (Section 15.7). Furthermore, these models are consistent with a range of empirical evidence.

Situations in which a single agent/animal interacts with nature, which was the original application of reinforcement learning in biology, are simpler relative to multi-player interactions. For instance, if players know that others are trying to learn something about their behavior, then they might alter their behavior. One must address similar feedback considerations in multi-agent learning situations. Learning might fail despite players being fully rational in the classical sense and despite there being infinite repetitions of the game. This “impossibility of learning” result was demonstrated by Foster and Young (2001).²⁸ Paradoxically, relaxing rationality can facilitate learning as we shall see below; see, for instance, the account in Young (2004) that we draw upon below.

We first illustrate the idea of *regret matching*. Consider the two-player coordination game in Figure 15.19 that has two pure strategy coordination Nash equilibria and one in mixed strategies. None of the players knows the actual game being played. But they can observe the history of play and the payoffs for each combination of actions played. Suppose that the row player observes the sequence of actions and own-payoffs shown in Table 15.9.

The undiscounted sum of payoffs to the row player, over the ten periods, is 3. Suppose that players are interested in minimizing their *ex-post regret*, which can be computed as follows. Fix the column player’s period by period strategy over the ten periods as in Table 15.9. Now give the row player an opportunity to play a fixed strategy, either s_1 or s_2 , in each of the ten periods. If the row player had chosen s_1 (respectively, s_2) in every period, keeping fixed the strategy of the

²⁷ Under the partition $\{\{1, 4\}, \{2, 5\}, \{3, 6\}\}$, under BRB, in response to the observed frequencies of play of the column players, row players should choose A or B in category $\{1, 4\}$, C in category $\{2, 5\}$ and A in category $\{3, 6\}$. The best responses under the partition $\{\{1, 4\}, \{2, 3, 5, 6\}\}$ depend on the frequencies of the various games.

²⁸ See also Binmore (1987) and Jordan (1992, 1993) for results of a similar nature.

	s_1	s_2
s_1	1, 1	0, 0
s_2	0, 0	1, 1

Figure 15.19 A 2×2 coordination game.

Source: *Strategic Learning and Its Limits* by H. Peyton Young (2014): table on p. 6 (The Soda Game) and table on p. 7. By permission of Oxford University Press.

Table 15.9 Sequence of actions and own-payoffs in a two-player coordination game.

Actions	s_1, s_2	s_2, s_1	s_1, s_2	s_1, s_1	s_1, s_2	s_2, s_2	s_1, s_2	s_1, s_2	s_2, s_1	s_2, s_1
Payoff	0	0	0	1	0	1	0	0	0	1

Source: Young, H. P. (2004). *Strategic Learning and Its Limits*. Oxford: Oxford University Press.

column player, he would have received a payoff of 4 (respectively, 6). Thus, the player's forgone payoffs from not playing s_1 and s_2 , respectively, throughout, are $4 - 3 = 1$ and $6 - 3 = 3$.

Regret is defined as the average forgone payoff over a given time period. Hence, regret from not playing s_1 throughout is $\frac{1}{10} = 0.1$, while the regret from not playing s_2 throughout is $\frac{3}{10} = 0.3$. While the past cannot be undone, the idea behind *regret learning* is that since not playing s_2 has been relatively more regrettable, in future the player should play s_2 more often.

The idea of a state of play which leaves *no regret* in the mind of the player goes back to Hannan (1957).²⁹ Consider a single player playing a game against nature. The player can choose actions from the finite set $A = (a_1, a_2, \dots, a_n)$. Nature has the finite action set $B = (b_1, b_2, \dots, b_m)$. The payoff function of the player is given by $\pi : A \times B \rightarrow \mathbb{R}$. The action profile at time $\tau = 1, 2, \dots$ is denoted by the tuple (a^τ, b^τ) . Hence, at any time T , the history of play is given by

$$h^T = [(a^1, b^1), (a^2, b^2), \dots, (a^T, b^T)]. \quad (15.74)$$

The regret from not playing action $a_i \in A$ at time T , denoted by r_i^T , is defined as

$$r_i^T = \bar{u}_i^T - \bar{u}^T, \quad (15.75)$$

where $\bar{u}^T = \frac{1}{T} \sum_{\tau=1}^T \pi(a^\tau, b^\tau)$ is the average payoff to the player at time T , and $\bar{u}_i^T = \frac{1}{T} \sum_{\tau=1}^T \pi(a_i, b^\tau)$ is the hypothetical average payoff to the player, had he played the action a_i in all past periods. Denote by $r^T = (r_1^T, r_2^T, \dots, r_n^T)$, the regret vector that summarizes the regret from each $a_i \in A$.

Definition 15.21 (No-Regret over histories) A history of play h^T has no-regret if

$$\text{as } T \rightarrow \infty, \max_{a_i \in A} r_i^T \leq 0. \quad (15.76)$$

²⁹ See also Foster and Vohra (1993) and Fudenberg and Levine (1995).

Suppose that we have a behavioral rule for the player $\phi : h^T \rightarrow \Delta A$, which describes a probability distribution over the actions of the player at time $T + 1$, given the history of play upto that time, h^T .

Definition 15.22 (No-Regret over behavioral rules) *A behavioral rule for the player, ϕ , satisfies “no regret” if almost surely, all realizations of histories generated by ϕ , satisfy (15.76).*

The next definition formalizes the basic idea in *regret matching*, namely, that in each period, players play their pure strategies with probabilities that equal the relative regret from the strategies. So strategies associated with greater regret are played with higher probabilities.

Definition 15.23 (Regret Matching (RM); Hart and Mas-Colell, 2000) *Suppose that an individual maps his vector of regrets r^T , available at the end of period T , into a vector of probabilities of play next period, $p^{T+1} : r^T \rightarrow [0, 1]^n$, by using the following rule*

$$p_i^{T+1} = \frac{\max\{r_i^T, 0\}}{\sum_i \max\{r_i^T, 0\}}, \quad (15.77)$$

where p_i^{T+1} is the probability of playing action $a_i \in A$ at time $T+1$. When $\sum_i \max\{r_i^T, 0\} = 0$, then we can choose p_i^{T+1} to be any number. If an individual plays his pure strategies with probabilities that satisfy (15.77), then he is said to be following regret matching (RM).

Theorem 15.1 (Hart and Mas-Colell, 2000) *Consider a finite normal form game in which all players use regret matching. Then, almost surely, in the limiting case, RM is a behavioral strategy that has no regret against any sequence of play by the opponent.*

From (15.75), RM can be viewed as a form of *aspiration level learning* where the aspiration level is given by the historical average payoff.

Foster and Vohra (1993) propose *regret matching with ε -experimentation*, in which a player only needs to know the actual payoffs in the past, but need not know the strategies of the other player. The main difference from RM is that with some small probability ε , players make mistakes, or trembles, or they simply experiment and choose randomly among all their actions with equal probability. With probability $1 - \varepsilon$, players use RM. It can then be shown that for small enough noise, ε , *regret matching with ε -experimentation*, leads to extremely small regret for all histories of play. Foster and Vohra (1998) allow for a time dependent noise, i.e., $\varepsilon(t)$. They show that if $\varepsilon(t) \rightarrow 0$ at a sufficiently slow rate, then there is no regret in the long run.

The central result with RM is that the frequency of joint play is related to the set of *coarse correlated equilibria*.

Theorem 15.2 (Hart and Mas-Colell, 2000) *Suppose that all players use RM. Then the empirical frequency of joint play of actions of all players converges (almost surely) to the set of coarse correlated equilibria.*

A encouraging feature of these results is that rules such as RM are computationally simple and require no more computations than are needed, say, for reinforcement learning. However, these rules do well only in the very long run of play. There is no guarantee that the outcome of such rules over short fixed periods of play converges to the coarse correlated equilibrium, or to anything at all.

A natural question in the context of learning models is: do players learn, and what does it mean to learn? Consider a two-player game. Suppose that players are *Bayesian rational* in the following sense. (1) They form posterior Bayesian beliefs about the likelihood of the future play of their

opponent, conditional on the history of play so far. (2) At each stage, players choose strategies to maximize their intertemporal payoffs, given their forecasts of what the other player is expected to do.

Suppose that the history of play at time t is given by h^t (see (15.74)). Conditional on the history, the beliefs of a player about the other player's behavior for the next period are given by the probability distribution p_{-i}^{t+1} . Suppose that the actual behavior of the opponent in period $t + 1$ is given by q_{-i}^{t+1} .

Definition 15.24 (*Predictive error*) The predictive error is given by the distance between p_{-i}^{t+1} and q_{-i}^{t+1} , i.e., $d(p_{-i}^{t+1}, q_{-i}^{t+1})$, where d is the Euclidean metric.

Definition 15.25 (*Learning to predict in the strong sense*) A player is said to learn to predict in the strong sense if, almost surely, $\lim_{t \rightarrow \infty} d(p_{-i}^{t+1}, q_{-i}^{t+1}) = 0$.

Definition 15.26 (*Learning to predict in the weak sense*) A player is said to learn to predict in the weak sense if, almost surely, the mean square error goes to zero, i.e., $\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{s \leq t} d(p_{-i}^{s+1}, q_{-i}^{s+1})^2 = 0$.

Blackwell and Dubins (1962) gave the absolute continuity condition. Roughly speaking, players always put positive probability on histories that could have arisen under their joint play so far.³⁰ Kalai and Lehrer (1993) showed that under Bayesian rationality and absolute continuity, players will learn to predict in the strong sense and we get convergence to a Nash equilibrium.

Suppose that each player does not know the payoffs of the other player. This makes the forecasting of the future play of the opponent difficult. Foster and Young (2001) show that if there is enough ex-ante uncertainty about the payoffs of the opponent, then it is possible to get the following two results. (1) No learning takes place, even in a weak sense, and even in the long run. (2) Play does not converge to a Nash equilibrium. The contrast between these results and the Kalai–Lehrer results is what Young (2004) refers to as the *possible* and the *impossible* in multi-agent learning. In his words: “A consequence of this result is that there exist no general, model-based procedures for multi-agent learning when players are perfectly rational and they have sufficiently incomplete knowledge of their opponents’ payoff functions.”

Hart and Mas-Colell (2006) also give a negative result. They focus on learning rules that are *uncoupled*. An uncoupled learning rule does not depend on the payoffs of the opponent. Reinforcement learning and regret matching (RM) are examples of such a rule. They show that under certain conditions no uncoupled learning rules exist such that the period by period behavior of players converges almost surely to a Nash equilibrium of the game (or even to a ε -equilibrium of the game). These conditions include the finiteness of action spaces, and a state variable that uses the last τ periods of play. RM in which regret is computed from the last τ periods of play is an example.

We now provide a brief discussion of a learning rule, *hypothesis testing*, that is also uncoupled, but it does not satisfy all the conditions of the Hart and Mas-Colell (2006) result. This has desirable long-run properties in that it is able to achieve an ε -equilibrium of the game. Consider the following version of hypothesis testing in Foster and Young (2003) that draws on the standard theory of hypothesis testing in statistics.³¹ Consider two players A and B whose finite pure strategy sets are S_A and S_B respectively, with generic elements a_i and b_i . The model is not restricted

³⁰ For the formal definition of the absolute continuity condition, see Fudenberg and Levine (1998) and Young (2004).

³¹ Also see Arthur (1994) and Jehiel (1995, 1998) for earlier treatments.

to two players; we make this restriction purely for pedagogical convenience. Probability distributions over these sets are given by Σ_A and Σ_B , respectively; the respective typical elements of these sets are given by p_A and p_B .

Suppose that we consider everything from player A's point of view (the analogous arguments for player B are obvious). Then "a model" for player A at some time t is a probability distribution $p_B^t \in \Sigma_B$, so the space of possible models is given simply by the set Σ_B . At time t , player A believes that player B will play p_B^t from time t onwards. Denote the payoff to player A from playing his action a_i every period from period t onwards by $u_{At}(a_i, p_B^t)$. Player A then plays each of the pure strategies in S_A , according to a *quantal response function*:

$$q_A^t(a_i) = \frac{\exp u_{At}(a_i, p_B^t)}{\sum_j \exp u_{At}(a_j, p_B^t)}. \quad (15.78)$$

The steps involved in hypothesis testing are as follows:

Step I. (Model formulation and action choice) At time t , players A and B pick models of their opponent's play, respectively, $p_B^t \in \Sigma_B$ and $p_A^t \in \Sigma_A$. Each player then plays the respective mixed strategy q_A^t, q_B^t over his actions, which is given by the quantal response function. This feature of the model is similar to smooth fictitious play. However, while the model is continuously updated in smooth fictitious play, here it is only updated if it fails; see below.

Step II. (Data collection phase) If a player is not collecting data, then he initiates data collection from the next period with a probability $\frac{1}{\tau}$, $\tau > 0$. Should a data collection phase begin, it lasts over τ periods.

Step III. (Hypothesis testing) Once data is collected, then over the τ periods of the sample, the player computes the empirical frequency of play of the opponent $\bar{p}_{-i}$. Players then specify a rejection region, choosing to reject their model in Step I if $d(\bar{p}_{-i}, p_{-i}^t) > \gamma_i$ (where $d(\bar{p}_{-i}, p_{-i}^t)$ is the Euclidean distance between $\bar{p}_{-i}$ and p_{-i}^t). If $d(\bar{p}_{-i}, p_{-i}^t) \leq \gamma_i$, the model is accepted. If a model is rejected, then with some probability ρ , the same model is retained, but with probability $1 - \rho$, a new model is chosen from the feasible set.

Foster and Young (2003) show that under certain plausible conditions, hypothesis testing leads to an ε -equilibrium of the game. These conditions are related to the sizes of the sample, the test tolerance, and a normalization of the utility function; for details, see Foster and Young (2003). While this result is similar to Kalai and Lehrer (1993), however, in this case the rationality of players is limited. They do not, for instance, need to know the payoffs of other players. Nor is information always used optimally: Players can get by without sampling at all for several periods and when they sample, they hold fixed their model of the other player. Also players have no repeated game strategies. Hence, paradoxically, relaxing rationality induces a good limiting outcome for the players.

CHAPTER 16

Stochastic Social Dynamics

16.1 Introduction

In the study of learning under deterministic dynamics (Chapters 14, 15), we have used the concept of an *asymptotically stable steady state*, which has two important features. First, when there are multiple steady states, the outcome is history-dependent, i.e., it depends on initial conditions. Second, shocks/mutations are rare events, so once the system reaches a steady state, it stays there forever.

In many cases, however, there could be multiple sources of *persistent randomness*, such as continuous mutation, taste shocks, and play of mixed strategies. Asymptotic stability is no longer the appropriate notion of stability in this case. Starting from an asymptotically stable state, persistent randomness can gradually, and over time, tip the system out of the basin of attraction of one asymptotically stable state, into the basin of attraction of another. The appropriate modification to asymptotic stability, in this case, is to consider the long-run likelihood of the system remaining in various possible steady states.

If, in the long run, the system is likely to stay in the neighborhood of a particular steady state, then it is called a *stochastically stable state* (Foster and Young, 1990). If such a state is unique, then the initial state of the system does not matter. Such a state is said to be *ergodic*. This framework can be used to understand the evolution of institutions and norms as the result of gradual, small, adjustments arising from adaptively learning individuals.

Section 16.2 introduces the basic concepts and considers two introductory applications (Kandori et al., 1993; Young, 1993a). In order to derive the outcome under stochastic dynamics, we introduce some noise into the system. Any player when faced with an opportunity to revise his strategy, chooses the best response with high probability and, with some small complementary probability, mutates by randomizing among all his choices.

In Section 16.2.1, we consider a 2×2 coordination game that has two Nash equilibria; both are asymptotically stable. Under stochastic dynamics, the multiplicity of the steady states is narrowed down to a unique stochastically stable state that is also the risk dominant equilibrium. The risk dominance property of the stochastically stable states need not hold for higher dimensional games.

Section 16.2.2 uses the method of directed graphs to derive the stochastically stable steady states in a model of technology adoption for a 3×3 game. In this model, strategies correspond to a choice of technology and the stochastically stable state corresponds to the coordinated outcome with the lowest joint payoffs for both players.

Section 16.3 begins our formal study of *perturbed Markov dynamics*. Players have a memory of the play in the last T time periods. They take a sample of size $\tau \leq T$ from memory, and play

a best response to the sampled frequencies of the play of opponents, with some high probability. However, with a low complementary probability, ε , they randomly choose among their pure strategies. The set of states in this model are all possible T period histories of the strategies played by the players.

This structure gives rise to a Markov process on the set of states, denoted by $P^{T,\tau,\varepsilon}$. We may also refer to this as adaptive play with memory length T , sample size τ , and error rate ε . The associated transition matrix on the set of states is given by P^ε . It gives the probability of transiting from any state to any other possible state. We begin by stating a fundamental result that helps us determine the unique stationary distribution as well as the stochastically stable states (Young, 1993a). We then consider some applications to 2×2 games.

Section 16.3.1 gives an application to symmetric 2×2 games in which we allow for inertia, and replace best response by *imitate the best* (Alós-Ferrer, 2000). Under imitate the best, the stochastically stable state may allow for the play of strictly dominated strategies, although this result does not arise when players play best replies. Section 16.3.2 shows that in 2×2 games, the stochastically stable state is also the risk dominant state (Kandori et al., 1993; Young, 1993a). Finally, in Section 16.3.2, we consider several extensions of the basic model, such as players who use different sample sizes, heterogeneity among players, and alternative noise structures (Young, 1998).

Section 16.4 considers results from 3×3 games (Young, 1993a, 1998). Many of the results of 2×2 coordination games do not generalize to these games. Not only might there be no strict risk dominant equilibrium for 3×3 games, but even if it exists, there is no guarantee that it coincides with stochastic stability. However, for a class of games, acyclic games, several of the results generalize. For instance, in this class of games, the unperturbed dynamics converge to one of the conventions, a Nash equilibrium, irrespective of the initial state.

Section 16.4.1 notes that the low rationality adaptive play environment of perturbed dynamics often generates a Nash equilibrium. It does so without requiring the stringent requirements of the epistemic conditions for a Nash equilibrium. However, in Section 16.4.2, we show that stochastic dynamics might pick an inefficient norm. In particular, stochastic dynamics will have a tendency of picking a convention that is distributionally more fair relative to conventions that are associated with extreme payoff inequities.

Section 16.4.3 briefly considers applications of stochastic dynamics to public goods games (Myatt and Wallace, 2005, 2008). We consider public goods games that exhibit threshold externalities, i.e., they require a certain minimum number of players to contribute, otherwise no public good materializes. The method of analysis is again to calculate the least cost of moving between the basins of attractions of two equilibria. In one equilibrium, the public good materializes, and in another, nobody contributes. While the applications are novel and interesting, they leave open the question of how one might use them to explain the results of experiments on public goods games with and without punishment (see Part 2 of the book).

Section 16.5 considers games played by boundedly rational players whose direct interaction is local, but indirect interaction is global (Blume, 1993, 1995). These are games played on graphs in which each individual is located at the vertex of a tree and is influenced by the behavior of the immediate neighbors. We characterize the features of long-run frequency of play and give an application to agrarian contracts (Young and Burke, 2001).

16.2 Fixing ideas in stochastic dynamics

In deterministic situations, random shocks are treated as one-off events, isolated and rare. Hence, once the deterministic system is in some steady state, the concept of *asymptotic stability* requires

that if a rare, but small, deviation occurs, the system reverts back to the asymptotically stable state. However, in many strategic situations, several possible sources of *persistent randomness* exist. For instance, players might play mixed strategies, players might be randomly matched with each other, random mutations might alter the types of players, payoffs might be subjected to random shocks, and there might be immigration, or emigration of players within distinct populations.

The behavior of a system subjected to persistent random shocks, possibly from several sources, is fundamentally different from deterministic dynamics. The long-run behavior of such a system cannot simply be understood by taking expectations of the relevant random variables. Under persistent randomness, the system never really settles down. It is ever-changing, and in a constant state of flux. These persistent shocks, even if they arise with a small probability, can accumulate gradually over time and force the system out of the basin of attraction of an asymptotically stable steady state. The focus under persistent randomness, then, is on the long-run frequency with which various strategies are played. It is possible, of course, as in other dynamic models, that the long run might entail an interval of time that is unrealistically long. However, when players interact locally, or in small geographic neighborhoods, the speed of convergence may be fast (Ellison, 1993; Young, 1998). Moreover, when the noise level is small, but non-vanishing, the process can reach a neighborhood of the stochastically stable equilibrium quite quickly (Kreindler and Young, 2013, 2014).

The basic ingredients of this approach are as follows. First, players have *limited information*, or *limited cognitive ability*, thus, their memory goes back only a finite number of past periods. Second, players are *boundedly rational* in the sense that although they learn from the limited history of play by the opponent, they play *myopic best replies* (as under fictitious play). Third, players play their myopic best replies with a probability $1 - \varepsilon$, where ε is small. With the complementary probability ε , players *mutate*, or make *mistakes* by randomly choosing any pure strategy in their strategy set, S . *Mutation* is taken here as a reduced form for the unmodeled behavioral biases of players. Relative to the classic case of full Bayesian rationality, this is a relatively low rationality framework.

In the absence of mutations, we get deterministic dynamics and the system might settle down to some long-run asymptotically stable state. When there are multiple asymptotically stable states, then history plays a powerful role in choosing the eventual outcome. However, with persistent noise created by ε , however small, the system might eventually *tip out* of the basin of attraction of an asymptotically stable state, and over time, enter into the basin of attraction of another asymptotically stable state. This gives rise to a *punctuated equilibrium*.

In a punctuated equilibrium, a long period of stay in one asymptotically stable state is followed by a tipping into another asymptotically stable state. In this case, we would like to know which of the asymptotically stable states have a greater likelihood of being observed in the long run, as the noise $\varepsilon \rightarrow 0$. We describe such states as *stochastically stable states*. Certain stochastically stable states might eventually be reached with great likelihood, *irrespective of the initial conditions*; these states are said to be *ergodic*. The basic machinery of analysis is described in Appendix C.

How does this framework help us to understand the evolution of institutions? One theory is that institutions arise through the actions of important, or pivotal individuals, or events. For instance, the importance of Napoleon to the legal code, Bismark to social security, and Gandhi to non-violence. However, the thrust of stochastic dynamics is that small individuals, who use adaptive learning, and their small mutations play a key role. Whilst individually small, such mutations can accumulate and tip an equilibrium from one state of the system to another and one of the states might be stochastically stable, or even ergodic.

	s_1	s_2
s_1	1, 1	0, 0
s_2	0, 0	2, 2

Figure 16.1 A 2×2 coordination game.

Source: Wallace and Young (2014).




Figure 16.2 Simulated plot of the transitions between the two states. Fraction of players who play s_1 on the vertical axis and time (step size 0.05) on horizontal axis.

Source Reprinted from *Handbook of Game Theory with Economic Applications*, Young, P. and Zamir, S. (eds.), "Stochastic evolutionary game dynamics," Wallace, C. and Young, P. Volume 4, pages 327–80. Copyright © 2015, with permission from Elsevier.

Wallace and Young (2014) illustrate the basic method of analysis for the simple 2×2 coordination game shown in Figure 16.1. The two *evolutionary stable strategies* (see Definition 14.4) are the two monomorphic equilibria in which all play s_1 , or all play s_2 .

Let x_t be the proportion of players in the single population who play the strategy s_1 at time t . Starting from the monomorphic population in which all play s_1 , suppose that there is an invasion by a small fraction $1 - x_t$ of mutants who play s_2 . Let $\bar{\sigma}_t$ be the average strategy in the population at time t , given by $\bar{\sigma}_t = x_t s_1 + (1 - x_t) s_2$. Then the replicator dynamics in (14.25) are given by $\dot{x}_t = [\pi(s_1, \bar{\sigma}_t) - \pi(\bar{\sigma}_t, \bar{\sigma}_t)] x_t$, where $\pi(s_1, \bar{\sigma}_t) = x_t$ and $\pi(\bar{\sigma}_t, \bar{\sigma}_t) = x_t^2 + 2(1 - x_t)^2$. Under stochastic dynamics, we can write the replicator dynamics as a stochastic differential equation

$$\dot{x}_t = [x_t - x_t^2 - 2(1 - x_t)^2] x_t + \varepsilon dW_t, \quad (16.1)$$

where W is a Wiener process such that $W_t - W_0 \sim N(0, dt)$.

Figure 16.2 shows a simulated plot of (16.1) for a relatively high value of $\varepsilon = 0.6$ and $x_0 = 0.5$. It is clear that the system spends most of its time in the monomorphic state s_2 , but very little in the monomorphic state s_1 . This becomes even more stark as the noise term ε shrinks further. In

	s_1	s_2
s_1	9, 9	0, 7
s_2	7, 0	8, 8

Figure 16.3 A symmetric two-player coordination game.

the limit as $\varepsilon \rightarrow 0$, the system spends all its time in the neighborhood of the monomorphic state s_2 , irrespective of the history of play. Thus, the monomorphic state s_2 is the stochastically stable state. In the rest of the chapter, we outline various applications of this basic idea for discrete time models, based on the frameworks developed simultaneously by Kandori et al. (1993) and Young (1993a).

16.2.1 Stochastic dynamics in a coordination game

Let us fix ideas on the difference between deterministic and stochastic dynamics in a concrete example. Consider the symmetric two-player coordination game in Figure 16.3. There are two Nash equilibria in the game, (s_1, s_1) and (s_2, s_2) .

Consider the deterministic dynamics of the game using replicator dynamics as an illustration. Thus, assume, as in evolutionary game theory, that individuals are *hardwired* to play either s_1 or s_2 . Suppose that initially the entire population plays s_1 , but a small fraction x of the population mutates and begins to play s_2 . Individuals in the population are randomly paired to play the coordination game. Thus, the average strategy in the population is given by $\bar{\sigma} = xs_2 + (1-x)s_1$.

If the population is large, it is *as if* each individual meets an opponent who plays the mixed strategy $\bar{\sigma}$. Thus, the payoff of individuals from playing, respectively, the strategies s_1 and s_2 , is given by, $\pi(s_1, \bar{\sigma}) = 9(1-x)$ and $\pi(s_2, \bar{\sigma}) = 7+x$. The payoff of the average strategy against itself is given by $\pi(\bar{\sigma}, \bar{\sigma}) = x\pi(s_2, \bar{\sigma}) + (1-x)\pi(s_1, \bar{\sigma})$, or $\pi(\bar{\sigma}, \bar{\sigma}) = 10x^2 - 11x + 9$. The replicator dynamics are described by $\dot{x} = x[\pi(s_2, \bar{\sigma}) - \pi(\bar{\sigma}, \bar{\sigma})]$, or

$$\dot{x} = f(x) = -2x(5x^2 - 6x + 1). \quad (16.2)$$

The steady state is found by solving $\dot{x} = 0$. There are three solutions: $x^* = 0$, $x^* = 1$, and $x^* = 1/5$. The solutions $x^* = 0$ and $x^* = 1$ correspond to a monomorphic equilibrium in which everyone plays, respectively, s_1 and s_2 . The solution $x^* = 1/5$ corresponds to a polymorphic equilibrium in which the strategies s_1 and s_2 are played by, respectively, 80% and 20% of the population.

To check for asymptotic stability, we evaluate $f'(x) = -30x^2 + 24x - 2$ at the three steady states to get: $f'(0) = -2 < 0$, $f'(1) = -8 < 0$, and $f'(1/5) = 1.6 > 0$. Hence, there are two asymptotically stable steady states, $x^* = 0$, $x^* = 1$, while $x^* = 1/5$ is asymptotically unstable. Once the deterministic system achieves a state where all play s_2 , or all play s_1 , then the system never escapes that state. In the terminology of Appendix C, the two *recurrent classes* are composed of the *absorbing states* s_1 and s_2 .

Let us now consider the outcome under stochastic dynamics. Suppose that in a large population, players are randomly matched with each other to play the coordination game in Figure 16.3 across many periods. In any period, only one randomly selected player, the *experimenting player*, gets an opportunity to choose his strategy in the coordination game. Before the experimenting

player chooses his strategy, he takes a random sample of size m of the population and receives an honest report of the strategies that each sampled player intends to play. Suppose that it is found that y players intend to play s_1 , while the remaining $m - y$ players intend to play s_2 .

The experimenting player then infers that a fraction $\frac{y}{m}$ of the population plays s_1 and the remaining, $1 - \frac{y}{m}$, plays s_2 . The experimenting player then plays a best response to these frequencies with some high probability $1 - \varepsilon$, where $\varepsilon > 0$ is small. However, with probability ε , the player chooses, for whatever reason (lack of deliberation, bounded rationality, mutation), to randomize equally between the two actions, s_1 and s_2 . We no longer have a deterministic process. Instead, we have a *perturbed process with noise* ε . In any period, behavior will be unpredictable. However, we can find the stationary distribution of play (see Appendix C), as follows.

1. Identify the *recurrence classes*, C_1 and C_2 , of the unperturbed system in which players play their best response with probability 1. In this example, the recurrence classes are the singleton sets $C_1 = \{s_1\}$, $C_2 = \{s_2\}$. Singleton recurrence classes are also known as the *absorbing states*.
2. A *tree* that is rooted at a node belonging to the j th recurrence class is known as a j -tree. The second step is to find all j -trees. There are two trees in this example; a 1-*tree* that is rooted at s_1 and a 2-*tree* that is rooted at s_2 ; see Figure 16.4 where the root is indicated with an arrow (s_1 in the left panel and s_2 in the right panel).
3. Suppose that we start with a 1-tree and its terminal node at which all individuals play s_2 ; see left panel in Figure 16.4. We now want to find the *directed path of least resistance* from s_2 to s_1 . This is called a 2,1 path and the resistance along this path is denoted by ρ_{21} . We then repeat an identical calculation for a 2-tree, rooted at s_2 . In this case, starting with the terminal node at s_1 , we are interested in calculating the least resistance, ρ_{12} along the directed path 1,2.

To calculate ρ_{21} , suppose that initially all players play s_2 . The best response is also to play s_2 ; see Figure 16.3 in which (s_2, s_2) is a Nash equilibrium. How can play move in the direction of s_1 ? This occurs when an experimenting player who has a chance to choose his strategy, mutates with probability ε and randomizes equally between s_1 and s_2 . Hence, s_1 is played with a probability $\frac{\varepsilon}{2}$. In particular, we want to find the path of least resistance from s_2 to s_1 . Hence, we wish to consider a series of successive mutations in the direction of s_1 , each of which happens with probability $\frac{\varepsilon}{2}$ in each period. We are interested in having a certain critical minimum percentage of players who play s_1 such that, thereafter, no further mutations are needed and the best response of each player is to play s_1 . Should this happen, the system naturally proceeds towards s_1 without any further mutations. Such a point is sometimes known as a *tipping point*.

What fraction of the players need to play s_1 for the tipping point to be achieved? Strategy s_1 becomes a best response by an experimenting player if

$$9\frac{y}{m} \geq 7\frac{y}{m} + 8(1 - \frac{y}{m}). \quad (16.3)$$

Solving for the cutoff value of $\frac{y}{m}$ from (16.3) we get $(\frac{y}{m})^* = \frac{4}{5}$, so 80% of the players must play s_1 for s_1 to be a best response. The number of such players in a sample of size m is $y^* = \frac{4}{5}m$. Starting from a situation where all play s_2 we require y^* successive mutations in the direction of s_1 to arrive at a tipping point; this gives the path of least resistance. Each mutation is independent of the other and has an identical probability of $\frac{\varepsilon}{2}$. Thus, $y^* = \frac{4}{5}m$ successive mutations have a probability of $(\frac{\varepsilon}{2})^{\frac{4}{5}m}$, hence, $\rho_{21} = \frac{4}{5}m$. Proceeding in an analogous manner, in the case of the




Figure 16.4 A 1-tree and 2-tree for the coordination game.

2-tree that is rooted at s_2 and starting from its terminal node s_1 we get that $\rho_{12} = 1 - \gamma^* = \frac{1}{5}m$. This situation is depicted in Figure 16.4.

4. The *stochastic potential* of the 1-tree is $\frac{4}{5}m$, while that of the 2-tree is $\frac{1}{5}m$.¹
5. The *stochastically stable state* is contained in the tree with the minimum stochastic potential. Since $\frac{1}{5}m < \frac{4}{5}m$, the stochastically stable state is s_2 . This is because the resistance (or cost) of going from state s_1 to state s_2 is lower than the resistance in the opposite direction. In other words, the equilibrium that has a wider basin of attraction is more difficult to escape from. Hence, it is the one that is stochastically stable.

Under deterministic dynamics, s_1 and s_2 are both asymptotically stable steady states. However, under stochastic dynamics, the multiplicity of the steady states is narrowed down to the unique stochastically stable state s_2 . Using Figure 16.3, and the definition of risk dominance for symmetric 2×2 games we get

$$(9 - 7)(9 - 7) < (8 - 0)(8 - 0).$$

Hence, the stochastically stable steady state picks out the risk dominant equilibrium in this symmetric 2×2 coordination game. This result was shown in Kandori et al. (1993) and Young (1993a). For a discussion of this result under more general assumptions, see Blume (2003).

16.2.2 A model of technology adoption

Consider the following model of technology adoption that is adapted from Young (2006). There is a large number of individuals and three possible technologies, s_1 , s_2 , and s_3 . In any period, a single player (the “experimenting player”) is chosen randomly to make a technology choice. As in the coordination game in Section 16.2.1, the experimenting player asks a sample of m other individuals about their existing choice, who report back honestly. The number of other individuals found to be using technology s_i is y_i , $i = 1, 2, 3$. Thus, it is *as if* the experimenting player is “playing the field” where the proportion of s_i players is given by $\frac{y_i}{m}$. The payoff matrix in pairwise

¹ If there are several nodes in a recurrence class then the stochastic potential for a j -tree is the minimum value of resistance across all the possible nodes of the j -tree (see the model of technology adoption below and Appendix C).

	s_1	s_2	s_3
s_1	7, 7	0, 0	0, 0
s_2	0, 0	3, 3	0, 0
s_3	0, 0	0, 0	5, 5

Figure 16.5 Payoff matrix for the technology adoption game.




Figure 16.6 1-trees (first row), 2-trees (second row), 3-trees (third row) for the technology adoption game.

bilateral matchings is shown in Figure 16.5. There are three Nash equilibria: all adopt s_1 , all adopt s_2 , and all adopt s_3 .

The experimenting player then plays a best response to the observed frequencies of play with some high probability $1 - \varepsilon$, where $\varepsilon > 0$ is small. With probability ε , the player experiences a mutation, or makes a mistake, and randomizes equally between the three technology choices, s_1, s_2, s_3 . Hence, we have a perturbed process with noise ε . In any period, behavior is unpredictable.

The aim is to find the stationary distribution of play. We proceed by identifying the recurrence classes of the unperturbed dynamics in which players play a best response with probability 1. There are three recurrence classes, $C_1 = \{s_1\}$, $C_2 = \{s_2\}$, and $C_3 = \{s_3\}$; each is a singleton and contains one absorbing state of the system. We now draw the relevant j -trees. There are three trees: A 1-tree that is rooted at s_1 , a 2-tree that is rooted at s_2 , and a 3-tree that is rooted at s_3 . In each case, there are three possible ways of reaching a root, starting from one of the terminal nodes. For instance, for the 1-tree, the possibilities are shown in the first row of Figure 16.6. These are: (i) $s_3 \rightarrow s_2 \rightarrow s_1$, (ii) $s_3 \rightarrow s_1, s_2 \rightarrow s_1$, and (iii) $s_2 \rightarrow s_3 \rightarrow s_1$. The second and third rows in Figure 16.6 show the possibilities for a 2-tree and a 3-tree.

Next, we calculate the *path of least resistance* at each directed edge of each tree. For instance, consider the left most tree in the first row in Figure 16.6. Suppose that initially all players play s_3 and we have consecutive mutations that involve the play s_2 . What fraction of players, $\frac{\gamma}{m}$, do we need to play s_2 , so that s_2 becomes a best response? This is given by the solution to $3\frac{\gamma}{m} \geq$

$5(1 - \frac{y}{m})$ or $\frac{y}{m} \geq \frac{5}{8}$, thus, $\rho_{32} = \frac{5}{8}$, as shown in Figure 16.6. Now suppose that all players play s_2 . What fraction of players $\frac{y}{m}$ should play s_1 , so that s_1 becomes a best response? Using the required computation $7\frac{y}{m} \geq 3(1 - \frac{y}{m})$, we get $\frac{y}{m} \geq \frac{3}{10}$, thus, $\rho_{21} = \frac{3}{10}$, which completes the calculations for the leftmost tree in the first row of Figure 16.6. One can proceed in this manner to compute all the numbers in Figure 16.6 for the minimum resistance, ρ_{ij} , of going from recurrence class i to the recurrence class j .

We now calculate the *stochastic potential* of each tree. For each tree we calculate the lowest resistance of an ij -path. Consider first the three 1-trees in the first row of Figure 16.6. The resistance of the path to s_1 for the three trees, from left to right, is $\frac{5}{8} + \frac{3}{10}$, $\min\{\frac{5}{12}, \frac{3}{10}\}$, and $\frac{3}{8} + \frac{5}{12}$, or 0.925, 0.30, and 0.79. Clearly the lowest stochastic potential, 0.30, among the trees in the first row of Figure 16.6, is found for the middle tree.

Finally, in order to find the stochastically stable state, we choose the tree with the least stochastic potential. This turns out to be the case for the 1-tree in the middle (stochastic potential 0.30). Hence, the root of that tree, i.e., the technology s_1 , is the stochastically stable state. Notice also that it is also the only symmetric Nash equilibrium that is Pareto optimal.

16.3 Perturbed Markov dynamics in 2×2 games

Consider a symmetric game in which the set of players is $N = \{1, \dots, n\}$. The model is consistent with each player drawn from a different population of players. Player $i \in N$ has the pure strategy set S_i and let $S = S_1 \times \dots \times S_n$. The payoff function is given by $\pi_i : S \rightarrow \mathbb{R}$, $i \in N$. All players know that the game is symmetric with respect to strategies and payoffs. Suppose that we are in some baseline time period, t . Then the profile of strategies played by all players, τ time periods ago (relative to the baseline time) is denoted by $s^{t-\tau} = (s_1^{t-\tau}, s_2^{t-\tau}, \dots, s_n^{t-\tau}) \in S$; $s^{t-\tau}$ is called the τ -period lagged population strategy profile. We denote by $\pi_i^{t-\tau}$, the payoff of the i th player, τ periods ago when he choose the strategy $s_i^{t-\tau}$. Denote by Σ_i , the set of mixed strategy profiles for player i . The essential features of the game, G , are as follows.

1. *Limited memory*: All players have a finite memory of $T \geq 0$ past periods. If $T = 0$, then players have no memory, except for current strategies/payoffs. A *state* of the system, ω^t , at the end of period t (or, just at the beginning of period $t + 1$), is the current strategy profile and a vector of T -period lagged population strategy profiles, i.e.,

$$\omega^t = [s^t, s^{t-1}, \dots, s^{t-T+1}]. \quad (16.4)$$

Suppose that the initial state is ω^0 . Ω^T is the set of all possible states of the form given in (16.4), or the *state space*; notice that each state lasts over T periods. Payoffs in any period are uniquely determined by the population strategy profile in that period. This information is known to all players.

2. *Sampling from memory*: At any time $t + 1$, players look back at their T period memory of the game, as in (16.4), and take a random sample of observations from the last $\tau \leq T$ periods. In the special case of no memory, $T = 0$; this rule simply reduces to imitate the best current action by any player.
3. *Behavioral strategies*: Each player takes the sample of observations in step 2 as the frequencies with which others will play in the current period. Each player then plays a best reply

to these frequencies with a high probability $1 - \varepsilon$, and with a low probability $\varepsilon > 0$ chooses to randomize among the set of pure strategies.

In conjunction, these rules describe a Markov process on Ω^T , denoted by $P^{T,\tau,\varepsilon}$; sometimes this is referred to as adaptive play with memory length T , sample size τ , and error rate ε . The Markov transition matrix in this case, P^ε , has elements of the form $P_{\omega,\omega'}$, which give the probability of moving from one T -period history, ω , to another T -period history, ω' , from one period to the next. There is some probability of going from any one state to any other state in finite time, hence, the Markov chain is irreducible. The typical probability vector at time t is denoted by the row vector $\lambda_t = [\lambda_t(s_1), \lambda_t(s_2), \dots, \lambda_t(s_n)]$, where $(s_1, \dots, s_n) \in S$, i.e., $\lambda_t(s_i)$ is the probability that at time t , player i plays his strategy s_i . Denote the probability vector corresponding to the Markov process $P^{T,\tau,\varepsilon}$ by $\lambda^{T,\tau,\varepsilon}$. We shall mainly use the following proposition to find the stochastically stable states of the underlying stochastic process.

Proposition 16.28 (Young, 1993a): Consider a regular, perturbed, Markov process $P^{T,\tau,\varepsilon}$ (Definition 16.63 in Appendix C) with the associated transition matrix P^ε . Let $P^{T,\tau,0}$ have the stationary distribution λ^* , i.e., $\lambda^* = \lambda^* P^0$. Then there exists a unique stationary distribution $\lambda^{T,\tau,\varepsilon}$ such that $\lim_{\varepsilon \rightarrow 0} \lambda^{T,\tau,\varepsilon} = \lambda^*$. The set of stochastically stable states (those with $\lambda^*(s) > 0$) are the recurrent classes (Definition 16.56 in Appendix C) of $P^{T,\tau,\varepsilon}$ that have minimum stochastic potential

We now consider two applications of this framework.

16.3.1 Dominated strategies as stochastically stable states

The example in this section is drawn from Alós-Ferrer (2000). The results are sensitive to the choice of behavioral rule followed. In order to expose the reader to a variety of behavioral rules, here we illustrate the solution when instead of best replies, players play *imitate the best*.² Clearly, the results differ when players play best replies. The assumptions are as follows.

1. *Limited memory*: All players have a finite memory of $T \geq 0$ past periods and the state, or history, at the end of time t is given by (16.4).
2. *Bounded rationality and imitation*: Players do not optimize in the classical sense. They look back at their T period memory of the game, and *imitate the best*. In other words, at time $t + 1$, they choose the strategy of the player that generated the highest payoff in periods $t, t - 1, \dots, t - T + 1$. Denote this behavioral rule for the i th player by $b_i(\omega^t)$. Denote by $\pi_j^{t-\tau}$, $0 \leq \tau \leq T + 1$, the payoff of player $j \in N$ from his strategy s_j , τ time periods ago, then

$$s_j \in b_i(\omega^t) \Leftrightarrow \pi_j^{t-\tau} \geq \pi_k^{t-t'} \text{ for all } k = 1, \dots, n \text{ and } \tau, t' = 0, 1, \dots, T + 1. \quad (16.5)$$

In the special case of no memory, i.e., $T = 0$, this rule simply reduces to imitate the best current action by any player. A range of different decision rules is used in several papers.³

² Under the decision rules specified in Kandori et al. (1993) and Rhode and Stegeman (1996), one gets the surprising result that coordination on strictly dominated strategies is a stochastically stable state (see also below). However, under alternative behavioral rules, as, for instance, in Sandholm (1998), convergence to Nash equilibrium is guaranteed.

³ See, for instance, Robson and Vega-Redondo (1996), Vega-Redondo (1997), Kandori et al. (1993).

	A	B
A	4, 4	1, 2
B	2, 1	0, 0

Figure 16.7 A symmetric, 2×2 normal form game.

Source: Alós-Ferrer (2000).

3. *Exogenous inertia*: In any period, each player gets an opportunity to revise his strategy with probability $1 - \theta$. With the complementary probability $\theta > 0$, he is unable to revise. The probability of a revision opportunity is identical and independent across players. The no-inertia case is denoted by $\theta = 0$.
4. *Mutation*: When a player gets a chance to revise his strategy, he chooses the behavioral rule in (16.5) with a probability $1 - \varepsilon$, where $\varepsilon > 0$ is small. And with probability ε , he mutates by randomly choosing any other pure strategy in his strategy set, S .

Consider the symmetric payoff matrix in the 2×2 case, shown in Figure 16.7. We will show that the long-run play might involve strictly dominated strategies. Using iterated deletion of strictly dominated strategies, there is a unique Nash equilibrium (A, A) . The only two recurrent classes are the monomorphic equilibria (A, A) and (B, B) . We will use Ellison's (2000) cost function approach for the proofs in this section (see Appendix C for Ellison's method).

Definition 16.27 Denote by $\text{mon}(A, T)$ the monomorphic state in which both players play A for T periods. Likewise $\text{mon}(B, T)$ is the monomorphic play of B for T periods.

Recall from Appendix C that the cost function $c(x, y)$ gives the minimum number of mutations required to transit from the state $\text{mon}(x, T)$ to the state $\text{mon}(y, T)$.

Lemma 16.2 Under the assumptions made above, and for the game shown in Figure 16.7, $c(B, A) = 2$ independently of T, θ .

Proof: Suppose that in some period, the state is $\text{mon}(B, T)$. We are interested in reaching a *tipping point*, such that no further mutations or noise is needed to reach $\text{mon}(A, T)$. “No further noise is needed” in the sense that, once the tipping point is reached, whenever players get a revision opportunity, they choose to imitate the best, but do not mutate ($\varepsilon = 0$).

Suppose that starting from $\text{mon}(B, T)$, only one player mutates. In this case, we show that the tipping point is not reached. In any period, suppose that one of the two players, the experimenting player, gets an opportunity to revise her strategy (a probability $1 - \theta$ event). Conditional on the revision opportunity, the experimenting player mutates with probability ε and plays A . In the following period, the mutating player gets a payoff of $1 > 0$, while the other player gets $2 > 1 > 0$. With no further mutations, players rely only on imitate the best rule, which prescribes B even to player 1.

Can two simultaneous mutations lead to $\text{mon}(A, T)$, starting from $\text{mon}(B, T)$? Suppose that both players were to simultaneously get a chance to revise their strategies (a probability $(1 - \theta)^2$ event). If both players mutate simultaneously, they both get a payoff of 4; a probability ε^2 event, conditional on the revision opportunity arising for both players. This achieves a tipping point, as no further mutations are needed for all players to play the strategy A . The reason is that in the absence of mutations, players rely on imitate the best rule. Whenever any player gets an

opportunity to revise they will choose A because in their memory, it has led to the highest payoff. Hence, two successive mutations are sufficient, so $c(B, A) = 2$. ■

To apply Ellison's (2000) approach, we need to compute the cost $c(A, B)$ and then compare it to $c(B, A)$.

Proposition 16.29 *If $T = 0$ and $\theta \geq 0$, then the only stochastically stable strategy is B .*

Proof: Recall from (16.4) that memory stretches back to T periods. So $T = 0$ implies that only the current payoffs are remembered in the next period. When $T = 0$, then, $\text{mon}(A, 0) = (A, A)$, i.e., the row and the column players play A . How can we move to the equilibrium $\text{mon}(B, 0)$? A single player mutation to strategy A will not work. Suppose that when a player has a chance to mutate (with ex-ante probability $(1 - \theta)\epsilon$), he chooses B , so that either of the two states (A, B) or (B, A) is achieved.⁴ The mutant who plays B gets a payoff of $2 > 1$, while the non-mutant gets $1 < 2$. Can one arrive now at the state $\text{mon}(B, 0)$ without further mutations, i.e., are we at a tipping point?

The current payoffs corresponding to (A, B) or (B, A) are remembered in the next period ($T = 0$), just prior to taking one's action, and the payoff from playing B is greater. Thus, in the absence of mutation, and relying purely on imitation, the individual not playing B will also switch to B . This argument holds for $\theta \geq 0$, hence, $c(A, B) = 1$. Since, $c(A, B) < c(B, A)$, it is easier to escape from the recurrent class (A, A) than from the recurrent class (B, B) , hence, using the result in Ellison (2000) (see Appendix C) the only stochastically stable strategy is B . ■

The next proposition shows that the only stochastically stable strategy can be in strictly dominated strategies.

Proposition 16.30 *If $T > 0$ and $\theta > 0$, then the only stochastically stable state is $\text{mon}(B, T)$.*

Proof: Suppose that we begin with $\text{mon}(A, T)$. Let a player who has a chance to revise his strategy mutate and play B instead. The mutant obtains 2, while the non-mutant obtains 1. Either of the two states (A, B) or (B, A) is achieved, depending on whether the mutant was a row, or a column player. Is this a tipping point? If yes, then in the absence of further mutations, the state (B, B) can be achieved.

(i) Suppose that subsequent to the mutation, state (A, B) occurs. Then, it is possible that with inertia ($\theta > 0$), no player actually gets an opportunity to revise their strategy for the next T periods. So it is possible that the resulting T -period profile is $(A, B), \dots, (A, B)$. After the T periods have elapsed, the memory of the payoffs (A, A) would have been erased and given an opportunity to revise, a player would mimic the play of B (because $2 > 1$).

(ii) Suppose that subsequent to the mutation, state (B, A) occurs. The argument now proceeds as in (i) because the game is symmetric.

Hence, in each case, only one mutation is needed to achieve the state (B, B) starting from $\text{mon}(A, T)$, so $c(A, B) = 1 < c(B, A) = 2$. Thus, the only stochastically stable strategy is B . ■

Proposition 16.31 *If $T \geq 2$ and $\theta = 0$, then, the only stochastically stable strategy is A . But if $T = 1$, then both A and B are stochastically stable.*

Proof: Suppose that we start at $\text{mon}(A, T)$. As before, let a player who has a chance to revise his strategy mutate and play B instead. The mutant obtains 2, while the non-mutant obtains 1; either

⁴ To reiterate, (A, B) is the state in which the row player plays A and the column player plays B . (B, A) represents the converse situation.

of the two states (A, B) or (B, A) is achieved. We need to check if this is a tipping point, i.e., in the absence of further mutations will one achieve the state (B, B) ?

Since $\theta = 0$ there is no inertia. Thus, at the next opportunity to revise (next period), the highest remembered payoffs are from the state $mon(A, T)$, so any player will choose A . Therefore, one round of mutation is not sufficient to transit to the state (B, B) . Nor will two mutations work. The only possible transition to the state (B, B) is a sequence of mutations such that we have the play (A, B) or (B, A) for T periods, so that the payoffs from the profile (A, A) are erased from memory. In this case, when a player gets an opportunity to revise, he/she would then choose B (because $2 > 1$). Thus, for A to be the stochastically stable state, we require $c(A, B) = T \geq c(B, A) = 2$. Check that if $T = 1$, then $c(A, B) = c(B, A) = 2$, so both states A, B are stochastically stable. ■

16.3.2 Risk dominance under perturbed dynamics

Consider the Markov process $P^{T, \tau, \varepsilon}$ described above and suppose that the game G has several strict Nash equilibria; let a Nash equilibrium be denoted by $s^* = (s_1^*, s_2^*, \dots, s_n^*)$. Consider the state $mon(s^*, T)$, i.e., the Nash equilibrium s^* is played for the last T periods, i.e., for as long as the players can remember. Now suppose that player i wishes to choose one of his strategies in any period. All $\tau \leq T$ sample points will reveal the play of s_{-i}^* by the opponents to which the best response of player i is s_i^* . And this holds true for all players. So, in the absence of any mutations, $\varepsilon = 0$, all players again play the population profile s^* . And they do it every period. Hence, we get the following definition.

Definition 16.28 (Conventions): A convention is a state of the system that is self-perpetuating. As far back as anyone can remember in any period, people have always played it.

If we consider the Markov process without noise, i.e., $P^{T, \tau, 0}$, then we can show that the conventions are the only *absorbing states*; starting from an absorbing state, we cannot reach another state in finite time. Each of the conventions constitutes a *recurrent class*, i.e., the Markov process is not ergodic and is not irreducible. We now ask which conventions does the long-run play converge to? The answer for the unperturbed system is clear—it depends on the history of the game. However, for the perturbed system, the risk dominant equilibrium plays a central role.

Proposition 16.32 Consider a 2×2 coordination game, G , with the associated Markov process $P^{T, \tau, \varepsilon}$, then:

- (a) The unperturbed Markov process $P^{T, \tau, 0}$ converges to one of the conventions.
- (b) The stochastically stable states of the perturbed Markov process $P^{T, \tau, \varepsilon}$ are the risk-dominant conventions.

Proof: Consider the 2×2 coordination game, G , shown in Figure 16.8. By definition, G has two strict Nash equilibria (L, L) and (R, R) .

	L	R
L	a_{11}, b_{11}	a_{12}, b_{12}
R	a_{21}, b_{21}	a_{22}, b_{22}

Figure 16.8 A 2×2 coordination game.

Recall from (16.4) that a state in this model is an entire T -period history of play. Thus, in terms of the two conventions, we can define the two states as the following T -period population profiles,

$$\omega_1 = [(L, L), (L, L), \dots, (L, L)], \omega_2 = [(R, R), (R, R), \dots, (R, R)], \quad (16.6)$$

or we may denote them, respectively, by $\text{mon}(L, T)$ and $\text{mon}(R, T)$. To prove (a) we need to show that the only recurrent classes correspond to the two conventions. In other words, we want to show that the union of the two basins of attractions $B(\omega_1)$ and $B(\omega_2)$ comprises the entire state space. Consider the unperturbed Markov process $P^{T, \tau, 0}$ and begin from any convention. The best response of any player to a sample of size τ , given that the Nash equilibrium is strict, is to play her part of the Nash equilibrium profile in every period.

(b) Denote by ρ_{12} , the minimum resistance of all paths that begin in the state ω_1 and terminate in the basin of attraction of the other state, $B(\omega_2)$. Conversely, let ρ_{21} be the minimum resistance of all paths that start in the state ω_2 and terminate in the basin of attraction $B(\omega_1)$. Once a path reaches the basin of attraction of any state, then no further noise, or mutations, are needed to converge to that state; best responses, by themselves, guarantee such convergence to that state. For any real number x let $\lceil x \rceil$ denote the least integer greater than or equal to x ; so if $x = 1.7$ then $\lceil x \rceil = 2$, while if $x = 3$ then $\lceil x \rceil = 3$.

Suppose that the row player thinks that the column player will play R with probability q , then it is optimal to play R if

$$a_{21}(1 - q) + a_{22}q \geq a_{11}(1 - q) + a_{12}q,$$

or,

$$q \geq \underline{q} = \frac{a_{11} - a_{21}}{a_{11} + a_{22} - a_{21} - a_{12}}. \quad (16.7)$$

Suppose that we are in the state ω_1 where all play L . Then, for the row player to decide to switch from L to R , he must have a sample from the last τ periods such that there are $\lceil \underline{q}\tau \rceil$ instances of play of R by the column player, where $\underline{q}$ is given by (16.7).

Let there be at least $\lceil \underline{q}\tau \rceil$ successive plays of R by the column player. So the probability assigned to the column player playing R is at least $\frac{\lceil \underline{q}\tau \rceil}{\tau}$, which is greater than or equal to $\underline{q}$. Thus, it would be optimal for the row player to play R . How can the row player observe such a sample starting from the state ω_1 ? The probability of making one mutation is ε , so the probability of making $\lceil \underline{q}\tau \rceil$ mutations is $\varepsilon^{\lceil \underline{q}\tau \rceil}$ (because mutations are independent across time periods).

Analogously, starting in state ω_1 , we ask what will induce the column player to play R ? Suppose that the column player believes that the row player plays R with probability p . Proceeding as in the case of the row player, it must be that the row player makes a mutation and plays R for $\lceil \underline{p}\tau \rceil$ successive instances in the sample chosen by the column player, where

$$\underline{p} = \frac{b_{11} - b_{21}}{b_{11} + b_{22} - b_{21} - b_{12}}. \quad (16.8)$$

The probability of this happening is $\varepsilon^{\lceil \underline{p}\tau \rceil}$. Since we are interested in the path of least resistance, define

$$\rho_{12} = \min \left\{ \lceil \underline{q}\tau \rceil, \lceil \underline{p}\tau \rceil \right\}. \quad (16.9)$$

Why is this sufficient to ensure that play enters $B(\omega_2)$? Given (16.9), at least one of the players plays R . Suppose, for the sake of argument, that this is the row player who has observed $\lceil \underline{q}\tau \rceil$ instances of play by the column player in his sample. The continuing play of R by the row player leads the column player to gradually assign a greater probability to the row player playing R . When the row player is observed to play R for $\lceil \underline{p}\tau \rceil$ successive instances (recall that we are interested in the path of least resistance), it becomes optimal for the column player to also play R .

By a completely analogous argument, one can find that

$$\rho_{21} = \min \left\{ \lceil (1 - \underline{q})\tau \rceil, \lceil (1 - \underline{p})\tau \rceil \right\}. \quad (16.10)$$

From Appendix C, we know that ω_1 is stochastically stable if, and only if, it has relatively greater minimum resistance. In other words, we need a relatively larger number of successive mutations to escape from the basin of attraction of state ω_1 as compared to state ω_2 . Thus, ω_1 is stochastically stable if, and only if,

$$\rho_{12} \geq \rho_{21}. \quad (16.11)$$

Substituting (16.9), (16.10) in (16.11), we get

$$\min \left\{ \lceil \underline{q}\tau \rceil, \lceil \underline{p}\tau \rceil \right\} \geq \min \left\{ \lceil (1 - \underline{q})\tau \rceil, \lceil (1 - \underline{p})\tau \rceil \right\}. \quad (16.12)$$

From the Appendix on game theory we know that (L, L) is the risk dominant equilibrium if

$$\min \left\{ \underline{q}, \underline{p} \right\} \geq \min \left\{ 1 - \underline{q}, 1 - \underline{p} \right\}. \quad (16.13)$$

Risk dominance of (L, L) (see (16.13)) implies (16.12) for a large enough sample, τ . ■

16.3.3 Extensions of the basic model

In the models considered so far, there was homogeneity with respect to (i) the length of samples drawn by each player, and (ii) the rate of mutation of each player, as captured in a common term, ε . Furthermore, if there was a separate population of row players and column players, then there was within-population homogeneity of payoffs. All these assumptions can be relaxed. See Young (1998) for a nice exposition of these issues that we draw on below.

(1) *Heterogeneity in sample size*: Suppose that the row player uses a sample of length τ_R , while the column player uses a sample of length τ_C . The players have the same mutation rate, ε . We say that a player is *better informed* if he has a longer sample. An important consideration here is to see if better informed players receive a relatively higher payoff. For the generic 2×2 coordination

	L	R
L	8, 6	0, 0
R	0, 0	5, 7

Figure 16.9 A coordination game with a *better informed* player.

game, G , in Proposition 16.32, the analog of (16.12), which guarantees that (L, L) is stochastically stable, is given by:

$$\min \left\{ \left\lceil \underline{q} \tau_R \right\rceil, \left\lceil \underline{p} \tau_C \right\rceil \right\} \geq \min \left\{ \left\lceil (1 - \underline{q}) \tau_R \right\rceil, \left\lceil (1 - \underline{p}) \tau_C \right\rceil \right\}. \quad (16.14)$$

For illustrative purposes, consider the game in Figure 16.9.

Define the two states as $\omega_1 = \text{mon}(L, T)$ and $\omega_2 = \text{mon}(R, T)$, and let the sample sizes of the row and column players be, respectively, τ_R and τ_C . Using the definitions of $\underline{q}, \underline{p}$ in (16.7), (16.8) we get

$$\underline{q} = \frac{8}{13}; \underline{p} = \frac{6}{13}; 1 - \underline{q} = \frac{5}{13}; 1 - \underline{p} = \frac{7}{13}. \quad (16.15)$$

Substituting (16.15) in (16.14) we get

$$\min \left\{ \left\lceil \frac{8}{13} \tau_R \right\rceil, \left\lceil \frac{6}{13} \tau_C \right\rceil \right\} \geq \min \left\{ \left\lceil \frac{5}{13} \tau_R \right\rceil, \left\lceil \frac{7}{13} \tau_C \right\rceil \right\}. \quad (16.16)$$

For large values of x , $x \approx \lceil x \rceil$. Suppose that τ_R, τ_C are large. One set of circumstances in which (L, L) is stochastically stable is the following. Suppose that (i) $\frac{8}{13} \tau_R \leq \frac{6}{13} \tau_C$, or equivalently $\frac{\tau_R}{\tau_C} \leq \frac{3}{4}$, and (ii) $\frac{7}{13} \tau_C \geq \frac{5}{13} \tau_R$, so that $\frac{\tau_R}{\tau_C} \leq \frac{7}{5}$. Let $\frac{\tau_R}{\tau_C} \leq \min \left\{ \frac{3}{4}, \frac{7}{5} \right\} = \frac{3}{4}$, thus, the row player takes a sample no larger than 75% of the sample taken by the column player. Then, for the game shown in the game in Figure 16.9, (L, L) is stochastically stable.

The row player prefers the equilibrium (L, L) to the equilibrium (R, R) . Paradoxically, having less information helps the row player to achieve a better outcome. The intuition is that less information imparts greater flexibility to the row player in reacting to errors by the opponent, which are to his advantage. However, “less information is better” is by no means a robust result. It can be easily overturned by altering the payoffs. The drawback of less information is that one is too quick to react to errors from the opponent, which are unfavorable.

(2) *Heterogenous payoffs and alternative noise specifications*: Consider two distinct populations of row and column players. In each period, there are pairwise interactions between row and column players. There is within-population heterogeneity. Suppose that there are n_r row players and n_c column players. Suppose that the sample size of the r th row player is $\tau(r)$, $r = 1, \dots, n_r$ and the sample size of the c th column player is $\tau(c)$, $c = 1, \dots, n_c$. The payoff matrix of the r th row player is given by $[a_{ij}(r)]_{2 \times 2}$ and that of the c th column player is given by $[b_{ij}(c)]_{2 \times 2}$, where the subscripts i, j refer, respectively, to the i th row and the j th column. This gives rise to within-population heterogeneity of payoffs.

For the r th row player, compute $\underline{q}(r)$, and for the c th column player compute $\underline{p}(c)$. Then the type of a row player can be denoted by $(\tau(r), \underline{q}(r))$, while the type of a column player can be denoted by $(\tau(c), \underline{p}(c))$. Consider the coordination game in Proposition 16.32 and let (L, L) and (R, R) be the two strict Nash equilibria. The main result is that convention (L, L) is stochastically stable if the following analog of (16.12) holds:

$$\min \left\{ \left\lceil \underline{q}\tau(r) \right\rceil, \left\lceil \underline{p}\tau(c) \right\rceil : \forall r, c \right\} \geq \min \left\{ \left\lceil (1 - \underline{q})\tau(r) \right\rceil, \left\lceil (1 - \underline{p})\tau(c) \right\rceil : \forall r, c \right\}. \quad (16.17)$$

The proof is along the same lines as Proposition 16.32.

So far we have modeled noise as follows. When a player gets an opportunity to revise his strategy, then with some small probability ε , he randomly chooses one of his strategies with equal probability. What if, instead, the player has a preference for particular strategies and, when mutating, chooses among his pure strategies with unequal probabilities. To fix ideas, suppose that we consider the generic 2×2 game, as in Proposition 16.32.

Suppose that each player, when mutating, chooses strategy L with probability $\rho \neq 0.5$ and strategy R with probability $1 - \rho$.⁵ Define the states ω_1, ω_2 as in (16.6). Young (1998) shows that the probability of transiting from state ω_1 to ω_2 is proportional to $[(1 - \rho)\varepsilon]^{\min\{\lceil \underline{q}\tau_R \rceil, \lceil \underline{p}\tau_C \rceil\}}$ and the probability of transiting in the opposite direction is proportional to $[\theta\varepsilon]^{\min\{\lceil (1 - \underline{q})\tau \rceil, \lceil (1 - \underline{p})\tau \rceil\}}$. As an example, suppose that the players have a strong preference to play R when they mutate, i.e., ρ is very close to zero. In this case, the probability of transiting from ω_1 to ω_2 is relatively higher than that in the opposite direction. State ω_2 would then be the stochastically stable state.

16.4 Memory, conventions, and risk dominance in 3×3 games

Not all of the results for 2×2 coordination games that we have considered above, hold for higher dimensional games. The main result is that the correspondence between risk dominance and stochastic stability disappears. However, for a certain class of games, the set of *acyclic games* (defined below), we can speak about the relation between the set of Nash equilibria, minimum stochastic potential, and stochastic stability.

Consider the generic 3×3 coordination game in Figure 16.10. The three strict Nash equilibria are (L, L) , (M, M) , (R, R) ; these are also the three coordination equilibria of the game.

	L	M	R
L	a_{11}, b_{11}	a_{12}, b_{12}	a_{13}, b_{13}
M	a_{21}, b_{21}	a_{22}, b_{22}	a_{23}, b_{23}
R	a_{31}, b_{31}	a_{32}, b_{32}	a_{33}, b_{33}

Figure 16.10 A generic 3×3 coordination game.

⁵ For other specifications of the noise structure, see Blume (1993, 1995).

	<i>L</i>	<i>M</i>	<i>R</i>
<i>L</i>	5, 8	0, 1	1, 0
<i>M</i>	1, 0	6, 7	0, 1
<i>R</i>	0, 1	1, 0	7, 6

Figure 16.11 No strict risk dominant equilibrium.

Source: Republished with permission of Princeton University Press, from *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions* by H. Peyton Young (1998), permission conveyed through Copyright Clearance Center, Inc.

Definition 16.29 (*Risk dominance for 3×3 coordination games*): Suppose that the indices $i, j \in \{L, M, R\}$. We say that equilibrium (i, i) risk dominates equilibrium (j, j) if

$$(a_{ii} - a_{ji})(b_{ii} - b_{ij}) > (a_{jj} - a_{ij})(b_{jj} - b_{ji}); i \neq j; i, j \in \{L, M, R\}. \quad (16.18)$$

Definition 16.30 (*Strict risk dominance*) An equilibrium is strictly risk dominant if it risk dominates all the other coordination equilibria.

Example 16.9 There might be no strict risk dominant equilibrium; see, e.g., the game in Figure 16.11. For this game, we get the following cycle: (L, L) risk dominates (R, R) , (R, R) risk dominates (M, M) , (M, M) risk dominates (L, L) . Hence, there is no strict risk dominant equilibrium.

The problem raised in Example 16.9 is compounded by the fact that, for 3×3 coordination games (and games of higher order) risk dominance might not coincide with stochastic stability. We demonstrate this by using the payoff matrix in Figure 16.12.

Since all the off-diagonal elements in Figure 16.12 are zero, (16.18) reduces to the following statement: Equilibrium (i, i) risk dominates equilibrium (j, j) if:

$$a_{ii}b_{ii} > a_{jj}b_{jj}; i \neq j; i, j = L, M, R. \quad (16.19)$$

In this special case, risk dominance simply requires us to check for the highest product of individual payoffs when they play identical strategies. Hence, the risk dominant equilibrium is (L, L) . We now show that this equilibrium is not stochastically stable. We consider the perturbed Markov process $P^{T, \tau, \varepsilon}$, i.e., each player has a memory of T periods, uses a sample size of τ periods and

	<i>L</i>	<i>M</i>	<i>R</i>
<i>L</i>	5, 60	0, 0	0, 0
<i>M</i>	0, 0	7, 40	0, 0
<i>R</i>	0, 0	0, 0	100, 1

Figure 16.12 A 3×3 coordination game.

Source: Republished with permission of Princeton University Press, from *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions* by H. Peyton Young (1998), permission conveyed through Copyright Clearance Center, Inc.

mutates (i.e., chooses each pure strategy with equal probability) with probability ε ; otherwise, with probability $1 - \varepsilon$, he plays his best response to the frequencies of observed play of the opponents in the sample.

Define the three states as follows.

$$\omega_1 = \text{mon}(L, T), \omega_2 = \text{mon}(M, T), \omega_3 = \text{mon}(R, T).$$

Thus, *each state is a T -period history* in which one of the three conventions, L , M , or R , is played by each player. Suppose that we are in some state $\omega_j = \text{mon}(j, T)$, $j = L, M, R$. How can the system then tip into the basin of attraction of another state $\omega_k = \text{mon}(k, M)$, $j \neq k$, $k = L, M, R$?

Since we are interested in the path of least resistance, we proceed as follows. What is the minimum number of times the column player must consecutively play strategy k such that when the row player samples the past τ periods, he finds it a best response to play strategy k as well? Suppose that the required frequency of play by the column player is q^* ; q^* equals the minimum number of required consecutive plays of k by the column player, divided by τ . Since we began in the state $\omega_j = \text{mon}(j, T)$, in the relevant memory of the row player, strategies k and j are played with respective frequencies q^* and $1 - q^*$. Thus, for the row player, q^* can be found as the solution to

$$q^* a_{kk} = (1 - q^*) a_{jj} \Leftrightarrow q^* = \frac{a_{jj}}{a_{jj} + a_{kk}}. \quad (16.20)$$

Starting from state ω_j , we now find the minimum number of times the row player must play strategy k such that using his sample of past τ observations, the column player also finds it optimal to play strategy k . Denote the resulting frequency of play of k by the row player by p^* . Analogous to (16.20), we can find p^* as follows.

$$p^* b_{kk} = (1 - p^*) b_{jj} \Leftrightarrow p^* = \frac{b_{jj}}{b_{jj} + b_{kk}}. \quad (16.21)$$

Since we are interested in the path of least resistance, as in (16.9), we find that the minimum resistance, ρ_{jk} , of moving from state ω_j to state ω_k is given by

$$\rho_{jk} = \min \{ \lceil q^* \tau \rceil, \lceil p^* \tau \rceil \}. \quad (16.22)$$

It is important to note that for 3×3 games, the formula in (16.22) holds only when the off-diagonal elements are zero. Such a formula always holds for 2×2 games, even if the off-diagonal elements are non-zero. For 3×3 games on the other hand, the path of least resistance from state ω_j to state ω_k might actually involve the intermediate play of some other state ω_i . For games of even higher dimension, the path of least resistance from one state to another might involve the intermediate play of several other states. Hence, higher dimensional games become even more difficult to handle.

Once we are able to compute the minimum resistance in moving from any one state to another, using (16.22), the stochastically stable state has the minimum stochastic potential (see Proposition 16.28). As usual, we draw, in Figure 16.13, the rooted 3-trees (see Appendix C and Figure 16.6, above) to compute the minimum stochastic potential. This involves the computation of the minimum resistance for each of the 3-trees in each row, and then the choice of the row with the minimum resistance across all rows. A simple calculation shows that the stochastically stable state is ω_2 , while the risk dominance criterion chooses state ω_1 .




Figure 16.13 Non-equivalence between risk dominance and stochastic stability

The main results for games with dimension higher than 2×2 are for the class of *acyclic games* (Young, 1993a). Essentially, these are games such that the best reply graph has no directed cycles. Another method of identifying acyclic games is described below.

Let the set of players be $N = \{1, \dots, n\}$ and let all players have identical strategy sets. Player $i \in N$ has the pure strategy set S and the possible strategy profiles of players belong to $S = S \times \dots \times S$. The utility function is given by $u_i : S \rightarrow \mathbb{R}$, $i \in N$. The game is described by $G(N, S, u)$.

Definition 16.31 (*Weighted potential game*) G is a weighted potential game if for $i \in N$ there exists weights $w_i \in \mathbb{R}_+$ and a potential function $\phi : S \rightarrow \mathbb{R}$, such that the following restriction is satisfied.

$$w_i u_i(s'_i, s_{-i}) - w_i u_i(s_i, s_{-i}) = \phi(s'_i, s_{-i}) - \phi(s_i, s_{-i}) \quad \forall s'_i, s_i \in S_i, s_{-i} \in S_{-i} \text{ and } \forall i \in N.$$

The idea behind the weighted potential game is that the utility functions of players can be rescaled such that the change in own utility arising from a unilateral change in strategy equals the change in potential. Every weighted potential game has at least one Nash equilibrium in pure strategies.

Example 16.10 Consider the 2×2 symmetric game in Figure 16.7. The potential function for this game is

$$\phi(A, A) = 2, \phi(A, B) = \phi(B, A) = 0, \phi(B, B) = -1.$$

When $w_i = 1$ for all $i \in N$, then we call G in Definition 16.31 simply a potential game. A generalization of this game is the ordinal potential game.

Definition 16.32 (*Ordinal potential game*) G is an ordinal potential game if there is an ordinal potential function $\phi : S \rightarrow \mathbb{R}$, such that the following restriction is satisfied.

$$\text{sign}[u_i(s'_i, s_{-i}) - u_i(s_i, s_{-i})] = \text{sign}[\phi(s'_i, s_{-i}) - \phi(s_i, s_{-i})] \quad \forall s'_i, s_i \in S_i, s_{-i} \in S_{-i} \text{ and } \forall i.$$

Definition 16.33 (*Acyclic game*) If G is an ordinal potential game such that no player is indifferent between any two distinct strategies, then G is also an acyclic game.

Coordination games are examples of acyclic games. The reason is that once we reach any of the coordination equilibria, then there are no directed edges pointing away from such equilibria. This eliminates the possibility of directed cycles. There are two main results for acyclic games. (1) The unperturbed process, $P^{T,\tau,0}$, converges to one of the conventions *from any arbitrary initial state*. (2) For small enough noise, ε , and small $\frac{\tau}{T}$, in a perturbed Markov process, $P^{T,\tau,\varepsilon}$, play converges with arbitrarily high probability to the convention which minimizes the stochastic potential.

16.4.1 Boundedly rational play and fully rational outcomes

Unlike the textbook fully rational players, players in models of stochastic evolution exhibit limited memory, myopia, and adaptive learning (e.g., forming expectations based on a sample of past observations). Should we expect the nature of economic outcomes to be different (ignoring issues of the time taken to achieve equilibrium)? Not necessarily. We have shown that the risk dominant equilibrium arises as a stochastically stable outcome for 2×2 games. When the off-diagonal elements are zero, then, this outcome corresponds to the Pareto optimal outcome in coordination games. In higher dimensional games, similar results arise for the class of acyclic games.

Despite being myopic, if players play static best replies, then they do not play strictly dominated actions.⁶ However, this framework rules out dominated actions based on a low rationality argument, which stands in contrast to classical justifications based on common knowledge and infinite regress of beliefs. Models of stochastic evolution are also able to show that the asymmetric Nash bargaining solution arises from a low rationality environment (Young, 1993b). As Young (1998) puts it: “Interpreted in this way, the evolutionary approach is a means of reconstructing game theory with minimal requirements about knowledge and rationality.”

16.4.2 Of norms and social evolution

Conventions are rules of individual behavior that coordinate expectations on some action profile of the participants. For instance, in meetings between two people, there are a range of conventions, e.g., shaking hands, hugging, and kissing on the cheeks. Conventions typically arise when there are multiple equilibria and one needs a mechanism to coordinate expectations on a particular equilibrium. A *norm* is a slightly stronger form of convention that carries with it some social sanction, if violated.

Norms can be enforced through a variety of mechanisms. Norms are self-enforcing when the need for a coordination motive makes it optimal for individuals to follow them. In some cases, it is the threat of punishment, e.g., social ostracization of some sort, that makes it incentive compatible for players to comply with norms. Other mechanisms for norm enforcement are feelings of shame, guilt, or a tug on one’s conscience, that arise when a norm is not followed (Elster, 1989, 1999; Coleman, 1990). In this spirit, Hume ([1739]1978) used the concept of a virtuous or right action. In many cases, however, the boundaries between norms and conventions are quite blurred, so we use norms to refer to both situations.⁷

How do norms come about? One simplistic answer is that they simply come about through the legal framework, such as rules, laws, official edicts, and judicial rulings. We considered a powerful argument for social norms in Part 4 based on the work of Gintis (2009b) in which social norms

⁶ Under alternative behavioral rules such as *imitate the best*, players may play dominated strategies, see Section 16.3.1.

⁷ For readers interested in pursuing these issues further, see Sugden (1986), Posner (2000), and Bicchieri (2006).

play the role of a correlating device to unify societal expectations. The stochastic evolutionary approach argues that norms come about by the adaptive learning efforts of a large number of individually small players. Furthermore, in many cases, laws and edicts follow the establishment of norms and simply codify what already exists. It is quite straightforward to apply the insights developed in this section to some simple examples of norms.

It turns out that norms that are also stochastically stable states, need not always be efficient (Young, 1998). To illustrate this result, suppose that there is a single population. Individuals are matched each period to engage in a bilateral transaction. Suppose there are two possible commodities, gold (G) and silver (S), that can be used as a medium of exchange. Each period, one individual, termed as the experimenting player, or the *experimenter*, is drawn at random. He chooses either to carry gold or silver. The selected individual has memory T , sample size τ , and he plays with some noise ε . This results in a perturbed Markov process, $P^{T,\tau,\varepsilon}$, over the set of all T period histories.

The payoff matrix is given by the sum of two different payoffs. If two randomly paired individuals carry the same currency, they can engage in a transaction and get transaction utility T_g (if they both carry gold) or T_s (if they both carry silver); otherwise, transaction utility is zero. Hence, based on transaction utility alone, we have a coordination game. Each metal also provides an adornment value that is independent of the transaction value. The adornment value from carrying the two metals G, S is, respectively, A_g, A_s . The matrix of payoffs that represents adornment utility is not a coordination game. Finally, total utility is the sum of transaction and adornment utility, as shown in Figure 16.14.

What norm, G or S, do we expect under stochastic evolution? Efficiency would dictate a choice of the norm, which maximizes

$$T_j + A_j, j = G, S. \quad (16.23)$$

We have seen above that for 2×2 games, the outcome is the risk dominant equilibrium. Using the condition for a risk dominant equilibrium, equilibrium j is the risk dominant equilibrium relative to equilibrium k , $k \neq j$ if and only if $(T_j + A_j - A_k)^2 \geq (T_k + A_k - A_j)^2$, $j, k = G, S$. Suppose that the rightmost payoff matrix in Figure 16.14 describes a coordination game, so $T_j + A_j > A_k$ and $T_k + A_k > A_j$. Then j is the risk dominant equilibrium if

$$T_j + 2A_j \geq T_k + 2A_k. \quad (16.24)$$

From (16.23), efficiency requires maximization of $T_j + A_j$, but from (16.24), stochastic evolution picks out the equilibrium that maximizes $T_j + 2A_j$. Hence, the two need not be the same. In fact, stochastic evolution will pick out the risk dominant equilibrium in this 2×2 game, which

Transaction utility				Adornment utility				Total utility			
G		S		G		S		G		S	
G	T_g, T_g	0, 0	+	G	A_g, A_g	A_g, A_s	=	G	$T_g + A_g, T_g + A_g$	$T_g + A_g, T_s + A_s$	
S	0, 0	T_s, T_s		S	A_s, A_g	A_s, A_s		S	A_s, A_g	$T_s + A_s, T_s + A_s$	

Figure 16.14 Description of utilities from carrying gold (G) or silver (S).

	X	Y	Z
X	5, 1	0, 0	0, 0
Y	0, 0	3, 3	0, 0
Z	0, 0	0, 0	1, 5

Figure 16.15 A 3×3 coordination game in which the two players are drawn from different populations.

Source: Young (2008). Steven Durlant, *The New Palgrave Dictionary of Economics*, published 2008, Palgrave Macmillan UK, reproduced with permission of Palgrave Macmillan.

maximizes $T + 2A$. This equilibrium puts twice the weight on adornment, which is not efficient. Another example of a norm that is not efficiency enhancing is the norm of etiquette, which probably has limited welfare consequences, as used by economists. For a very readable account of this and other issues, see Young (2008).

Consider the 3×3 coordination game in Figure 16.15. Suppose that the row and the column players are drawn from different populations. For instance, the row player could be the principal and the column player, the agent. There are three norms, (X, X) , (Y, Y) , (Z, Z) , that are the Nash equilibria of the game. The distributional implications of the three equilibria are quite distinct. The two equilibria (X, X) and (Z, Z) represent extreme allocations that heavily benefit one party, while (Y, Y) represents a compromise allocation. None of the equilibria Pareto dominates the other.

Each of the three equilibria represents a potential norm. Which gets chosen? Consider, as before, a perturbed stochastic process, $p^{T, \tau, \varepsilon}$. A simple calculation, that the reader should by now be able to work out, shows that the stochastically stable state is the compromise solution (Y, Y) . It can be shown more generally that in similar games, the outcome of stochastic evolution is a compromise solution; see Young (1998). This is suggestive of the endogenous evolution of fairness.

16.4.3 A brief selection of further applications

In an early application of the concept of stochastic stability Young (1993b) asked: what sort of bargaining norm would arise from limited rationality models, as, for instance, in adaptive play $p^{T, \tau, \varepsilon}$? In this model players are matched randomly from two populations and play a Nash demand game. The two matched players demand respective shares s_1 and s_2 of a pie of fixed size 1. If $s_1 + s_2 \leq 1$, then both players get their demands, otherwise they get nothing.

Young (1993b) showed that when the demands by players are sufficiently divisible and T, τ are sufficiently large, then the stochastically stable bargaining norms are arbitrarily close to the classic Nash bargaining solution. The method of proof should now be standard to the reader. One begins with a history where a particular norm, s_1 or s_2 , is prevalent; one of the norms is the Nash bargaining solution. One then calculates the least resistance required to reach the basin of attraction of another norm and infers the stochastic stability of alternative norms.

Myatt and Wallace (2005) consider public goods games with threshold externalities. Suppose that we have $N = \{1, \dots, n\}$ players. Player $i \in N$ chooses a contribution $g_i \in \{0, 1\}$. If at least $m \geq 2$ players contribute a unit amount, then the public good materializes, otherwise not. Suppose that the benefit to each player from the public good is v and the cost of contribution is c ;

we assume $v > c$. If the public good materializes, then the respective utilities of contributors and non-contributors are $u = v - c$ and $u = v$. If fewer than m players contribute, then the respective utilities of contributors and non-contributors are $u = -c$ and $u = 0$.

In the spirit of adaptive dynamics, $P^{T,\tau,\varepsilon}$, the state at any time t is the vector of contributions of all players. One of the players is chosen at random to revise his contribution. With a high probability $1 - \varepsilon$ he plays a best response to his observed sample of past states with memory T and a sample size $\tau \leq T$. However, with probability $\varepsilon > 0$ he makes a mistake and randomly chooses one of his feasible contribution levels.

The main result is that m players will contribute if $v > mc$. The intuition can be conveyed as follows. Start from an equilibrium where nobody contributes; this is a Nash equilibrium. We now calculate the path of least resistance to an equilibrium where m individuals contribute. This involves $m-1$ consecutive mistakes over $m-1$ periods that take the form of making a contribution even when it is not a best response. In the final, m th, period, the individual who is chosen to revise his strategy will find it a best response to contribute, so no further mistakes are needed. Since each contributing player incurs a cost c of contributing, thus, the total cost of this move between the two equilibria is $(m-1)c$.

Now we need to calculate the cost of moving in the opposite direction, starting from the Nash equilibrium in which m players contribute. Suppose that in the next period, a contributing player is given the choice (experimenter), and makes a mistake with probability ε , so chooses not to contribute. Thus, we are left with $m-1$ contributors, so the public good does not materialize. The cost of this to the experimenter is $v - c$. Suppose that in the next period one of the remaining $m-1$ contributors gets a revision opportunity.⁸ It is now a best response not to contribute ($-c < 0$). Thus, no further mistakes are needed. The total cost of the move from an equilibrium with “ m players contribute” to “nobody contributes” is $v - c$.

Putting together the two costs of moving between the states, it is less costly to move in the direction of m players contributing, starting from nobody contribute if $(m-1)c < v - c$, or $v > mc$, as claimed.

Myatt and Wallace (2008) extend this analysis to introduce valuation and cost heterogeneity into the problem. Suppose now that player $i \in N$ has player-specific valuation v_i and cost c_i . The lowest cost of moving from the state where nobody contributes to all-contribute is that the first $m-1$ lowest cost contributors contribute. Ordering the players from lowest cost to highest cost gives the least cost of moving to a contributions equilibrium to be $\sum_{i=1}^{m-1} c_i$. In moving towards the opposite direction (nobody contribute) newer considerations come into play that were absent in the symmetric case.

The most obvious individual who should drop out from the set of m contributors is the one with the lowest value of $v_i - c_i$ (utility forgone from not contributing, when a contributing individual gets a chance to revise his strategy). However, such an individual may not be part of the team of m contributors. Furthermore, it is possible that it may be less costly to exit upwards ($m+1$ contributors) rather than downwards. This, for instance, would be the case if $\min_{1 \leq i \leq m} (v_i - c_i) > c_{m+1}$. It might be cheaper to transit downwards towards the nobody contributes state by first going up to $m+1$ contributors, then shedding the contributor with the highest $v_i - c_i$ difference among the set of $m+1$ contributors. This enables the public good to be provided during these intermediate steps, but reduces the average cost of public good provision, which, in turn, reduces the cost of moving towards the nobody-contribute state.

⁸ Note that we are interested in the path of least resistance in moving to an equilibrium where nobody contributes. Along this path, we must assume at this step that one of the remaining $m-1$ contributors gets a revision opportunity.

The applications to public goods are important and help build our intuition further about using stochastic dynamics in different problems. Fehr and Schmidt (2006) stress the importance of explaining the extensive empirical evidence that has now accumulated on the pattern of public goods contributions with and without punishment (see Part 2 of the book for details). It is an open question if an analysis of public goods based on stochastic dynamics helps us to explain the relevant evidence. It would appear to be a difficult task because considerations of inequity aversion and reciprocity, for which there is good empirical evidence, help us to track closely the evidence, even in the absence of any stochastic components. Furthermore, adaptive dynamics are backward looking, while the possibility of future punishment in public goods games requires a form of forward induction. These concerns also apply to other settings in which one might apply stochastic dynamics to the results from experimental games that involve forward induction.

16.5 Social networks

In many cases, we would like to get closed-form solutions to the limiting distribution of play under stochastic dynamics; this can be obtained in an interesting case (Blume, 1993, 1995) that we now consider.

So far, we have assumed that individuals learn by imitating, or observing strategies based on a sample drawn randomly from the population. In many cases of actual interest, however, individuals learn from other individuals in their vicinity, e.g., geographical neighbors, network of friends, or colleagues, or trusted role models; for a fuller treatment of social networks, see Jackson (2010). Consider n players. To model individuals in the “vicinity” consider the notion of a *social network* (Γ, V, E) , where Γ is a graph, V the set of its vertices and E the set of its directed edges. Each vertex denotes a player.

The directed edge from vertex i to vertex j is denoted by (i, j) . This directed edge denotes the *social influence* of individual j on individual i . The extent of such influence is measured by a weight $\omega_{ij} \geq 0$. In the discussion below, ω_{ij} is considered to be the *geographical proximity* of the two individuals and, so, in this interpretation, $\omega_{ij} = \omega_{ji}$.

Suppose that player i chooses some action s_i^t at time t from his strategy set S_i and $S = S_1 \times \dots \times S_n$. Denote by $s^t = (s_1^t, s_2^t, \dots, s_n^t) \in S$ the action profile at time t ; call this the *state of the system* at time t . Suppose that the opportunity to revise one’s strategies in any period is a Poisson random variable θ , which is distributed independently and identically across players.

The payoff of each player is composed of two parts. The first part is the utility arising from one’s own choice, $v_i(s_i^t)$. The second part is the externality arising to player i from the choices made by others. The externality is given by

$$\chi_i(s^t) = \sum_{j, j \neq i} \omega_{ij} u(s_i^t, s_j^t), \quad (16.25)$$

where $u(s_i^t, s_j^t)$ is the unweighted externality to player i from player j . This is weighted by the geographical proximity of the two players and summed across all other players to compute the total externality on player i . The total utility of player i is then written as

$$V_i(s^t) = v_i(s_i^t) + \chi_i(s^t) + \varepsilon_i^t, \quad (16.26)$$

where ε_i^t is a random preference shock that is distributed independently across players and has an extreme value distribution.

Each player who gets an opportunity to choose a strategy, chooses s_i^t to maximize (16.26), fixing the strategies of the other players. It can then be shown (see Appendix B) that from an observer's point of view, player i chooses pure strategies in his choice set S_i with probabilities given by a logistic function, i.e.,

$$P(s_i^t | s_{-i}^t) = \frac{\exp \beta \{v_i(s_i^t) + \chi_i(s_i^t, s_{-i}^t)\}}{\sum_{s_j^t \in S_i} \exp \beta \{v_i(s_j^t) + \chi_i(s_j^t, s_{-i}^t)\}}, \quad (16.27)$$

where β is the single parameter of the extreme value distribution.

By reacting as if the actions of the other players are fixed (i.e., reacting to the current state) each player who gets an opportunity to revise chooses according to (16.27). This induces a new state of the system, which depends only on the previous state. Hence, this is a Markov process. Whenever, any player gets an opportunity to revise, (16.27) indicates that each of the available strategies is chosen with a positive probability. Hence, beginning from any state in S we can arrive at any other state in S in finite time. And so, the Markov process is irreducible. We know from Appendix C that, in this case, the Markov process is ergodic and has a unique invariant distribution $\lambda^*(s)$, which is the probability with which the process visits state s in the long run.

Define the *potential* of any state $s = (s_1, \dots, s_n)$ by the sum of all individual payoffs plus half of all externality payoffs, i.e.,

$$\rho(s) = \sum_{i=1}^n v_i(s_i) + \frac{1}{2} \sum_{i=1}^n \chi_i(s). \quad (16.28)$$

Under these assumptions, $\lambda^*(s)$ has a simple distribution, the Gibbs distribution (see Blume, 1993, 1995), which is given by

$$\lambda^*(s) = \frac{\exp \{\beta \rho(s)\}}{\sum_{s' \in S} \exp \{\beta \rho(s')\}}, \quad \beta \geq 0. \quad (16.29)$$

For any two distinct states, s and $\hat{s}$, the log of the likelihood ratio is given by

$$\log \frac{\lambda^*(s)}{\lambda^*(\hat{s})} = \beta [\rho(s) - \rho(\hat{s})].$$

Hence, in the long run, the process is more likely to be in that state whose stochastic potential is higher. Suppose that the only two states are s and $\hat{s}$. Then, for very large β , $\rho(s) > \rho(\hat{s})$ implies that in the long run, the process spends virtually all the time in state s . This intuition motivates the following result.

Proposition 16.33 (Blume, 1993, 1995): *Irrespective of the initial state, the long-run probability of being in state s is proportional to $\exp \{\beta \rho(s)\}$. Furthermore, for large enough β , there is a probability very close to one that the process is in the state which maximizes the potential $\rho(s)$.*

16.5.1 An application to agrarian contracts

Young and Burke (2001) use Blume's framework to explain the pattern of agricultural contracts in Illinois. In particular, they look at the pattern of *share contracts*. In a share contract, the total output is shared in some fixed, agreed upon, proportions between a tenant and a landlord. On any farm, assume that there is one landlord and one tenant. Let the tenant's share on farm i be s_i , so the landlord's share is $1 - s_i$. Denote the yield per acre on the farm to be y_i . So if the farm has a total area of h_i hectares, then the total output is $h_i y_i$. We normalize $h_i \equiv 1$.

Controlling for other agricultural inputs, yields depend primarily on things like soil productivity. In general, one would expect large variations in the tenant's share s_i across farms. However, one observes that in 98% of the cases, the tenant's share is either $\frac{1}{2}$, $\frac{2}{3}$, or $\frac{3}{5}$. In North Illinois, the predominant share is $\frac{1}{2}$, while in the South it is predominantly $\frac{2}{3}$ or $\frac{3}{5}$. The uniformity in any region occurs despite substantial differences in soil quality and productivity. Conversely, across very similar farms there can be large differences in s_i . Young and Burke (2001) provide an explanation.

Consider each farm as a vertex in a graph. All vertices are joined to their respective geographical neighbors. Along each edge, the social influence weight is assumed to be identical. So, for instance, if a farm has eight geographical neighbors, it is equally influenced by all eight neighbors. A tenant on farm i has a reservation wage w_i . In any time period, each farm gets a random contract revision opportunity. The arrival rate of contract revision opportunities on farms is independent and follows a Poisson distribution. When such an opportunity arises, the landlord makes a share offer s_i to the tenant, which must respect the participation constraint of the tenant, $s_i y_i \geq w_i$.

Omitting time superscripts, the monetary payoff of the landlord is given by

$$v_i(s_i) = (1 - s_i)y_i. \quad (16.30)$$

Recall that the externality is given by (16.25). Since each neighbor has exactly the same influence, so if there are n neighbors, $\omega_{ij} = 1/n$. Furthermore, define $u(s_i, s_j)$ to be the following indicator function δ_{ij} :

$$u(s_i, s_j) = \delta_{ij} = \begin{cases} 1 & \text{if } i, j \text{ are neighbors and } s_i = s_j \\ 0 & \text{otherwise} \end{cases}. \quad (16.31)$$

Motivated by (16.26), the total utility of the landlord is given by $[v_i(s_i) + \gamma \chi_i(s)] + \varepsilon_i$, where γ , referred to as a *conformity parameter*, is the relative weight of social influence on a farm, relative to its own monetary payoff. Denote by N_i , the set of neighbors of farm i . Using (16.30), (16.31), and $\omega_{ij} = 1/n$ we get that the utility function of the landlord in the i th farm is given by

$$V_i(s) = \left[(1 - s_i)y_i + \gamma \sum_{j \in N_i} \delta_{ij} \right] + \varepsilon_i. \quad (16.32)$$

If ε_i has an extreme value distribution that is independent across farms, then we can use Proposition 16.33 by defining an appropriate potential function $\rho(s)$.

$$\rho(s) = \sum_{i=1}^n (1 - s_i)y_i + \frac{\gamma}{2} \sum_i \sum_{j \in N_i} \delta_{ij},$$




Figure 16.16 Hypothetical distribution of farms.

Source: Young and Burke (2001) with permission from the American Economic Association.

where $\sum_{i=1}^n (1-s_i)y_i = R(s)$ is the total rent from land and $\frac{1}{2} \sum_i \sum_{j \in N_i} \delta_{ij} = e(s)$ is the *aggregate externality*. Hence, we can write $\rho(s) = R(s) + \gamma e(s)$. From Proposition 16.33, the main feature of the limiting distribution $\lambda^*(s)$ is that

$$\lambda^*(s) \propto \exp \beta [R(s) + \gamma e(s)]. \quad (16.33)$$

If we specify values for β , γ , y_i , and the reservation wage w_i , then we can find the long-run probability distribution of the states of the system. Suppose that there are soils of three qualities: high, medium and low. If there are m farms and each farm has eight geographical neighbors, then there is a total of $4m$ edges. These are shown in Figure 16.16. Each farm is denoted at a vertex by its soil quality, H (high), M (medium), and L (low). Each vertex also has eight directed edges pointing away from it, indicating the social influence of neighboring farms.

Consider the following numerical example. Suppose that $\gamma = 0$. Then, the long-run distribution is one that simply maximizes total rent $R(s)$. Suppose that: (1) Labor market conditions are competitive, and the reservation wage is given by $w_i = 32$. (2) Yields on the soils of the three qualities, high, medium, and low, are respectively 85, 70, 60. (3) Contract negotiations can be concluded only in multiples of 10, so $s_i = 0, 0.10, 0.20, \dots, 1.00$. Since tenants must be given their reservation utilities, the shares of tenants, s_i , on the three quality types, high, middle, and low, are 40%, 50%, and 60%.

In the spirit of models of stochastic social evolution, these long-run states in an ergodic system do not depend on the initial state. By varying the sizes of β and γ , Young and Burke (2001) explain the pattern of contracts in Illinois. In particular, for intermediate values of the conformity parameter, γ , the authors show that stochastic potential is maximized by a set of contracts such that there is *local uniformity*, but *global diversity*, i.e., inter-regional heterogeneity and intra-regional homogeneity, which matches the Illinois case.

APPENDICES

Appendix A: Ordinary differential equations

We review some elementary results in the theory of *ordinary differential equations* (ODE) in this section. There are many good sources (Coddington, 1961; Hirsch and Smale, 1974). A simple treatment for economists can be found in Simon and Blume (1994).

Suppose that we have a vector of real valued functions of time $t \in \mathbb{R}$

$$\mathbf{x}(t) = (x_1(t), x_2(t), \dots, x_n(t)) \in \mathbb{R}^n,$$

where $x_i(t) \in \mathbb{R}$ is a function of time that has continuous partial derivatives. $\mathbf{x}(t)$ can be thought of as a curve in an n -dimensional space. However, in many cases we do not know the exact curve, but rather the rate of change and direction of change of the curve in some region of $\mathbb{R}^n$.

Let a single dot over a variable represent the first time derivative, two dots represent the second time derivative, and so on. It is sometimes more convenient to denote the k th time derivative of some time dependent variable $x_i(t)$ by $\frac{d^k x_i(t)}{dt^k}$. The k th time derivative of the vector $\mathbf{x}(t)$ is to be understood as taking the k th time derivative, for each element of the vector, i.e., for any particular $t \in \mathbb{R}$

$$\frac{d^k \mathbf{x}(t)}{dt^k} = \left(\frac{d^k x_1(t)}{dt^k}, \frac{d^k x_2(t)}{dt^k}, \dots, \frac{d^k x_n(t)}{dt^k} \right) \in \mathbb{R}^n.$$

Definition 16.34 (*Differential equation*): A differential equation describes, at time t , the relation between a vector of n variables, and its time derivatives. This is written compactly as

$$g(t, \mathbf{x}(t), \dot{\mathbf{x}}(t), \ddot{\mathbf{x}}(t), \dots) = 0. \quad (5.A1)$$

The order of a differential equation is the highest time derivative in (5.A1).

We are primarily interested in *ordinary differential equations* (ODE) where the order $k = 1$. In contrast, if in addition to time, t , there was at least one more variable in (5.A1) with respect to which we differentiate, then we would have *partial* instead of *ordinary* differential equations.

Consider the following system of n first order differential equations.

$$\begin{aligned} \dot{x}_1(t) &= f_1(t, \mathbf{x}(t)), \\ \dot{x}_2(t) &= f_2(t, \mathbf{x}(t)), \\ &\vdots \\ \dot{x}_n(t) &= f_n(t, \mathbf{x}(t)), \end{aligned} \quad (5.A2)$$

where $f_i(t, \mathbf{x}(t)) : \mathbb{R}^{n+1} \rightarrow \mathbb{R}$. The system in (5.A2) is compactly written as

$$\dot{\mathbf{x}}(t) = \mathbf{f}(t, \mathbf{x}(t)), \quad (5.A3)$$

where $\dot{\mathbf{x}}(t)$ and $\mathbf{f}$ are column vectors whose i th elements are, respectively, $\dot{x}_i(t)$ and $f_i(t, \mathbf{x}(t))$. An *autonomous ODE* is one where time does not separately enter the function $\mathbf{f}$, i.e.,

$$\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t)). \quad (5.A4)$$

Conversely, the system of ODE in (5.A3) can be termed as *non-autonomous* or time dependent. The vector $\mathbf{x}(t)$ is called the *state* of the system. The set of all states is called the state space. We assume that the set of all time–state pairs $(t, \mathbf{x})$ forms a connected open subset, X , of $\mathbb{R}^{n+1}$. The function $\mathbf{f}$, which equals the column vector $(\dot{x}_1(t), \dot{x}_2(t), \dots, \dot{x}_n(t))$ at time t , is called a *vector field*. It specifies at any instant of time t , the *direction* and the *velocity* of the state of the system.

An ODE is said to be *linear* if the function $\mathbf{f}(\cdot)$ is linear. The majority of differential equations that we will consider are *linear, autonomous, first order differential equations*. In a *scalar* differential equation, $n = 1$, while $n \geq 2$ implies that we have a *system of differential equations*.

SOLUTION OF AN ODE

A *solution* to an ODE is a function $\mathbf{x}(t)$ that satisfies (5.A3). However, there is typically a whole family of solutions. For instance, consider the linear scalar autonomous first order differential equation, $\dot{x}(t) = ax$. This has the general solution $x^*(t, k) = ke^{at}$ where k is some constant; check that this is the solution by differentiating both sides of the solution with respect to t , to get the original ODE. Set $t = 0$ in the solution to get $x^*(0, k) = k$, hence, the solution is

$$x^*(t, k) = x^*(0, k)e^{at}.$$

More generally, we write a solution to the ordinary differential equation (5.A3) that passes through $\mathbf{x}$ at time 0 as:

$$\mathbf{x}^*(t, \mathbf{x}), \quad \text{where} \quad \mathbf{x}^*(0, \mathbf{x}) = \mathbf{x} \quad (5.A5)$$

$\mathbf{x} \times (0, \mathbf{x}) = \mathbf{x}$ is called the *initial condition*. We are interested in several aspects of the solution. Does it exist? Is it unique? Is it local or global? We begin with the last of these three questions.

Definition 16.35 (*Local and global solutions*): Consider the solution (5.A5) to the ordinary differential equation (5.A3), where $\mathbf{x}^*(t, \mathbf{x}) \in X$ for $t \in T$. If T is the entire real line, $\mathbb{R}$, then $\mathbf{x}^*(t, \mathbf{x})$ is a *global solution*. Otherwise, it is a *local solution*. T could be so short that the local solution defined on T might have little interest. On the other hand, if $T = [0, \text{infinity})$, then the local solution defined on T might be adequate for all practical purposes.

We will use the notation C^1 to indicate that a function has continuous first order partial derivatives. The next theorem provides the existence and uniqueness results for the solution to ODE.

Definition 16.36 (*Lipschitz continuity*): Let $f : X \rightarrow X$. Then $\mathbf{f}$ is said to be *Lipschitz continuous* if for all $\mathbf{x}, \mathbf{y} \in X$, there exists some $\lambda \in \mathbb{R}_+$ such that

$$d(f(\mathbf{x}), f(\mathbf{y})) \leq \lambda d(\mathbf{x}, \mathbf{y})$$

where $d(\mathbf{x}, \mathbf{y})$ is the distance between points $\mathbf{x}, \mathbf{y}$. An example is the Euclidean metric in $\mathbb{R}^n$:

$$d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\| = \sqrt{\sum_{i=1}^n (x_i - y_i)^2}.$$

Theorem 16.3 Consider the problem of solving the system of ODE in (5.A4), subject to the initial conditions in (5.A5). Suppose that each $f_i(\cdot)$, $i = 1, \dots, n$ is a continuous function at the point $(0, \mathbf{x}_i(0))$.

- (a) *Existence:* There exists some C^1 function $\mathbf{x}^*(t, \mathbf{x}(0))$ defined over some open time interval T containing $t = 0$ such that it satisfies (5.A4) for all $t \in T$ and $\mathbf{x}^*(0, \mathbf{x}(0)) = \mathbf{x}(0)$.
- (b) *Uniqueness:* Suppose in addition that each $f_i(\cdot)$, $i = 1, \dots, n$ is C^1 at the point $(0, \mathbf{x}_i(0))$, then $\mathbf{x}^*(t, \mathbf{x}(0))$ is unique.

In Theorem 16.3(b), the restriction on $f(\cdot)$ to be C^1 can be weakened to *Lipshitz continuity*. Any C^1 function is Lipshitz continuous and Lipshitz continuity in turn is stronger than the notion of continuity.

Consider a compact subset C of the state space X . Time is drawn from the real line $\mathbb{R}$. For some initial value $\mathbf{x}(0) \in C$, the solution to the ODE in (5.A4) is denoted by $\mathbf{x}^*(t, \mathbf{x}(0))$. This allows us to define a *dynamic system*.

Definition 16.37 (Dynamic system): A *dynamic system* is a pair $D = (X, \mathbf{x}^*)$, where $X \subset \mathbb{R}^{n+1}$ is open and connected with $(0, \mathbf{x}) \in X$ for some $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{x}^* : X \rightarrow \mathbb{R}^n$ is a *flow*, i.e., $\mathbf{x}^*$ satisfies:

1. *Initial value condition:* $\mathbf{x}^*(0, \mathbf{x}) = \mathbf{x}$ for all $\mathbf{x} \in \mathbb{R}^n$ such that $(0, \mathbf{x}) \in X$.
2. *Time additivity:* $\mathbf{x}^*(t, \mathbf{x}^*(s, \mathbf{x})) = \mathbf{x}^*(s + t, \mathbf{x})$ for all $s, t, \mathbf{x}$ such that $(s, \mathbf{x}), (s + t, \mathbf{x}), (t, \mathbf{x}^*(s, \mathbf{x})) \in X$.
3. *Differentiability:* $\mathbf{x}^*(t, \mathbf{x})$ is continuously differentiable in t for all $(t, \mathbf{x}) \in X$.

An important concept in gaining some geometric intuition about many problems, particularly those that cannot be explicitly solved, is the concept of a *trajectory*.

Definition 16.38 (Trajectory): A *trajectory*, $\varphi(s, \mathbf{x})$, is the graph of the solution that passes through $\mathbf{x}$ at time s , i.e.,

$$\varphi(s, \mathbf{x}) = \{(t, \mathbf{y}) \in X : \mathbf{y} = \mathbf{x}^*(t, \mathbf{x})\}.$$

We might only be interested in the path taken by the solution of a dynamic system, without reference to the time at which a particular state is reached. This leads to the concept of an *orbit*.

Definition 16.39 (Orbit): The *orbit*, $\gamma(s, \mathbf{x})$, of the trajectory, $\varphi(s, \mathbf{x})$, is the set of points (states) that the trajectory passes through, i.e.,

$$\gamma(s, \mathbf{x}) = \{\mathbf{y} : (t, \mathbf{y}) \in \varphi(s, \mathbf{x}) \text{ for some } t\}.$$

An important concept in dynamical systems is that of a *steady state*. One can define the steady state in terms of the induced solution mapping $\mathbf{x}^*(t, \mathbf{x}(0))$ as a state, $\bar{\mathbf{x}}$, such that

$$\mathbf{x}^*(t, \bar{\mathbf{x}}) = \bar{\mathbf{x}}.$$

Thus, in the steady state, if we begin at the state $\bar{\mathbf{x}}$, the dynamic system stays at $\bar{\mathbf{x}}$ over time, so evaluated at $\bar{\mathbf{x}}$, we have $\dot{\mathbf{x}}(t) = 0$.

It is not always possible to find an *explicit solution* to a differential equation. However, explicit solutions can be found in some special cases.

1. Suppose that

$$\dot{x}(t) = ax,$$

where a is some constant. We have already noted above, that the solution to this differential equation is $x^*(t, x(0)) = x(0)e^{at}$.

2. Suppose that

$$\dot{x}(t) = ax + b, \quad (5.A6)$$

where a, b are constants. The general solution is

$$x(t) = \frac{-b}{a} + ke^{at},$$

which can be differentiated to recover the differential equation in (5.A6). The steady state can be found by setting $\dot{x}(t) = 0$ in (5.A6) and solving out for $x = \bar{x}$, which gives $\bar{x} = \frac{-b}{a}$.

We also have $x(0) + \frac{b}{a} = k$ as the value of k .

Any differential equation of order 2 or above can be written as a system of first order differential equations. Consider, for instance, the second order differential equation.

$$\ddot{x}(t) = f(t, x(t), \dot{x}(t)) \quad (5.A7)$$

This can be converted into a system of two first order differential equations. Define

$$\dot{x}(t) = y(t). \quad (5.A8)$$

Then (5.A7) can be written as

$$\dot{y}(t) = f(t, x(t), y(t)). \quad (5.A9)$$

Equations (5.A8), (5.A9) give us a system of two first order ordinary differential equations. In general, any m th order differential equation can be converted into a system of m first order ordinary differential equations. Hence, there is no loss of generality in considering a system of first order ordinary differential equations.

SYSTEM OF LINEAR DIFFERENTIAL EQUATIONS

Often we would like to focus on a system of n linear first order differential equations, expressed in the compact matrix notation

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t), \quad (5.A10)$$

where $\mathbf{x}(t)$ is a $n \times 1$ column vector with the i th element given by $x_i(t)$. Similarly $\dot{\mathbf{x}}(t)$ is a $n \times 1$ column vector with i th element given by $\dot{x}_i(t)$. A is a $n \times n$ matrix with typical element given by a_{ij} . Denote the eigenvalues of the matrix A by $r_1, r_2, \dots, r_n$ and the eigenvectors by $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$. Then the solution to (5.A10) for various cases is given by the following theorems. We omit the proofs; the reader can find these in Hirsch and Smale (1974).

Theorem 16.4 (*Distinct real eigenvalues*) Suppose that $r_1, r_2, \dots, r_n$ are distinct and real. Then the solution to (5.A10) is given by

$$\mathbf{x}(t) = \sum_{i=1}^n c_i e^{r_i t} \mathbf{v}_i. \quad (5.A11)$$

The constants $c_1, \dots, c_n$ are determined from n initial conditions, such as the one given in (5.A5).

Theorem 16.4 is valid even if we have multiple eigenvalues, so long as any eigenvalue with multiplicity k has k linearly independent eigenvectors. If not, then we have the following case, which we give for the case of a 2×2 matrix A .

Theorem 16.5 (*Multiple eigenvalues, $n = 2$*) Suppose that the two eigenvalues of the 2×2 matrix A are equal, i.e., $r_1 = r_2 = r$. Let $\mathbf{v}$ be the one independent eigenvector corresponding to the eigenvalue, r . Let $\mathbf{u}$ be the generalized eigenvector of A .¹ Then the general solution to (5.A10) is given by

$$\mathbf{x}(t) = e^{rt} (c_1 \mathbf{v} + c_2 \mathbf{u}) + t e^{rt} c_2 \mathbf{v}, \quad (5.A12)$$

where the two constants c_1, c_2 are determined by using the initial conditions.

Theorem 16.6 (*Complex eigenvalues, $n = 2$*) Suppose that the two eigenvalues of the 2×2 matrix A are complex, i.e., $\alpha \pm i\beta$, where $\alpha, \beta \in \mathbb{R}$. The corresponding eigenvectors are also complex, i.e., $\mathbf{u} \pm i\mathbf{v}$. Then the general solution to (5.A10) is given by

$$\mathbf{x}(t) = e^{\alpha t} \cos \beta t (c_1 \mathbf{u} - c_2 \mathbf{v}) - e^{\alpha t} \sin \beta t (c_2 \mathbf{u} + c_1 \mathbf{v}), \quad (5.A13)$$

and the two constants c_1, c_2 are determined by using the initial conditions.

STEADY STATES AND STABILITY

Consider the system of n ODE in (5.A3), $\dot{\mathbf{x}}(t) = \mathbf{f}(t, \mathbf{x}(t))$. The *steady state*, also known as a *stationary point* or an *equilibrium*, of an ODE, is a state $\bar{\mathbf{x}}$, that satisfies

$$\mathbf{f}(\bar{\mathbf{x}}, t) = 0.$$

Thus, in a steady state, $\dot{\mathbf{x}}(t) = 0$.

Definition 16.40 (*Liapunov stability or neutral stability*): Let $\mathbf{x}^*(t, \mathbf{x}(0))$ be the solution to (5.A3). A steady state $\bar{\mathbf{x}}$ is *Liapunov stable* if for some neighborhood U of $\bar{\mathbf{x}}$, there exists a neighborhood V of $\bar{\mathbf{x}}$ such that for all $t \geq 0$, and for any $\mathbf{x}(0) \in V$, $\mathbf{x}^*(t, \mathbf{x}(0))$ stays in U .

¹ The generalized eigenvector, $\mathbf{u}$, solves $(A - rI)\mathbf{u} = \mathbf{v}$, where I is the 2×2 identity matrix and $\mathbf{v}$ is the one independent eigenvector corresponding to the eigenvalue, r .

Definition 16.41 (Dynamic or asymptotic stability): Let U, V be as defined in Definition 16.40. A steady state $\bar{x}$ is asymptotically stable if (i) it is Liapunov stable, and (ii) for any $x(0) \in V$

$$\lim_{t \rightarrow \infty} x^*(t, x(0)) = \bar{x}. \quad (5.A14)$$

Definition 16.42 (Unstable steady states): If $\bar{x}$ is neither asymptotically stable nor neutrally stable, then it is said to be unstable.

The two concepts of stability and asymptotic stability are shown in Figure 5.A1, which is self-explanatory in light of the definitions. Asymptotic stability is the more stringent of the two stability conditions, but it does not necessarily imply Liapunov stability. The intuition is that a trajectory might have to stray outside the set U before it eventually converges to $\bar{x}$, but that would contradict Liapunov stability.

Another useful concept is that of a *basin of attraction* and an *attractor*. Heuristically, a basin of attraction of some set A is the set of all initial states such that the trajectory of the dynamic system stays in A for large enough values of t .

Definition 16.43 (Basin of attraction): Let X be a compact set of states and let $x \in X$. The basin of attraction of a closed set $A \subset X$ is the following set:

$$\{x(0) \in X : \lim_{t \rightarrow \infty} x^*(t, x(0)) \in A\}.$$

Definition 16.44 (Attractors) A set $A \subset X$ is an attractor if its basin of attraction is a neighborhood of A .

The analysis of stability is substantially simplified by examining the size and sign of the eigenvalues corresponding to the matrix A . We will use the solutions for the various cases in (5.A11), (5.A12), (5.A13). We only state the results that we need for this appendix, some for the 2×2 case; the interested reader can consult Hirsch and Smale (1974) for the more general case.

Theorem 16.7 (Stability of linear systems): Consider the system of n linear first order differential equations in (5.A10), $\dot{x}(t) = Ax(t)$.




Figure 5.A1 Stability of steady states.

- (a) Suppose that all eigenvalues $r_1, r_2, \dots, r_n$ are distinct and real. If each $r_i < 0$, $i = 1, 2, \dots, n$, then the steady state $\bar{x} = 0$ is globally asymptotically stable.
- (b) Suppose that in the 2×2 case, the eigenvalues are complex, i.e. $\alpha \pm i\beta$. If $\alpha < 0$, then $\bar{x} = 0$ is globally asymptotically stable. If $\alpha = 0$, i.e., the eigenvalues are imaginary numbers and a complete set of independent eigenvectors does not exist, then $\bar{x} = 0$ is asymptotically unstable.
- (c) Suppose that in the 2×2 case, the eigenvalues are equal, i.e., $r_1 = r_2 = r$. Then if $r < 0$, $\bar{x} = 0$ is asymptotically stable. If $r \geq 0$, then $\bar{x} = 0$ is asymptotically unstable.

Now let us focus on the stability of non-linear systems. Consider an ODE in one variable.

$$\dot{x}(t) = f(x). \quad (5.A15)$$

Suppose that the graph of this function is given in Figure 5.A2. In the Figure, the steady state is shown at $\bar{x}$. When $x > \bar{x}$, then $f(x) < 0$, so $\dot{x}(t) < 0$, thus, x decreases. On the other hand, when $x < \bar{x}$, then $f(x) > 0$, so $\dot{x}(t) > 0$, thus, x increases. In other words, any trajectory of the system converges to $\bar{x}$ as $t \rightarrow \infty$. Hence, $\bar{x}$ is asymptotically stable. Notice that $f'(\bar{x}) < 0$; this turns out to be a crucial condition for asymptotic stability.

Theorem 16.8 (Stability of non-linear systems, $n = 1$): Let $f(x)$ be a C^1 function and let $\bar{x}$ be a steady state of the system, so $f(\bar{x}) = 0$. If $f'(\bar{x}) < 0$, then $\bar{x}$ is an asymptotically stable equilibrium. If $f'(\bar{x}) > 0$, then $\bar{x}$ is an asymptotically unstable equilibrium.

Suppose that we have a system of $n > 1$ non-linear differential equations, $\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t))$. The extension of Theorem 16.8 to this case is straightforward with the derivative $f'(\bar{x})$ replaced by the Jacobian matrix of partial derivatives, $D\mathbf{f}(\bar{\mathbf{x}})$.

$$D\mathbf{f}(\bar{\mathbf{x}}) = \begin{bmatrix} \frac{\partial f_1(\bar{\mathbf{x}})}{\partial x_1} & \frac{\partial f_1(\bar{\mathbf{x}})}{\partial x_2} & \cdot & \cdot & \frac{\partial f_1(\bar{\mathbf{x}})}{\partial x_n} \\ \frac{\partial f_2(\bar{\mathbf{x}})}{\partial x_1} & \frac{\partial f_2(\bar{\mathbf{x}})}{\partial x_2} & \cdot & \cdot & \frac{\partial f_2(\bar{\mathbf{x}})}{\partial x_n} \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \frac{\partial f_n(\bar{\mathbf{x}})}{\partial x_1} & \frac{\partial f_n(\bar{\mathbf{x}})}{\partial x_2} & \cdot & \cdot & \frac{\partial f_n(\bar{\mathbf{x}})}{\partial x_n} \end{bmatrix}$$




Figure 5.A2 Stability of the steady state in a first order non-linear differential equation.

Theorem 16.9 (Stability of non-linear systems, $n > 1$): Consider the system of $n > 1$ non-linear autonomous differential equations, $\dot{\mathbf{x}}(t) = \mathbf{f}(\mathbf{x}(t))$, where $\mathbf{f}(\cdot)$ is C^1 . Denote the steady state of the system by $\bar{\mathbf{x}}$.

- (a) If all eigenvalues of the matrix $D\mathbf{f}(\bar{\mathbf{x}})$ are negative, or have a negative real part in the case of complex numbers, then $\bar{\mathbf{x}}$ is an asymptotically stable steady state.
- (b) If at least one eigenvalue of $D\mathbf{f}(\bar{\mathbf{x}})$ is positive, or has a positive real part in the case of complex numbers, then $\bar{\mathbf{x}}$ is an asymptotically unstable steady state.
- (c) If some of the eigenvalues are zero, or some have purely imaginary parts in the case of complex numbers, then the Jacobian matrix cannot be used to determine the stability of the steady state. In this case more information is needed about the state of the system.

In many cases, we cannot find the solution to an ODE. In this case, we have to construct a *phase diagram* of the system. Essentially, this is a plot of the vector field of the system (described above). Often it helps to plot the vector field by fixing the value of $\dot{\mathbf{x}}$, e.g., $\dot{\mathbf{x}} = 0$. The vector field corresponding to a particular value of $\dot{\mathbf{x}}$ is called an *isocline*. The stability of the system can then be inferred from a plot of the isoclines for $\dot{\mathbf{x}} = 0$ and the point of intersection of the isoclines.

Appendix B: A note on choice under stochastic utility

In most learning models that we shall consider, individuals choose among their pure strategies in a probabilistic manner that is consistent with either the conditional logit model, or the multinomial logit model.² The foundation for this behavior lies in models of stochastic utility, as we now explain.

Suppose that $N = \{1, \dots, n\}$ is the set of players. The pure strategy set of player $i \in N$ is given by $S_i = \{s_{i1}, s_{i2}, \dots, s_{in_i}\}$ and the pure strategy profile is given by $s = (s_1, s_2, \dots, s_n) \in S = S_1 \times \dots \times S_n$. From the perspective of player i we may write the strategy profile as (s_i, s_{-i}) where s_{-i} is the strategy profile of all players other than player i . Let ε_i be a random variable, specific to player i . The utility function of player i is given by $u_i : S \times \varepsilon_i \rightarrow \mathbb{R}$. In particular suppose that

$$u_i(s_i, s_{-i}, \varepsilon_i) = v_i(s_i, s_{-i}) + \varepsilon_i(s_i, s_{-i}). \quad (5.A16)$$

This is the general formulation of a stochastic utility model. Suppose that any two random variables ε_i and ε_j for two different players are independent. We assume that ε_i is an Extreme Value Type I random variable (sometimes referred to as a Weibull distribution). The distribution function of ε_i is given by

$$F(c) = \Pr(\varepsilon_i < c) = e^{-e^{-(c+\alpha_i)}}, \quad (5.A17)$$

where α_i is the single parameter of the distribution, sometimes also called the *location parameter*. The next proposition gives some important results about Extreme Value Type I distribution as well as about the utility function, u_i , in (5.A16).

Proposition 16.34

- (a) The density of an Extreme Value Type I random variable, ε_i , is given by

$$f(\varepsilon_i) = e^{-(\varepsilon_i + \alpha_i)} \left(e^{-e^{-(\varepsilon_i + \alpha_i)}} \right).$$

² See Amemiya (1981), McFadden (1981), and Maddala (1986) for this framework.

- (b) *Independently distributed Extreme Value Type I random variables, $\varepsilon_1, \dots, \varepsilon_n$ are closed under maximization, i.e., the distribution function $\Pr(\max_i \{\varepsilon_i\} \leq \varepsilon)$ belongs to an Extreme Value Type I random variable.*
- (c) *u_i defined in (5.A16) is an Extreme Value Type I random variable.*

Proof: (a) From (5.A17), $F(\varepsilon_i) = e^{-e^{-(\varepsilon_i + \alpha_i)}}$, hence, the density is given by $f(\varepsilon_i) = \frac{\partial F(\varepsilon_i)}{\partial \varepsilon_i}$. Let $x = e^{-e^{-(\varepsilon_i + \alpha_i)}}$, so $\log x = -e^{-(\varepsilon_i + \alpha_i)}$ and $\frac{\partial x}{\partial \varepsilon_i} = x \frac{\partial}{\partial \varepsilon_i} (-e^{-(\varepsilon_i + \alpha_i)}) = x e^{-(\varepsilon_i + \alpha_i)}$, thus,

$$f(\varepsilon_i) = e^{-(\varepsilon_i + \alpha_i)} \left(e^{-e^{-(\varepsilon_i + \alpha_i)}} \right). \quad (5.A18)$$

(b) Suppose that $\varepsilon_1, \dots, \varepsilon_n$ are independently distributed Extreme Value Type I random variables. Then

$$\begin{aligned} \Pr(\max_i \{\varepsilon_i\} \leq \varepsilon) &= \prod_{i=1}^n \Pr(\varepsilon_i \leq \varepsilon) \text{ [By definition]} \\ &= \prod_{i=1}^n e^{-e^{-(\varepsilon + \alpha_i)}} \text{ [Using (5.A17)]} \\ &= e^{-e^{-\varepsilon} \sum_{i=1}^n e^{-\alpha_i}} \text{ [By simplification]}. \end{aligned} \quad (5.A19)$$

Let $\sum_{i=1}^n e^{-\alpha_i} = e^{-\alpha}$. Taking logs on both sides, we get $-\alpha = \log \sum_{i=1}^n e^{-\alpha_i}$. Substituting in (5.A19) we get

$$\Pr(\max_i \{\varepsilon_i\} \leq \varepsilon) = e^{-e^{-\varepsilon} e^{-\alpha}} = e^{-e^{-(\varepsilon + \alpha)}}. \quad (5.A20)$$

Comparing (5.A17) and (5.A20) it follows that the random variable $\max_i \{\varepsilon_i\}$ also has an Extreme Value Type I distribution.

(c) Denote the distribution function of u_i by F_{u_i} , then

$$\begin{aligned} F_{u_i}(\varepsilon) &= \Pr(u_i \leq \varepsilon) \\ &= \Pr(v_i + \varepsilon_i \leq \varepsilon) \text{ [Using (5.A16)]} \\ &= \Pr(\varepsilon_i \leq \varepsilon - v_i) \\ &= e^{-e^{-(\varepsilon - v_i + \alpha_i)}} \text{ [Using (5.A17)]}. \end{aligned} \quad (5.A21)$$

Thus, u_i has an Extreme Value Type I distribution. ■

We can now use Proposition 16.34 to show that player $i \in N$ chooses his pure strategies with probabilities using the conditional logit, or multinomial model.

Suppose first that $S_i = \{s_{i1}, s_{i2}\}$. When does player i choose his first strategy over the second? Since the RHS of (5.A16) is a random variable, we can compute the probability of this event. Denote $u_i(s_{ij}, s_{-i}, \varepsilon_i)$ by u_{ij} , $v_i(s_{ij}, s_{-i})$ by v_{ij} , $\varepsilon_i(s_{ij}, s_{-i})$ by ε_{ij} , and $\Pr(s_{i1} \text{ is chosen})$ by P_{i1} , then

$$\begin{aligned} P_{i1} &= \Pr(u_{i1} > u_{i2}) \\ &= \Pr(v_{i1} + \varepsilon_{i1} > v_{i2} + \varepsilon_{i2}) \\ &= \Pr(\varepsilon_{i2} < v_{i1} - v_{i2} + \varepsilon_{i1}). \end{aligned} \quad (5.A22)$$

Since ε_{ij} has an Extreme Value Type I distribution, thus

$$\Pr(\varepsilon_{i2} < v_{i1} - v_{i2} + \varepsilon_{i1}) = \int_{-\infty}^{\infty} f(\varepsilon_{i1}) \left[\int_{-\infty}^{v_{i1} - v_{i2} + \varepsilon_{i1}} f(\varepsilon_{i2}) d\varepsilon_{i2} \right] d\varepsilon_{i1}. \quad (5.A23)$$

Using (5.A21), (5.A22), (5.A23) and we get

$$\Pr(s_{i1} \text{ is chosen}) = \int_{-\infty}^{\infty} e^{-e^{-(v_{i1} - v_{i2} + \varepsilon_{i1} + \alpha_2)}} f(\varepsilon_{i1}) d\varepsilon_{i1}. \quad (5.A24)$$

Substituting the value of f from (5.A18), $e^{-(\varepsilon_i + \alpha_i)} \left(e^{-e^{-(\varepsilon_i + \alpha_i)}} \right)$, into (5.A24)

$$\begin{aligned} P_{i1} &= \int_{-\infty}^{\infty} e^{-e^{-(v_{i1} - v_{i2} + \varepsilon_{i1} + \alpha_2)}} e^{-(\varepsilon_{i1} + \alpha_1)} \left(e^{-e^{-(\varepsilon_{i1} + \alpha_1)}} \right) d\varepsilon_{i1} \\ &= e^{-\alpha_1} \int_{-\infty}^{\infty} e^{-\varepsilon_{i1}} e^{-e^{-(v_{i1} - v_{i2} + \varepsilon_{i1} + \alpha_2)}} \left(e^{-e^{-(\varepsilon_{i1} + \alpha_1)}} \right) d\varepsilon_{i1} \\ &= e^{-\alpha_1} \int_{-\infty}^{\infty} e^{-\varepsilon_{i1}} e^{-(e^{-\varepsilon_{i1}})(e^{-\alpha_1 - e^{-(v_{i1} - v_{i2} + \alpha_2)}})} d\varepsilon_{i1}. \end{aligned}$$

Let $\mu = (e^{-\alpha_1} - e^{-(v_{i1} - v_{i2} + \alpha_2)})$, which is independent of ε_{i1} , then

$$P_{i1} = \frac{e^{-\alpha_1}}{\mu} \int_{-\infty}^{\infty} e^{-(e^{-\varepsilon_{i1}})\mu} (e^{-\varepsilon_{i1}}\mu) d\varepsilon_{i1}. \quad (5.A25)$$

Using the result that $\int_{-\infty}^{\infty} e^{g(x)} g'(x) dx = |e^{g(x)}|_{-\infty}^{\infty}$, we can rewrite (5.A25) as

$$P_{i1} = \frac{e^{-\alpha_1}}{\mu} \left| e^{-(e^{-\varepsilon_{i1}})\mu} \right|_{-\infty}^{\infty}.$$

Simplifying this expression gives

$$P_{i1} = \frac{e^{-\alpha_1}}{e^{-\alpha_1} + e^{-(v_{i1} - v_{i2} + \alpha_2)}}.$$

Multiplying and dividing the denominator by $e^{v_{i1}}$ we get the desired result

$$P_{i1} = \frac{e^{v_{i1} - \alpha_1}}{e^{v_{i1} - \alpha_1} + e^{v_{i2} - \alpha_2}}.$$

Such a probabilistic choice is called a conditional or multinomial logit model.

We can now use Proposition 16.34(b) to generalize this result to any finite strategy set $S_i = \{s_{i1}, s_{i2}, \dots, s_{in_i}\}$.

$$\begin{aligned} P_{i1} &= \Pr(u_{i1} > u_{i2}, u_{i1} > u_{i3}, \dots, u_{i1} > u_{in_i}) \\ &= \Pr\left(u_{i1} > \max_{2 \leq j \leq n_i} u_{ij}\right) \\ &= \Pr\left(v_{i1} + \varepsilon_{i1} > \max_{2 \leq j \leq n_i} (v_{ij} + \varepsilon_{ij})\right) \\ &= \frac{e^{v_{i1} - \alpha_1}}{e^{v_{i1} - \alpha_1} + e^{v_{i2} - \alpha_2} + \dots + e^{v_{in_i} - \alpha_{n_i}}}, \end{aligned} \quad (5.A26)$$

which is the result we were seeking to prove.

A special case arises by setting the location parameter $\alpha_1 = 0$. This case is the most commonly used in applied experimental work. In this case we can rewrite (5.A26) as

$$P_{ij} = \frac{e^{v_{ij}}}{\sum_{k=1}^{n_i} e^{v_{ik}}}.$$

Appendix C: A brief user's guide to Markov Chains

This section gives a brief, informal, cookbook guide to using Markov chains for Part 5 of the book. Readers interested in an introduction to the field may consult Karlin and Taylor (1998).

Suppose that time t is discrete, so $t = 0, 1, 2, \dots$. A discrete *random variable* x_t at time t takes possible realizations from a finite set, $S = \{s_1, s_2, \dots, s_n\}$. The set S is said to be the *state-space*.

Definition 16.45 (*Stochastic process*): A stochastic process $\{x_t\}_{t=0}^{\infty}$ is a sequence of random variables.

Definition 16.46 (*Markov property*): A stochastic process is said to have the Markov property if³

$$p(x_{t+k} \mid x_{t+k-1}, x_{t+k-2}, \dots, x_0) = p(x_{t+k} \mid x_{t+k-1}); \forall t, k \geq 2.$$

Thus, under the Markov property, in predicting the future, only the immediate past matters. Denote by $\Delta(S)$ the set of all probability vectors on the state space S . The typical probability vector at time t is denoted by the row vector $\lambda_t = [\lambda_t(s_1), \lambda_t(s_2), \dots, \lambda_t(s_n)]$, where $\lambda_t(s_i)$ is the probability that the random variable x_t takes the value s_i at time t , $s_i \in S$.

Definition 16.47 (*Transition matrix*): The transition matrix, P , is given by

$$P = \begin{bmatrix} p_{11} & p_{12} & \cdot & \cdot & p_{1n} \\ p_{21} & p_{22} & \cdot & \cdot & p_{2n} \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ p_{n1} & p_{n2} & \cdot & \cdot & p_{nn} \end{bmatrix}$$

The typical element of the transition matrix is denoted by

$$p_{ij} = P(s_i, s_j); i, j = 1, \dots, n,$$

where p_{ij} denotes the probability of being in state $s_j \in S$ at time $t+1$, conditional on being in state $s_i \in S$ at time t . For any row in P , the sum of all entries in any row is 1, i.e.,

$$\sum_{j=1}^n p_{ij} = 1; \forall i = 1, \dots, n.$$

For instance, starting from the state s_1 at time t , the first row of P gives the probability of transiting to all the n states in S at time $t+1$. The typical element in the first row, p_{1j} gives the probability of transiting to state $s_j \in S$ at time $t+1$. The transition matrix is a *stochastic matrix* because all entries are non-negative and all rows sum to 1.

³ We use the standard notation, $p(x \mid y)$, for the probability of x conditional on y .

Definition 16.48 (*Markov chain*): A Markov chain is given by the pair (S, P) , where S is the set of states, and P is the transition matrix, given the initial probability vector λ_0 , such that $\sum_{j=1}^n \lambda_0(s_j) = 1$. We will focus on finite Markov chains in which there is a finite number of elements in S .

The unconditional probability vector at time t , λ_t , can be found as follows

$$\lambda_t = \lambda_{t-1}P. \quad (5.A27)$$

From (5.A27), the probability of being in state s_i at time $t + 1$ is given by

$$\lambda_t(s_i) = \sum_{j=1}^n \lambda_{t-1}(s_j) p_{ji}. \quad (5.A28)$$

Iteratively substituting backwards, using (5.A27), we get $\lambda_t = (\lambda_{t-2}P)P = \dots = \lambda_0P^t$. Given

$$\lambda_t = \lambda_0P^t, \quad (5.A29)$$

we are usually interested in the limiting unconditional probability distribution

$$\bar{\lambda} = \lim_{t \rightarrow \infty} \lambda_0P^t. \quad (5.A30)$$

The limit in (5.A30) exists under some mild regularity conditions; see Karlin and Taylor (1998). If probabilities converge to the limiting distribution $\bar{\lambda}$, (5.A27) implies that $\bar{\lambda} = \bar{\lambda}P$, so

$$\bar{\lambda}(I - P) = 0. \quad (5.A31)$$

where $I, 0$ are, respectively, the identity matrix and the row vector of zeros. Taking the transpose on both sides of (5.A31) we get that:

$$(I - P^T)\bar{\lambda}^T = 0.$$

Hence, $\bar{\lambda}^T$ is the eigenvector corresponding to the matrix P^T whose columns each sum to unity. Since P is a stochastic matrix (all elements positive and rows summing to 1) there is at least one unit eigenvalue. This is easily seen by noting that $Pe = e$ where e is the unit column vector.

A question of particular interest is: does $\bar{\lambda}$ depends on the initial probability vector λ_0 ?

Definition 16.49 (*Unique invariant distribution*) If $\bar{\lambda}$ does not depend on the initial probability vector λ_0 , then it is said to be a unique invariant distribution of the underlying stochastic process (S, P) .

Definition 16.50 (*Asymptotic stationarity*): If the stochastic process (S, P) has a unique invariant distribution, it is said to be asymptotically stationary.

An example where a unique invariant distribution does not exist is when $P = I$, the identity matrix. In this case, $\bar{\lambda} = \lambda_0$, so there is a clear dependence of the limiting distribution on the initial conditions. On the other hand, suppose that $P = [0, 0, \dots, 1, \dots, 0]$, i.e., the i th column has all entries equal to 1, while all other entries in the matrix are 0. In this case, using (5.A28),

$\bar{\lambda} = (0, 0, \dots, 1, \dots, 0)$ for any initial probability vector λ_0 , whose entries sum up to 1. Hence, $\bar{\lambda}$ is a *unique invariant distribution*.

Definition 16.51 (Ergodicity): When the limit in (5.A30) exists and $\bar{\lambda}$ is independent of λ_0 , then the stochastic process (S, P) is said to be *ergodic*.

Definition 16.52 (Empirical frequency of play): Suppose that we begin with some state $s_j \in S$. Suppose we let the system evolve for t time periods and note down the frequency with which each of the n states, $s_1, s_2, \dots, s_n$, is visited. Then the empirical frequency with which the i th state is visited (in t time periods), conditional on the initial state being s_j is denoted by $f^t(s_i | s_j)$.

Theorem 16.10 (Convergence of the empirical frequency): Suppose that the stochastic process (S, P) is ergodic and the unique invariant distribution is given by $\bar{\lambda}$. Then, the empirical frequency distribution converges almost surely to $\bar{\lambda}$, independently of the initial state $s_j \in S$, i.e., for all $s_i \in S$

$$\lim_{t \rightarrow \infty} f^t(s_i | s_j) = \bar{\lambda}(s_i),$$

where $\bar{\lambda}(s_i)$ is the probability of being in state s_i in the unique invariant distribution.

Theorem 16.11 The stochastic process (S, P) is asymptotically stationary with a unique invariant distribution if $p_{ij} > 0$ for all $i, j = 1, 2, \dots, n$.

For the case of a stochastic matrix P^t with typical element p_{ij}^t (where t denotes time) this theorem requires that $p_{ij}^t > 0$ for all $i, j = 1, 2, \dots, n$ and for some $t \geq 1$.

Definition 16.53 (Basin of attraction): The basin of attraction of a state s_i , denoted by $B(s_i)$, is the set of all states such that there is a positive probability of moving from those states to the state s_i in a finite number of periods.

We now provide an alternative motivation of ergodicity, which will be useful in the analysis of perturbed Markov chains.

Definition 16.54 (Accessibility): State s_i is “accessible” from state s_j if starting from state s_j there is a positive probability of moving to state s_i in a finite number of periods. We write this as $s_j \rightarrow s_i$. The symbol $\nrightarrow$ denotes “not accessible.”

Definition 16.55 (Communication): If $s_j \rightarrow s_i$ and $s_i \rightarrow s_j$, then both states are accessible from each other. In this case, we say that states s_i, s_j “communicate” and we write this as $s_i \sim s_j$. A communication class is a set such that for any s_i, s_j in the class, we have $s_i \sim s_j$.

Being an equivalence relation, $\sim$, partitions the state space, S , into equivalence classes, or communication classes.⁴

⁴ Suppose that the equivalence relation, $\sim$, is defined over some set A and $a, b, c \in A$. Then the equivalence relation is reflexive ($a \sim a$), symmetric ($a \sim b \Leftrightarrow b \sim a$), and transitive ($a \sim b, b \sim c \Rightarrow a \sim c$). For example, let the set A be the set of birthdays for each individual in a city. Let $\sim$ be the binary relation “has the same birthday as.” Then the equivalence relation partitions the set A into disjoint equivalence classes and each class has all those individuals who share their birthday.

Definition 16.56 (Recurrence classes): A communication class is said to be a recurrence class if for any s_j in that class, $s_j \rightarrow s_i$, for any s_i outside that class.

Recalling the definition of accessibility, a recurrence class only requires that starting within it we cannot reach a state outside it in finite time. Furthermore, since the recurrence class is a communication class we can, in finite time, move between elements of the recurrence class.

Definition 16.57 (Recurrent and transient states): A state is said to be a recurrent state if it belongs to some recurrent class; otherwise it is said to be a transient state.

Definition 16.58 (Absorbing state): If a recurrent class has a single state in it, then the single state is called an absorbing state. In other words, for an absorbing state s_i , $p_{ii} = 1$.

Definition 16.59 (Irreducible Markov process): If a Markov process has a single recurrent class and it constitutes the entire state space, then the Markov process is said to be irreducible. Alternatively, a Markov process is irreducible if for the entire state space, there is a positive probability of moving from any state to any other in finite time.

Theorem 16.12 (Ergodicity): A finite Markov chain (S, P) is ergodic if it has a unique recurrence class.

Suppose that there is not a unique recurrent class. Then, there are several limiting distributions, each depending not only on the initial conditions (non-ergodic), but also on what happens along the path.

Theorem 16.13 (Alternative method for checking ergodicity): Consider a finite Markov chain (S, P) . Suppose that we can partition the set of states S into two disjoint subsets S_1, S_2 , such that the following hold.

- (1) There is a positive probability of moving in finite time from (a) a state in S_1 to a state in S_2 , (b) a state in S_2 to a state in S_1 .
- (2) There is a zero probability of moving in finite time from a state in S_2 to a state in S_1 .
Then ergodicity holds.

Definition 16.60 (Aperiodic Markov chain): Choose some $s_i \in S$. Let N_i be the set of positive integers $\{1, 2, \dots, n_i\}$ such that there is a positive probability that $s_i \rightarrow s_i$ in exactly n_i periods. If the greatest common divisor of N_i is unity for all $s_i \in S$, then the Markov chain is said to be aperiodic.

The next theorem makes even sharper predictions about asymptotic behavior.

Theorem 16.14 (Average and probabilistic behavior): Suppose that a finite Markov chain is aperiodic and ergodic (i.e., has a unique invariant distribution, $\bar{\lambda}$). Then,

- (i) The frequency of play of the pure strategies (or the average behavior) converges to $\bar{\lambda}$, i.e.,

$$\lim_{t \rightarrow \infty} f^t(s_i | s_j) = \bar{\lambda}(s_i); \forall s_i, s_j \in S.$$

- (ii) The probabilistic behavior also converges to $\bar{\lambda}$, i.e.,

$$\lim_{t \rightarrow \infty} p_{ij}^t = \bar{\lambda}(s_j),$$

where p_{ij}^t is an element of the matrix P^t .

PERTURBED MARKOV CHAINS

We now consider the basics of *perturbed Markov chains*; see, for instance, Kandori et al. (1993), Young (1993a, 1998), and Ellison (2000). Consider a family of Markov chains indexed by the parameter ε and denoted by $(S, P_\varepsilon)_{\varepsilon \in [0,1]}$. As before, S contains a finite number of states.

Definition 16.61 For each $\varepsilon \in [0,1]$, P_ε is a transition matrix and it is called the *perturbed transition matrix*. For $\varepsilon = 0$, P_0 is called the *unperturbed transition matrix*. All elements of P_ε are continuous in ε for all $\varepsilon \in [0,1]$.

For each $\varepsilon \in (0,1)$, the stochastic process (S, P_ε) is assumed to be ergodic with a unique invariant distribution given by $\bar{\lambda}_\varepsilon$. We do not assume ergodicity for the unperturbed transition matrix, P_0 .

Definition 16.62 (*Resistance or transition cost function*): The “resistance” or “transition cost function,” $c(s_i, s_j)$, captures the difficulty of moving from state s_i to s_j in finite time. One may think of $c(s_i, s_j)$ as the minimum number of steps required to make the transition from state s_i to state s_j .

Definition 16.63 (*Regular perturbed Markov process*): The Markov chain (S, P_ε) is said to be a regular perturbed Markov process if it satisfies two conditions.

- (i) P_ε is ergodic for all small enough $\varepsilon > 0$.
- (ii) The rate at which P_ε approaches P_0 is exponentially smooth. This requires two conditions. First, there is an element by element convergence of the perturbed matrix to the unperturbed matrix, as the noise in the system dies out, i.e.,

$$\lim_{\varepsilon \rightarrow 0^+} P_\varepsilon(s_i, s_j) = P_0(s_i, s_j); i, j = 1, 2, \dots, n.$$

Second, it requires that if for some $\varepsilon > 0$

$$P_\varepsilon(s_i, s_j) > 0, \text{ then } 0 < \lim_{\varepsilon \rightarrow 0^+} \frac{P_\varepsilon(s_i, s_j)}{\varepsilon^{c(s_i, s_j)}} < \infty.$$

For regular perturbed Markov processes, $\lim_{\varepsilon \rightarrow 0} \bar{\lambda}_\varepsilon$ always exists.

Definition 16.64 (*Limit invariant distribution*): Define the limit invariant distribution, λ^* , as one when the noise, ε , in the dynamical system goes to zero, i.e.,

$$\lambda^* = \lim_{\varepsilon \rightarrow 0} \bar{\lambda}_\varepsilon. \quad (5.A32)$$

Definition 16.65 (*Stochastically stable states*): The set of stochastically stable states, $S^* \subset S$, are those which are in the support of the limit invariant distribution, i.e.,

$$S^* = \{s \in S : \lambda^*(s) > 0\}. \quad (5.A33)$$

Thus, the stochastically stable states are those that are most likely to be observed in the long run as the noise in the system becomes very small. The concept was introduced by Foster and Young (1990).

Theorem 16.15 Every stochastically stable state lies in the invariant distribution of the unperturbed dynamics.

Proof: By definition, $\bar{\lambda}_\varepsilon$ is the unique invariant distribution of the stochastic process (S, P_ε) , hence, $\bar{\lambda}_\varepsilon = \bar{\lambda}_\varepsilon P_\varepsilon$.

Using (5.A32) we get $\lambda^* = \lim_{\varepsilon \rightarrow 0} \bar{\lambda}_\varepsilon = \lim_{\varepsilon \rightarrow 0} \bar{\lambda}_\varepsilon P_\varepsilon$, or

$$\lambda^* = \lim_{\varepsilon \rightarrow 0} \bar{\lambda}_\varepsilon \lim_{\varepsilon \rightarrow 0} P_\varepsilon.$$

Using (5.A32) and continuity of P_ε in ε we have

$$\lambda^* = \lambda^* P_0.$$

Thus λ^* (which is the *limit invariant distribution*) is the *unique invariant distribution* of the unperturbed transition matrix P_0 .

Remark 16.1 Suppose that we start with some initial state of the unperturbed system, which coincides with the stochastically stable state, λ^* . Then, the long-run probability of any state s_i in the support of λ^* is positive. This can be easily shown:

$$\left[\lim_{t \rightarrow \infty} \lambda^* P_0^t \right] (s_i) = \lambda^* (s_i) > 0.$$

Recall that the *transition cost function* $c(s_i, s_j)$ gives the minimum number of steps required to make the transition from state s_i to state s_j . The point of introducing perturbed Markov chains was to analyze cases where it is difficult or cumbersome to directly compute the unique invariant distribution. In this sense, Theorem 16.15 gives only limited help because it does not tell us how to compute $\bar{\lambda}_0$. The following methods are useful in this regard.

THE METHOD OF YOUNG (1993A)

Let the recurrence classes for the unperturbed process, $\{S, P_0\}$, be given by $C_1, C_2, \dots, C_M$. If we focus just on the unperturbed process, then, by definition of the recurrence classes, we can never move to a state outside a recurrence class once we are in it. However, if the process is perturbed, $\{S, P_\varepsilon\}$, then we can move from a state inside a recurrence class of $\{S, P_0\}$ to a state outside the recurrence class. The next theorem tells us how to find the recurrence classes of the unperturbed process.

Theorem 16.16 (*Recurrence classes of the unperturbed process*): Suppose that all players in the game play an identical strategy $s_i \in S$. This is called a *monomorphic equilibrium*. The recurrence classes of the unperturbed process are the singletons $s_i \in S$, $i = 1, 2, \dots, n$. Thus, the recurrence classes of the unperturbed processes are the *monomorphic equilibria*.

The central concept in Young's method is that of the *stochastic potential of a directed graph*. We give a series of definitions that lead to this concept.

Definition 16.66 (*ij-path*): For any two distinct recurrence classes, C_i, C_j , an *ij-path*, denoted by κ_{ij} , is a sequence of states $s_i \rightarrow s_{i+1} \rightarrow \dots \rightarrow s_{j-1} \rightarrow s_j$ where $s_i \in C_i$ and $s_j \in C_j$.

Definition 16.67 (*Resistance of an ij-path*): Given an *ij path*, κ_{ij} , and a transition cost function $c(s_i, s_j)$, the *resistance of κ_{ij}* , denoted by $R(\kappa_{ij})$, is defined as

$$R(\kappa_{ij}) = c(s_i, s_{i+1}) + c(s_{i+1}, s_{i+2}) + \dots + c(s_{j-1}, s_j).$$

Definition 16.68 (Minimum resistance of an ij -path): For all possible ij -paths that take us from a state in C_i to a state in C_j , define the minimum resistance of an ij -path, denoted by ρ_{ij} by

$$\rho_{ij} = \min R(\kappa_{ij}) \text{ over all } ij\text{-paths.}$$

Recall that we have M recurrence classes, $C_1, C_2, \dots, C_M$. We now construct a *directed graph* or *tree* over the set of M nodes, one for each of the M recurrence classes. A tree that is rooted at a node corresponding to the j th recurrence class is called a j -tree. So, for instance, a tree that is rooted at a node that corresponds to C_1 is called a 1-tree, and at a node that corresponds to C_6 , a 6-tree and so on. The shapes of some basic trees are shown in Figures 5.A3, 5.A4, below.

Definition 16.69 (j -tree): A j -tree is rooted at the j th node and it comprises a set of $M - 1$ directed edges such that (i) every node is the initial point of one and only one directed edge, and (ii) from every node different from j , there is a unique directed path to the root of the tree (that comprises a node corresponding to the C_j recurrence class). The weight on the directed edge $i \rightarrow j$ is the minimum resistance of the path, ρ_{ij} .

Example 16.11 (Two recurrence classes): Suppose that there are two recurrence classes C_1, C_2 and each class has exactly one state, respectively s_1, s_2 . Then there are two trees, a 1-tree (which is rooted at the first recurrence class) and a 2-tree (which is rooted at the second recurrence class). This is shown in Figure 5.A3.

Example 16.12 (Three recurrence classes): Suppose that there are three recurrence classes, C_1, C_2, C_3 and each class has exactly one state, respectively, s_1, s_2, s_3 . There are three trees in this case, a 1-tree, a 2-tree, and a 3-tree. These trees are shown, respectively, as the first, second, and third rows in Figure 5.A4. For instance, the 1-tree shows the root at the unique element, s_1 , of the recurrence class, C_1 . There are three possible ways of reaching this root, each shown by the first row in Figure 5.A4. These are: (i) $s_3 \rightarrow s_2 \rightarrow s_1$, (ii) $s_3 \rightarrow s_1$ and $s_2 \rightarrow s_1$, and (iii) $s_2 \rightarrow s_3 \rightarrow s_1$. The 2-trees and the 3-trees can be similarly read off from the second and third rows of the figure.




Figure 5.A3 j -trees for two recurrence classes with one state each.




Figure 5.A4 1-tree (first row), 2-tree (second row), 3-tree (third row) for the case of three recurrence classes with one state each.

Definition 16.70 (*Resistance of a j -tree*): The resistance of a j -tree is the sum across all $M - 1$ directed edges of ρ_{ij} .

Definition 16.71 (*Stochastic potential*): The stochastic potential of the j th recurrence class, C_j , denoted by ϕ_j , is the minimum value of resistance across all possible j -trees for the states in C_j .

Theorem 16.17 (*Stochastically stable states*): The stochastically stable states are those which are contained in the recurrence classes with the minimum stochastic potential (this includes the possibility that there might be more than one recurrence classes, which has the minimum stochastic potential). So essentially, the root of the tree with the minimum stochastic potential is a stochastically stable steady state.

Example 16.13 Consider Figure 5.A3. Using Definition 16.71, the stochastic potential of C_1 is ρ_{21} , while that of C_2 is ρ_{12} . We can now use Theorem 16.17. C_1 is the stochastically stable state if $\rho_{21} < \rho_{12}$, while C_2 is the stochastically stable state if $\rho_{21} > \rho_{12}$.

Example 16.14 Consider Figure 5.A4. We can find the resistance of any tree by adding together the two resistances along its two edges. For instance, for the leftmost tree in the first row, the resistance is $\rho_{32} + \rho_{21}$. The resistance of the 1-tree is given by the minimum of the resistance across the three trees in the first row. We can likewise compute the resistance of the 2-tree and the 3-tree (by taking, respectively, the minimum across the second and third rows). Finally, we can take the minimum resistance across the three trees. The corresponding recurrence class is the stochastically stable steady state.

THE METHOD OF ELLISON (2000)

Definition 16.72 (*Radius and co-radius*): Suppose $s_i \in S$. Then the radius of s_i , $R(s_i)$ is given by

$$R(s_i) = \min_{s_j \in S} c(s_i, s_j), s_i \neq s_j.$$

The co-radius $CR(s_i)$ is given by

$$CR(s_i) = \max_{s_j \in S} c(s_i, s_j), s_i \neq s_j.$$

Theorem 16.18 (Ellison, 2000): If $R(s_i) > CR(s_i)$, then s_i is the only stochastically stable state.

Exercises for Part 5

- Recall from Figure 14.2 that ESS is a subset of the set of Nash equilibria of a game. Consider the normal form game in Figure 5.E1. There are two pure strategy Nash equilibria (A, A) and (B, B). Show that A is not an ESS, but B is.
- (Gintis, 2009a, Theorems 10.1, 10.2): Consider the game $G(2, S, \pi)$ and $S = (s_1, s_2)$.
 - Suppose that $a_{11} \neq a_{21}$ and $a_{12} \neq a_{22}$ (see Section 14.2 for this notation; a_{ij} is the payoff of player i when he follows the strategy s_i and the opponent follows the strategy s_j). Then show that $G(2, S, \pi)$ has an evolutionarily stable strategy.
 - Show that $G(2, S, \pi)$ has a mixed strategy Nash equilibrium if $a_{11} < a_{21}$ and $a_{22} < a_{12}$. Show that this equilibrium is an ESS.
- Consider the stag-hunt game in Figure 5.E2.

Each player has two pure strategies. Hunt hare (H) or stag (S). Hunting stag requires a cooperative effort that is only successful if both parties engage in it; a successful stag hunt gives each player a payoff of $s > 0$. Hunting hare gives each player a payoff of $0 < h < s$, irrespective of what the other player does. Finally, hunting stag gives a payoff of 0 if the other player hunts hare.

 - Show that there are three Nash equilibria: (H,H), (S,S) and a mixed strategy equilibrium in which $\sigma(H) = \frac{s-h}{s}$ and $\sigma(S) = \frac{h}{s}$.
 - Show that each pure strategy Nash equilibrium is an ESS, but the mixed strategy Nash equilibrium is not an ESS.
- Consider the hawk-dove game in Figure 5.E3.

The strategies in this game are dove (D) and hawk (H), which are, respectively, the non-aggressive and aggressive strategies. Two players fight for a fixed prize that has value v . If both play aggressively, each shares the prize v , but each player also incurs a fixed cost $c > v$. If a dove meets a hawk, then the dove backs down and the hawk gets the prize.

 - Show that there are no pure strategy Nash equilibria, and in the unique mixed strategy equilibrium, $\sigma(D) = \frac{c-v}{c}$ and $\sigma(S) = \frac{v}{c}$.
 - Show that the mixed strategy Nash equilibrium is an ESS.
 - (Smith, 1982) How might this simple example account for ritualized fighting in animals?

	A	B
A	2, 2	2, 2
B	2, 2	4, 4

Figure 5.E1 A 2×2 game.

	H	S
H	h, h	h, 0
S	0, h	s, s

Figure 5.E2 The stag-hunt game.

	D	H
D	$v/2, v/2$	$0, v$
H	$v, 0$	$(v-c)/2, (v-2)/2$

Figure 5.E3 A hawk–dove game.

	A	B	C
A	1, 1	1, 1	0, 0
B	1, 1	0, 0	1, 0
C	0, 0	0, 1	0, 0

Figure 5.E4 Evolutionary stable strategies are not unbeatable.

Source: *Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Interaction* by Gintis, Herbert (2009). Reproduced with permission of Princeton University Press in the format Book via Copyright Clearance Center.

5. (Multiple mutations: Gintis, 2009a, Section 10.10) Consider the normal form game shown in Figure 5.E4.
 - (a) Show that the pure strategy A is an ESS.
 - (b) Starting from the established strategy A, consider two simultaneous mutations. A fraction $0 < \lambda_B < 1$ of the players play B (type-B mutants) and a fraction $\lambda_C > \lambda_B$ of players play C (type-C mutants). In the face of these simultaneous mutations, show that the type-B mutants have a higher payoff than players who play A. This is the sense in which the ESS, A, is not “unbeatable.”¹
6. Consider the following coordination game in Section 14.6.4:

	A	D
A	1, 1	0, 0
D	0, 0	1, 1

Suppose now that this game is played by players who are randomly matched from the same population. Show that the equilibrium in which one half of the population plays A, while the other half plays D is asymptotically stable.

7. Consider the following coordination game played between two populations of players who are randomly matched.

1\2	A	D
A	0, 0	2, 1
D	1, 2	0, 0

¹ As to the usefulness of this concept, Gintis (2009a, p. 240) writes: “But unbeatability is not a very important concept, because it has no interesting dynamic properties.”

Show that the polymorphic equilibrium in which each population plays strategy A with probability $\frac{2}{3}$ is unstable and that it is a saddle path. What are the stable equilibria of this game in replicator dynamics?

8. Consider the rock–scissors–paper game that we considered above in Section 14.6.3. Suppose that the payoffs are instead given by the payoff matrix:

	R	S	P
R	ε, ε	$1, -1$	$-1, 1$
S	$-1, 1$	ε, ε	$1, -1$
P	$1, -1$	$-1, 1$	ε, ε

- (i) Let $\varepsilon > 0$ be small. Show that the unique Nash equilibrium is still $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.
 (ii) Show that the Nash equilibrium is an unstable steady state of the replicator dynamics such that the trajectories spiral outwards from $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ without ever touching the boundaries of the simplex.
 (iii) Suppose that $\varepsilon < 0$ is small. Show that the Nash equilibrium in this case is an asymptotically stable steady state.
9. (van Damme, 1987) For many examples in the text, ESS coincided with the asymptotically stable states of replicator dynamics. From Figure 14.2, we know that the latter set is larger. To illustrate this point, consider the payoff matrix.

	A	B	C
A	$0, 0$	$1, -2$	$1, 1$
B	$-2, 1$	$0, 0$	$4, 1$
C	$1, 1$	$1, 4$	$0, 0$

- (i) Show that the game has a unique symmetric equilibrium $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.
 (ii) Show that $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is not an ESS and it can be successfully invaded by the strategy $(0, \frac{1}{2}, \frac{1}{2})$.
 (iii) Show that $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is asymptotically stable in the replicator dynamics.
10. Prove Proposition 14.18.
 11. Prove Proposition 14.19.
 12. (Gintis, 2011): Let g, c, e be, respectively, vectors of genetic variables, cultural variables, and environmental variables. Suppose that in an evolutionary model we have the following system of first order differential equations.

$$\begin{cases} \dot{g} = F(g, c, e) \\ \dot{c} = G(g, c, e) \\ \dot{e} = H(e) \end{cases} \quad (5.E1)$$

An argument is sometimes made that culture is an “effect of genes that can be factored out in the long run.” This is equivalent to the assertion that c can be eliminated from the equations in (5.E1). Based on the statement given, derive a contradiction using (5.E1).

13. Consider the following statement in Gintis et al. (2014, p.19): “Fire and cooking thus co-evolved with the emergence of a normative order and social organization based on ethical behavior.” Can you think of arguments to support this statement?
14. (Bowles and Gintis, 2011, pp. 6, 8) Consider the following influential views that strongly advocate that humans are essentially selfish.
- (i) In his celebrated book, Richard Dawkins (2006) writes:² “Let us try to teach generosity and altruism, because we are born selfish.”
 - (ii) Richard Alexander (1987) writes: “ethics, morality, human conduct, and the human psyche are to be understood only if societies are seen as collections of individuals seeking their own self-interest.”
 - (iii) Garrett Hardin (1968) gave an account of the tragedy of the commons and the importance of self-interest in the depletion of common property resources. He writes (p. 1246): “The tragedy cannot be solved by an appeal to conscience, for those who heeded the appeal would have fewer children, and by the heritability of the capacity of the conscience, this would lead to a less moral population.”

In light of these views how can we explain the central theme in Part 2 of the book that most humans have other-regarding preferences and possess character virtues?

15. Consider a technology adoption game in which there is a choice of two technologies s_1 , s_2 . There is a large population, and in every period a single player gets the opportunity to make a technology choice (perhaps the technology becomes obsolete). Before making the technology choice, the player samples m individuals about their choice of technology, who report back honestly. The payoff matrix is shown below.

	s_1	s_2
s_1	a, a	$0, 0$
s_2	$0, 0$	b, b

Consider a perturbed Markov process in which the individual plays a best response with probability $1 - \varepsilon$ ($\varepsilon > 0$ is small) and randomizes equally between the two technologies with equal probability (i.e., $\frac{\varepsilon}{2}$ each). Show that if $a > b$ then the stochastically stable choice is to invest in technology a , while if $a < b$, then the stochastically stable choice is b . Show that such a choice also satisfies the criterion of risk dominance.

16. (Young, 2006) Consider again a technology adoption game. The setup is as in Exercise 15. Suppose that the payoff from a technology arises from two sources. The first source is *networking* which arises from interaction between different technologies. The second source is *payoff from own use* which arises independently of networking effects. The payoff matrix is shown below.

	s_1	s_2
s_1	$a + a', a + a'$	$c + a', c + b'$
s_2	$c + b', c + a'$	$b + b', b + b'$

² This is a much caricatured view. See Dawkin’s clarification in the introduction to the anniversary edition (2006, p. ix).

The payoffs without primes arise from networking effects, while the payoffs with primes arise from own use. We assume that $a + a' > c + b'$ and $b + b' > c + a'$, so we have a symmetric coordination game. Consider the special case of $a = 4$, $b = 1$, $c = 1$, $a' = 1$, $b' = 3$. Show that the stochastically stable outcome is for everyone to adopt s_2 , but that this is not a Pareto optimal outcome.

17. (Young, 2006) Consider the same setup as in Exercise 15, but now there is a choice between three competing technologies, s_1 , s_2 , s_3 . The payoff matrix is given by

	s_1	s_2	s_3
s_1	5, 5	0, 0	0, 0
s_2	0, 0	4, 4	0, 0
s_3	0, 0	0, 0	3, 3

Show that in the stochastically stable state, technology s_1 is chosen.

18. (Young, 1998) Several different commodities have historically been used as a medium of exchange. Young (1998) gives a fascinating account. Consider the following stylized example in which there are two possible commodities, gold and silver, that can be used as a medium of exchange. There is a single population. Individuals are matched each period to engage in a transaction. Each period, one individual is drawn at random (the experimenter). He chooses either to carry gold or silver. Consider the Markov process, $P^{T,\tau,\varepsilon}$; thus, the selected individual takes a sample of size $\tau \leq T$ and plays with some noise ε . The payoff matrix is given by:

	Gold	Silver
Gold	3, 3	0, 0
Silver	0, 0	2, 2

Thus, of the individuals sampled, if a proportion p holds gold, then the payoff to the experimenter from holding gold is $3p$, while that from holding silver is $2(1-p)$. With probability, $1 - \varepsilon$ the experimenter plays a best response to these frequencies and with probability ε (noise) he chooses each of the two actions with equal probability. Show that the stochastically stable equilibrium is (Gold, Gold).

19. (Young, 1998) Along the lines of the discussion in Section 16.3.3, consider the following 2×2 game of chicken.

	Defensive	Aggressive
Defensive	10, 5	0, 0
Aggressive	7, 7	5, 10

Suppose that the row and the column players use different sample sizes. Greater information corresponds to a larger sample size. Show that more information is advantageous.

REFERENCES FOR PART 5

- Abbot, P., Abe, J., Alcock, J., et al. (2011). Inclusive fitness theory and eusociality. *Nature* 471: E1–E4.
- Alexander, R. D. (1987). *The Biology of Moral Systems*. New York: Aldine.
- Alger I. and Weibull, J. W. (2010). Kinship, incentives and evolution. *American Economic Review* 100: 1725–58.
- Alger I. and Weibull, J. W. (2013). Homo Moralis—preference evolution under incomplete information and assortative matching. *Econometrica* 81(6): 2269–302.
- Alós-Ferrer, C. (2000). Learning, memory, and inertia. University of Vienna, Department of Economics. Vienna Economics Papers 0003.
- Amaldoss, W., Chang, J.-K., and Ho, T.-H. (2005). Choosing the right pond: EWA learning in games with different group sizes. Working paper, Department of Marketing, University of California, Berkeley.
- Amemiya, T. (1981). Qualitative response models: a survey. *Journal of Economic Literature* 19(4): 483–536.
- Anderson, C. and Camerer, C. F. (2001). Experience-weighted attraction learning in sender-receiver signaling games. *Economic Theory* 16: 689–718.
- Andreoni, J. and Miller, J. H. (1993). Rational cooperation in the finitely repeated prisoner's dilemma: experimental evidence. *Economic Journal* 103: 570–85.
- Andreoni, J. and Miller, J. H. (2002). Giving according to GARP: an experimental test of the consistency of preferences for altruism. *Econometrica* 70(2): 737–53.
- Apesteguia J., Huck, S., Oechssler, J., and Weidenholzer, S. (2010). Imitation and the evolution of Walrasian behavior: theoretically fragile but behaviorally robust. *Journal of Economic Theory* 145(5): 1603–17.
- Arthur, W. B. (1991). Designing economic agents that act like human agents: a behavioral approach to bounded rationality. *American Economic Review Papers and Proceedings* 81: 353–9.
- Arthur, W. B. (1993). On designing economic agents that behave like human agents. *Journal of Evolutionary Economics* 3: 1–22.
- Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *American Economic Association Papers and Proceedings* 84: 406–11.
- Axelrod, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.
- Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science* 211: 1390–96.
- Barrett, H. C., Cosmides, L., and Tooby, J. (2007). The hominid entry into the cognitive niche. In S. Gangestad and J. Simpson (eds.), *The Evolution of Mind: Fundamental Questions and Controversies*. New York: Guilford, pp. 241–8.
- Battalio, R., Samuelson, L., and Van Huyck, J. (2001). Optimization incentives and coordination failure in laboratory stag hunt games. *Econometrica* 69: 749–64.
- Bednar, J., Chen, Y., Xiao Liu, T., and Page, S. (2012). Behavioral spillovers and cognitive load in multiple games: an experimental study. *Games and Economic Behavior* 74(1): 12–31.
- Bendor, J., Mookerjee, D., and Ray, D. (2001). Aspiration-based reinforcement learning in repeated interaction games: an overview. *International Game Theory Review* 3: 159–74.
- Benjamin, D. J., Cesarini, D., Chabris, C. F., et al. (2012). The promises and pitfalls of genoecconomics. *Annual Review of Economics* 4: 627–62.
- Bereby-Meyer, Y. and Erev, I. (1998). On learning to become a successful loser: a comparison of alternative abstractions of learning in the loss domain. *Journal of Mathematical Psychology* 42: 266–86.
- Bergstrom, T. (2003). The algebra of assortative encounters and the evolution of cooperation. *International Game Theory Review* 5: 211–28.
- Bergstrom, T. (2009). Ethics, evolution and games among neighbors. Working paper, UCSB.
- Bicchieri, C. (2006). *The Grammar of Society: The Nature and Dynamics of Social Norms*. New York: Cambridge University Press.
- Binmore, K. (1987). Modelling rational players (part I). *Economics and Philosophy* 3: 179–214.
- Binmore, K. and Samuelson, L. (1997). Muddling through: noisy equilibrium selection. *Journal of Economic Theory* 74: 235–65.
- Bisin, A., Topa, G., and Verdier, T. (2004). Religious intermarriage and socialization in the United

- States. *Journal of Political Economy* 112(3): 615–64.
- Bisin, A. and Verdier, T. (2011). The economics of cultural transmission and socialization. In J. Benhabib, A. Bisin, and M. Jackson (eds.), *Handbook of Social Economics*. Elsevier, pp. 339–416.
- Björnerstedt, J. (1995). Experimentation, imitation and evolutionary dynamics. Mimeo, Stockholm University.
- Björnerstedt, J. and Weibull, J. W. (1993). Nash equilibrium and evolution by imitation. DELTA (Ecole normale supérieure) Working Papers 93–23.
- Blackwell, D. and Dubins, L. (1962). Merging of opinions with increasing information. *Annals of Mathematical Statistics* 38: 882–6.
- Blume, L. E. (1993). The statistical mechanics of strategic interaction. *Games and Economic Behavior* 5: 387–426.
- Blume, L. E. (1995). The statistical mechanics of best-response strategy revision. *Games and Economic Behavior* 11: 111–145.
- Blume, L. E. (2003). How noise matters. *Games and Economic Behavior* 44: 251–71.
- Blume, A., DeJong, D. V., Kim, Y.-G., and Sprinkle, G. B. (2001). Evolution of communication with partial common interest. *Games and Economic Behavior* 37: 79–120.
- Börgers, T. and Sarin, R. (1997). Learning through reinforcement and replicator dynamics. *Journal of Economic Theory* 77: 1–14.
- Börgers, T. and Sarin, R. (2000). Naive reinforcement learning with endogenous aspirations. *International Economic Review* 41: 921–50.
- Bosch-Domènech, A. and Vriend, N. J. (2003). Imitation of successful behavior in Cournot markets. *Economic Journal* 113: 495–524.
- Botelho, A., Dinar, A., Pinto, L. M. C., and Rapoport, A. (2014). Time and uncertainty in resource dilemmas: equilibrium solutions and experimental results. *Experimental Economics* 1: 649–72.
- Bowles, S. and Gintis, H. (2011). *A Cooperative Species: Human Reciprocity and Evolution*. Princeton, NJ: Princeton University Press.
- Boyd, R., Gintis, H., and Bowles, S. (2010). Coordinated punishment of defectors sustains cooperation and can proliferate when rare. *Science* 328: 617–20.
- Boyd, R. and Richerson, P. J. (1985). *Culture and the evolutionary process*. Chicago: University of Chicago Press.
- Boyd, R., and Richerson, P. J. (1988). The evolution of reciprocity in sizable groups. *Journal of Theoretical Biology* 132: 337–56.
- Boyd, R., Richerson, P. J., and Henrich, J. (2011). The cultural niche: why social learning is essential for human adaptation. *Proceedings of the National Academy of Sciences (USA)* 108: 10918–25.
- Boylan, R. and El Gamal, M. (1993). Fictitious play: a statistical study of multiple economic experiments. *Games and Economic Behavior* 5: 205–22.
- Brandt, H. and Sigmund, K. (2004). The logic of reprobation: assessment and action rules for indirect reciprocation. *Journal of Theoretical Biology* 231(4): 475–86.
- Broseta, B., Costa-Gomes, M., and Crawford, V. P. (2000). Cognition and behavior in normal-form games: an experimental study. Department of Economics, UCSD. Paper 2000-02R.
- Brown, G. (1951). Iterative solution of games by fictitious play. In T. C. Koopmans (ed.), *Activity Analysis of Production and Allocation*. New York: John Wiley.
- Brown, J. S., Clark, F. R., and Stein, L. (1958). A new technique for studying spatial generalization with voluntary responses. *Journal of Experimental Psychology* 55: 359–62.
- Brunner, C., Camerer, C. F., and Goeree, J. K. (2011). Stationary concepts for experimental 2×2 games: comment. *American Economic Review* 101(2): 1029–40.
- Bush R. R. and Mosteller, F. (1951). A mathematical model for simple learning. *Psychological Review* 58: 313–23.
- Bush, R. R. and Mosteller, F. (1955). *Stochastic Models for Learning*. New York: John Wiley & Sons.
- Cachon, G. P. and Camerer, C. F. (1996). Loss-avoidance and forward induction in experimental coordination games. *Quarterly Journal of Economics* 111: 165–94.
- Camerer, C. F. (2003). *Behavioral Game Theory: Experiments in Strategic Interaction*. Princeton, NJ: Princeton University Press.
- Camerer, C. F. and Ho, T.-H. (1997). Experience-weighted attraction learning in games: a unifying approach. California Institute of Technology, Division of the Humanities and Social Sciences. Working Papers 1003.

- Camerer, C. F. and Ho, T.-H. (1999). Experience-weighted attraction learning in normal form games. *Econometrica* 67: 827–74.
- Camerer, C. F. and Ho, T.-H. (2014). Behavioral game theory experiments and modelling. In *Handbook of Game Theory*. Forthcoming.
- Camerer, C. F., Ho, T.-H., and Chong, J. K. (2001). Behavioral game theory: thinking, learning and teaching. Paper presented at the Nobel Prize symposium.
- Camerer, C. F., Ho, T.-H., and Chong, J. K. (2002). Sophisticated EWA learning and strategic teaching in repeated games. *Journal of Economic Theory* 104: 137–88.
- Camerer, C. F., Ho, T.-H., and Chong, J. K. (2003). Models of thinking, learning and teaching in games. *American Economic Review* 93: 192–5.
- Camerer, C. F., Ho, T.-H., and Wang, X. (1999). Individual differences and payoff learning in Games. University of Pennsylvania Working Paper.
- Camerer, C. F., Johnson, E., Rymon, T., and Sen, S. (1994). Cognition and framing in sequential bargaining for gains and losses. In K. Binmore, A. Kirman, and P. Tani (eds.), *Frontiers of Game Theory*. Cambridge, MA: MIT Press, pp. 27–47.
- Capra M., Goeree, J., Gomez, R., and Holt, C. (1999). Anomalous behavior in a traveler's dilemma. *American Economic Review* 89: 678–90.
- Cavalli-Sforza, L. L., Feldman, M. W., Chen, K. H., and Doubusch, S. M. (1982). Theory and observation in cultural transmission. *Science* 218: 19–27.
- Chen, Y. and Tang, F. (1998). Learning and incentive-compatible mechanisms for public goods provision: an experimental study. *Journal of Political Economy* 106: 633–62.
- Cheung, Y.-W. and Friedman, D. (1997). Learning in evolutionary games: some laboratory results. *Games and Economic Behavior* 19: 46–76.
- Cheung, Y.-W. and Friedman, D. (1998). A comparison of learning and replicator dynamics using experimental data. *Journal of Economic Behavior and Organization* 35: 263–80.
- Chmura, T., Goerg, S. J., and Selten, R. (2012). Learning in experimental 2×2 games. *Games and Economic Behavior* 76(1): 44–73.
- Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2009). Reinforcement learning and savings behavior. *Journal of Finance* 64: 2515–34.
- Clark, K. and Sefton, M. (1999). Matching protocols in experimental games. University of Manchester working paper.
- Coddington, E. A. (1961). *An Introduction to Ordinary Differential Equations*. Englewood Cliffs, NJ: Prentice-Hall.
- Coleman, J. S. (1990). *Foundations of Social Theory*. Cambridge, MA: Harvard University Press.
- Colman, A. M., Pulford, B. D., Omtzigt, D., and al-Nowaihi, A. (2010). Learning to cooperate without awareness in multiplayer minimal social situations. *Cognitive Psychology* 61: 201–27.
- Costa-Gomes, M., Crawford, V. P., and Broseta, B. (2001). Cognition and behavior in normal-form games: an experimental study. *Econometrica* 69: 1193–235.
- Cournot, A. ([1838]1927). *Researches into the Mathematical Principles of the Theory of Wealth* (translated by N.T. Bacon, 1897). New York: Macmillan.
- Crawford, V. P. (1990a). On the definition of an evolutionarily stable strategy in the “playing the field” model. *Journal of Theoretical Biology* 143: 269–73.
- Crawford, V. P. (1990b). Nash equilibrium and evolutionary stability in large- and finite-population “playing the field” models. *Journal of Theoretical Biology* 145: 83–94.
- Crawford, V. P. (1991). An evolutionary interpretation of Van Huyck, Battalio, and Beil's experimental results on coordination. *Games and Economic Behavior* 3: 25–59.
- Crawford, V. P. (1995). Adaptive dynamics in coordination games. *Econometrica* 63: 103–44.
- Crawford, V. P. (1997). Theory and experiment in the analysis of strategic interactions. In D. Kreps and K. Wallis (eds.), *Advances in Economics and Econometrics: Theory and Applications, Seventh World Congress, Volume I*. Cambridge: Cambridge University Press.
- Crawford, V. P. and Broseta, B. (1998). What price coordination? The efficiency-enhancing effect of auctioning the right to play. *American Economic Review* 88: 198–225.
- Dasgupta, P. and Maskin, E. (2005). Uncertainty and hyperbolic discounting. *American Economic Review* 95(4): 1290–9.
- Dawkins, R. (2006). *The Selfish Gene (30th Anniversary Edition)*. New York: Oxford University Press.

- Dekel, E., Ely, J. C., and Yilankaya, O. (2007). Evolution of preferences. *Review of Economic Studies* 74: 685–704.
- Duffy, J. and Feltovich, N. (1999). Does observation of others affect learning in strategic environments? An experimental study. *International Journal of Game Theory* 28(1): 131–52.
- Dufwenberg, M., Sundaram, R., and Butler, D. (2010). Epiphany in the Game of 21. *Journal of Economic Behavior and Organization* 75(2): 132–43.
- Dunbar, R. M. (1993). Coevolution of neocortical size, group size and language in humans. *Behavioral Brain Sciences* 16: 681–735.
- Durrett, R. and Levin, S. A. (1994). The importance of being discrete (and spatial). *Theoretical Population Biology* 46: 363–94.
- Ellison, G. (1993). Learning, local interaction and coordination. *Econometrica* 61: 1047–71.
- Ellison, G. (2000). Basins of attraction, long-run stochastic stability, and the speed of step-by-step evolution. *Review of Economic Studies* 67: 17–45.
- Ellison, G. and D. Fudenberg (1993). Rules of thumb for social learning. *Journal of Political Economy* 101: 612–43.
- Elster, J. (1989). Social norms and economic theory. *Journal of Economic Perspectives* 3: 99–117.
- Elster, J. (1999). *Alchemistries of the Mind*. Cambridge: Cambridge University Press.
- Ely, J. C. and Yilankaya, O. (2001). Nash equilibrium and the evolution of preferences. *Journal of Economic Theory* 97(2): 255–72.
- Erev, I. and Roth, A. E. (1998). Predicting how people play games: reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review* 88: 848–81.
- Erev, I. and Roth, A. E. (2007). Multi-agent learning and the descriptive value of simple models. *Artificial Intelligence* 171: 423–8.
- Fehr, E. and Joseph, H. (2004). Is strong reciprocity a maladaptation? On the evolutionary foundations of human altruism. In P. Hammerstein (ed.), *Genetic and Cultural Origins of Cooperation*. Cambridge, MA: MIT Press.
- Fehr, E. and Schmidt, K. M. (2006). The economics of fairness, reciprocity and altruism: experimental evidence and new theories. In S. Kolm and J. M. Ythier (eds.) *Handbook of the Economics of Giving, Altruism and Reciprocity*. Elsevier.
- Feltovich, N. (1999). Equilibrium and reinforcement learning in private-information games: an experimental study. *Journal of Economic Dynamics and Control* 23: 1605–32.
- Feltovich, N. (2000). Reinforcement-based vs. belief-based learning models in experimental asymmetric information games. *Econometrica* 68: 605–41.
- Fischbacher, U. and Föllmi-Heusi, F. (2013). Lies in disguise: an experimental study on cheating. *Journal of the European Economic Association* 11(3): 525–47.
- Foster, D. P. and Vohra, R. (1993). A randomized rule for selecting forecasts. *Operations Research* 41: 704–9.
- Foster, D. P. and Vohra, R. (1998). Asymptotic calibration. *Biometrika* 85: 379–90.
- Foster, D. P. and Young, H. P. (1990). Stochastic evolutionary game dynamics. *Theoretical Population Biology* 38(2): 219–32.
- Foster, D. P. and Young, H. P. (2001). On the impossibility of predicting the behavior of rational agents. *Proceedings of the National Academy of Sciences (USA)* 98(222): 12848–53.
- Foster, D. P. and Young, H. P. (2003). Learning, hypothesis testing, and Nash equilibrium. *Games and Economic Behavior* 45: 73–96.
- Fryer, R. and Jackson, M. (2008). Categorical cognition: a psychological model of categories and identification in decision making. *B.E. Journal of Theoretical Economics (Contributions)* 8(1) article 6.
- Fudenberg, D. and Kreps, D. (1990). Lectures on learning and equilibrium in strategic-form games. CORE lecture series.
- Fudenberg, D. and Kreps, D. (1993). Learning mixed equilibria. *Games and Economic Behavior* 5: 320–67.
- Fudenberg, D. and Levine, D. K. (1989). Reputation and equilibrium selection in games with a patient player. *Econometrica* 57: 759–78.
- Fudenberg, D. and Levine, D. K. (1995). Consistency and cautious fictitious play. *Journal of Economic Dynamics and Control* 19: 1065–89.
- Fudenberg, D. and Levine, D. K. (1998). *The Theory of Learning in Games*. Cambridge, MA: MIT Press.
- Fudenberg, D., Levine, D. K., and Maskin, E. (1994). The folk theorem with imperfect public information. *Econometrica* 62: 997–1039.

- Fudenberg, D. and Maskin, E. (1990). Evolution and cooperation in noisy repeated games. *American Economic Review* 80: 274–9.
- Gintis, H. (2009a). *Game Theory Evolving: A Problem Centered Introduction to Modeling Strategic Interaction*. Princeton, NJ: Princeton University Press.
- Gintis, H. (2009b). *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences*. Princeton, NJ: Princeton University Press.
- Gintis, H. (2011). Gene-culture coevolution and the nature of human sociality. *Philosophical Transactions of the Royal Society* 366: 878–88.
- Gintis, H. (2014). Inclusive fitness and the sociobiology of the genome. *Biology and Philosophy* 29: 477–515.
- Gintis, H. and Helbing, D. (2014). Homo Socialis: an analytical core for sociological theory. Mimeo.
- Gintis, H., van Schaik, C., and Boehm, C. (2014). The evolutionary origins of human political systems. Forthcoming in *Current Anthropology*.
- Gneezy, U., Rockenbach, B., and Serra-Garcia, M. (2013). Measuring lying aversion. *Journal of Economic Behavior and Organization* 93: 293–300.
- Gneezy, U., Rustichini, A., and Vostroknutor, A. (2010). Experience and insight in the race game. *Journal of Economic Behavior and Organization* 75(2): 144–55.
- Grimm, V. and Mengel, F. (2012). An experiment on learning in multiple game environments. *Journal of Economic Theory* 147: 2220–59.
- Güth, W. (1995). An evolutionary approach to explaining cooperative behavior by reciprocal incentives. *International Journal of Game Theory* 24: 323–44.
- Güth, W. and Yaari, M. E. (1992). Explaining reciprocal behavior in simple strategic games: an evolutionary approach. In U. Witt (ed.), *Explaining Process and Change*. Ann Arbor, MI: University of Michigan Press, pp. 23–34.
- Guttman, N. and Kalish, H. I. (1956). Discriminability and stimulus generalization. *Journal of Experimental Psychology* 51: 79–88.
- Hagen, E. and Hammerstein, P. (2006). Game theory and human evolution: a critique of some recent interpretations of experimental games. *Theoretical Population Biology* 69: 339–48.
- Haigh, J. (1975). Game theory and evolution. *Advances in Applied Probability* 7: 8–11.
- Haldane, J. B. S. (1932). *The Causes of Evolution*. New York: Harper Brothers.
- Haley, K. J. and Fessler, D. M. (2005). Nobody's watching? Subtle cues affect generosity in an anonymous economic game. *Evolution and Human Behavior* 26: 245–56.
- Hamilton W. D. (1963). The evolution of altruistic behavior. *The American Naturalist* 96: 354–6.
- Hamilton, W. D. (1964). The genetical evolution of social behaviour, I and II. *Journal of Theoretical Biology* 7: 1–16, 17–52.
- Hammerstein, P. (2003). Why is reciprocity so rare in social animals? A protestant appeal. In P. Hammerstein (ed.), *Genetic and Cultural Evolution of Cooperation*. Cambridge, MA: MIT Press, pp. 84–93.
- Hammerstein, P. and Selten, R. (1994). Game theory and evolutionary biology. In R. Aumann, and S. Hart (eds.), *Handbook of Game Theory, volume 2*, Amsterdam: North-Holland, pp. 929–93.
- Hannan, J. (1957). Approximation to Bayes risk in repeated play. In M. Dresher, A. W. Tucker, and P. Wolfe (eds.), *Contributions to the Theory of Games III (Annals of mathematics studies 39)*. Princeton, NJ: Princeton University Press, pp. 97–139.
- Hardin, G. (1968). The tragedy of the commons. *Science* 162(3859): 1243–8.
- Hart, S. and Mas-Colell, A. (2000). A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68: 1127–50.
- Hart, S. and Mas-Colell, A. (2006). Stochastic uncoupled dynamics and Nash equilibrium. *Games and Economic Behavior* 57: 286–303.
- Hauert, C. and Doebeli, M. (2004). Spatial structure often inhibits the evolution of cooperation in the snowdrift game. *Nature* 428: 643–6.
- Henrich, J., Boyd, R., Bowles, S., et al. (2001). Cooperation, reciprocity and punishment in fifteen small-scale societies. *American Economic Review* 91: 73–8.
- Hirsch, M. W. and Smale, S. (1974). *Differential Equations, Dynamical Systems, and Linear Algebra*. New York: Academic Press.
- Ho, T.-H., Camerer, C. F., and Chong, J.-K. (2002). Economic value of EWA lite: a functional theory of learning in games, working paper.
- Ho, T.-H., Camerer, C. F., and Chong, J. K. (2007). Self-tuning experience-weighted attraction

- learning in games. *Journal of Economic Theory* 133: 177–98.
- Ho, T.-H., Camerer, C. F., and Weigelt, K. (1998). Iterated dominance and iterated best response in experimental p-beauty contests. *American Economic Review* 88: 947–69.
- Ho, T.-H., Wang, X., and Camerer, C. F. (2008). Individual differences in the EWA learning with partial payoff information. *Economic Journal* 118: 37–59.
- Ho, T. and Weigelt, K. W. (1996). Task complexity, equilibrium selection, and learning. *Management Science* 42(5): 659–79.
- Hofbauer, J., Schuster, P., and Sigmund, K. (1979). A note on evolutionary stable strategies and game dynamics. *Journal of Theoretical Biology* 81: 609–12.
- Hofbauer, J. and Sigmund, K. (1988). *The Theory of Evolution and Dynamical Systems*. Cambridge: Cambridge University Press.
- Hofbauer, J. and Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge: Cambridge University Press.
- Hofbauer, J. and Weibull, J. (1996). Evolutionary selection against dominated strategies. *Journal of Economic Theory* 71: 558–73.
- Hopkins, E. (2002). Two competing models of how people learn in games. *Econometrica* 70(6): 2141–66.
- Hopkins, E. and Posch, M. (2005). Attainability of boundary points under reinforcement learning. *Games and Economic Behavior* 53: 110–125.
- Huck, S., Konrad, K. A., Müller, W., and Normann, H.-T. (2007). The merger paradox and why aspiration levels let it fail in the laboratory. *Economic Journal* 117: 1073–95.
- Huck, S., Normann, H., and Oechssler, J. (1999). Learning in Cournot oligopoly: an experiment. *Economic Journal* 109: C80–95.
- Huck, S., Normann, H., and Oechssler, J. (2004). Through trial and error to collusion. *International Economic Review* 45: 205–24.
- Hume, D. ([1739]1978). *A Treatise of Human Nature*. Oxford: Oxford University Press.
- Hyndman, K., Ozbay, E. Y., Schotter, A., and Ehrblatt, W. Z. (2012). Convergence: an experimental study of teaching and learning in repeated games. *Journal of the European Economic Association* 10: 573–604.
- Jackson, M. O. (2010). *Social and Economic Networks*. Princeton, NJ: Princeton University Press.
- Jehiel, P. (1995). Limited horizon forecast in repeated alternate games. *Journal of Economic Theory* 67: 497–519.
- Jehiel, P. (1998). Learning to play limited forecast equilibria. *Games and Economic Behavior* 22: 274–98.
- Jordan, J. S. (1992). Bayesian learning in games: a non-Bayesian perspective. Manuscript, University of Minnesota.
- Jordan, J. S. (1993). Three problems in learning mixed-strategy equilibria. *Games and Economic Behavior* 5: 368–86.
- Kalai, E. and Lehrer, E. (1993). Rational learning leads to Nash equilibrium. *Econometrica* 61: 1019–45.
- Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long-run equilibria in games. *Econometrica* 61: 29–56.
- Karandikar, R., Mookerjee, D., Ray, D., and Vega-Redondo, F. (1998). Evolving aspirations and cooperation. *Journal of Economic Theory* 80: 292–331.
- Karlin, S. and Taylor, H. M. (1998). *A First Course in Stochastic Processes*. 3rd edition. New York: Academic Press.
- Kocher, M. and Sutter, M. (2005). The decision maker matters: individual versus team behavior in experimental beauty-contest games. *Economic Journal* 115: 200–23.
- Kreindler, G. E. and Young, H. P. (2013). Fast convergence in evolutionary equilibrium selection. *Games and Economic Behavior* 80: 39–67.
- Kreindler, G. E. and Young, H. P. (2014). Rapid innovation diffusion in social networks. *Proceedings of the National Academy of Sciences* 111(3): 10881–8.
- Laslier, J. F., Topol, R., and Walliser, B. (2001). A behavioral learning process in games. *Games and Economic Behavior* 37: 340–66.
- Lieberman, E., Hauert, C., and Nowak, M. A. (2005). Evolutionary dynamics on graphs. *Nature* 433: 312–16.
- Luce, R. D. (1959). *Individual Choice Behaviour*. New York: Wiley.
- Lumsden, C. and Wilson, E. O. (1980). Translation of epigenetic rules of individual behavior into ethnographic patterns. *Proceedings of the National Academy of Sciences (USA)*. 77: 4382–6.
- Lumsden, C. and Wilson, E. O. (1981). *Genes, Mind, and Culture: The Coevolutionary Process*. Cambridge, MA: Harvard University Press.

- McAllister, P. H., (1991). Adaptive approaches to stochastic programming. *Annals of Operations Research* 30: 45–62.
- McFadden, D. (1981). Econometric analysis of probabilistic choice. In C. Manski and D. McFadden (eds.), *Structural Analysis of Discrete Data*. Cambridge, MA: MIT Press.
- McKelvey, R. D. and Palfrey, T. R. (1995). Quantal response equilibria for normal-form games. *Games and Economic Behavior* 7: 6–38.
- Maddala, G. S. (1986). *Limited Dependent and Qualitative Variables in Econometrics*. Cambridge: Cambridge University Press.
- Mak, V., Rapoport, A., Gisches, E. J., and Han, J. (2014). Consumer purchase decisions under dynamic pricing: an experimental study on the effects of scarcity. *Manufacturing and Service Operations Management* 16: 425–38.
- Maynard Smith, J. (1964). Group selection and kin selection. *Nature* 201: 1145–7.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.
- Maynard Smith, J. (1984). Game theory and the evolution of behavior. *Behavioral and Brain Sciences* 7: 95–101.
- Maynard Smith, J. and Price, G. R. (1973). The logic of animal conflict. *Nature* 246: 15–18.
- Mengel, F. (2008). Matching structure and the cultural transmission of social norms. *Journal of Economic Behavior and Organization* 67: 608–23.
- Mengel, F. (2012). Learning across games. *Games and Economic Behavior* 74(2): 601–19.
- Miyasawa, K. (1961). On the convergence of learning processes in a 2×2 nonzero-person game. Princeton Econometric Research Program, Research Memorandum 33.
- Monderer, D. and Shapley, L. S. (1996). Fictitious play property for games with identical interests. *Journal of Economic Theory* 68: 258–65.
- Mookerjee, D. and Sopher, B. (1994). Learning behavior in an experimental matching pennies game. *Games and Economic Behavior* 7: 62–91.
- Mookerjee, D. and Sopher, B. (1997). Learning and decision costs in experimental constant sum games. *Games and Economic Behavior* 19: 97–132.
- Morgan, J. and Sefton, M. (1999). An empirical investigation of unprofitable games. Princeton University Working Paper.
- Myatt, D. P. and Wallace, C. (2005). The evolution of teams. In N. Gold (ed.), *Teamwork: An Interdisciplinary Perspective*. New York: Palgrave Macmillan, pp. 78–101.
- Myatt, D. P. and Wallace, C. (2008). When does one bad apple spoil the barrel? An evolutionary analysis of collective action. *Review of Economic Studies* 75(1): 499–527.
- Nagel, R. and Tang, F. F. (1998). An experimental study on the centipede game in normal form: an investigation on learning. *Journal of Mathematical Psychology* 42: 356–84.
- Nowak, M. A. (2006). Five rules for the evolution of cooperation. *Science* 314(5805): 1560–3.
- Nowak, M. A. and Sigmund, K. (1992). Tit for tat in heterogeneous populations. *Nature* 355, 250–3.
- Nowak, M. A. and Sigmund, K. (1993). A strategy of win–stay, lose–shift that outperforms tit for tat in prisoner’s dilemma. *Nature* 364: 56–8.
- Nowak, M. A. and Sigmund, K. (1998). Evolution of indirect reciprocity by image scoring. *Nature* 393: 573–77.
- Nowak, M. A. and Sigmund, K. (2005). Evolution of indirect reciprocity. *Nature* 437: 1291–8.
- Nowak, M. A., Tarnita, C. E., and Wilson, E. O. (2010). The evolution of eusociality. *Nature* 466(7310): 1057–62.
- Nyarko, Y. and Schotter, A. (2002). An experimental study of belief learning using elicited beliefs. *Econometrica* 70: 971–1005.
- Ockenfels, A. and Selten, R. (2005). Impulse balance theory and feedback in first price auctions. *Games and Economic Behavior* 51: 155–70.
- Ockenfels, A. and Selten, R. (2014). Impulse balance in the news vendor game. *Games and Economic Behavior* 86: 237–47.
- Oechssler, J. and Schipper, B. (2003). Can you guess the game you are playing? *Games and Economic Behavior* 43(1): 137–52.
- Ok, E. and Vega-Redondo, F. (2001). On the evolution of individualistic preferences: an incomplete information scenario. *Journal of Economic Theory* 97(2): 231–54.
- Ostrom, E. (2000). Collective action and the evolution of social norms. *Journal of Economic Perspectives* 14(3): 137–58.
- Panchanathan, K. and Boyd, R. (2004). Indirect reciprocity can stabilize cooperation without the second-order free rider problem. *Nature* 432: 499–502.
- Pinker, S. (2010). The cognitive niche: coevolution of intelligence, sociality, and language. *Proceed-*

- ings of the National Academy of Sciences 107: 8893–9.
- Posner, E. (2000). *Law and Social Norms*. Cambridge MA: Harvard University Press.
- Price, G. R. (1972). Extension of covariance selection mathematics. *Annals of Human Genetics* 35: 485–90.
- Rapoport, A. and Amaldoss, W. (2000). Mixed strategies and iterative elimination of strongly dominated strategies: an experimental investigation of states of knowledge. *Journal of Economic Behavior and Organization* 42: 483–521.
- Rapoport, A., Daniel, T. E., and Seale, D. A. (1998). Reinforcement-based adaptive learning in asymmetric two-person bargaining with incomplete information. *Experimental Economics* 1: 221–53.
- Rapoport, A. and Erev, I. (1998). Coordination, magic, and reinforcement learning in a market entry game. *Games and Economic Behavior* 23: 146–75.
- Rapoport, A., Seate, D. A., and Winter, E. (2000). An experimental study of coordination and learning in iterated two-market entry games. *Economic Theory* 16: 661–87.
- Rapoport, A., Seale, D. A., and Winter, E. (2002). Coordination and learning behavior in large groups with asymmetric players. *Games and Economic Behavior* 39: 111–36.
- Rapoport, A., Stein, W. E., Parco, J. E., and Nicholas, T. E. (2003). Equilibrium play and adaptive learning in a three-person centipede game. *Games and Economic Behavior* 43: 239–65.
- Rhode, P. and Stegeman, M. (1996). A comment on “Learning, mutation, and long-run equilibria in games.” *Econometrica* 64: 443–9.
- Richerson, P. J. and Boyd, R. (2004). *Not By Genes Alone*. Chicago: University of Chicago Press.
- Robinson, J. (1951). An iterative method of solving a game. *Annals of Mathematics* 54: 296–301.
- Robson, A. J. and Samuelson, L. (2011). The evolutionary foundations of preferences. In J. Benhabib, A. Bisin, and M. O. Jackson (eds.), *Handbook of Social Economics*. Amsterdam: North-Holland, pp. 221–310.
- Robson, A. J. and Vega-Redondo, F. (1996). Efficient equilibrium selection in evolutionary games with random matching. *Journal of Economic Theory* 70: 65–92.
- Roth, A. and Erev, I. (1995). Learning in extensive-form games: experimental data and simple dynamic models in the intermediate term. *Games and Economic Behavior* 8: 164–212.
- Rustichini, A. (1999). Optimal properties of stimulus-response learning models. *Games and Economic Behavior* 29: 244–73.
- Samuelson, L. (1997). *Evolutionary Games and Equilibrium Selection*. Cambridge, MA: MIT Press.
- Samuelson, L. (2001). Analogies, adaptation, and anomalies. *Journal of Economic Theory* 97: 320–66.
- Samuelson, L. and Zhang, J. (1992). Evolutionary stability in asymmetric games. *Journal of Economic Theory* 57: 363–91.
- Samuelson, W. F. and Bazerman, M. H. (1985). The winner’s curse in bilateral negotiations. In V. L. Smith (ed.), *Research in Experimental Economics*, 3. Greenwich, CT: JAI Press.
- Sandholm, W. H. (1998). Simple and clever decision rules for a model of evolution. *Economics Letters* 61: 165–70.
- Sandholm, W. H. (2010). *Population Games and Evolutionary Dynamics*. Cambridge, MA: MIT Press.
- Sarin, R. and Vahid, F. (1999). Payoff assessments without probabilities: a simple dynamic model of choice. *Games and Economic Behavior* 28: 294–309.
- Sarin, R. and Vahid, F. (2001). Predicting how people play games: A simple dynamic model of choice. *Games and Economic Behavior* 34: 104–22.
- Selten, R., Abbink, K., and Cox, R. (2005). Learning direction theory and the winner’s curse. *Experimental Economics* 8(1): 5–20.
- Selten, R. and Chmura, T. (2008). Stationary concepts for experimental 2×2 games. *American Economic Review* 98(3): 938–66.
- Selten, R. and Hammerstein, P. (1984). Gaps in Harley’s argument on evolutionarily stable learning rules and in the logic of tit for tat. *The Behavioral and Brain Sciences* 7: 47–70.
- Selten, R. and Stöcker, R. (1986). End behavior in sequences of finite prisoner’s dilemma supergames: a learning theory approach. *Journal of Economic Behavior and Organization* 7: 47–70.
- Selten, R., Thorsten, C., and Goerg, S. J. (2011). Stationary concepts for experimental 2×2 games: reply. *American Economic Review* 101(2): 1041–4.

- Shapley, L. S. (1964). Some topics in two-person games. In M. Dresher, L. S. Shapley, and A. W. Tucker (eds.), *Advances in Game Theory*. Princeton, NJ: Princeton University Press, pp. 1–28.
- Simon, C. P. and Blume, L. (1994). *Mathematics for Economists*. New York: W. W. Norton.
- Smith, M. J. (1982). *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.
- Sozou, P. D. (1998). On hyperbolic discounting and uncertain hazard rates. *Proceedings of the Royal Society of London, Series B* 265(1409): 2015–20.
- Stahl, D. O. (1999). Sophisticated learning and learning sophistication. University of Texas at Austin working paper.
- Stahl, D. O. (2000). Rule learning in symmetric normal-form games: Theory and evidence. *Games and Economic Behavior* 32: 105–138.
- Stahl, D. O. (2001). Population rule learning in symmetric normal-form games: theory and evidence. *Journal of Economic Behavior and Organization* 45: 19–35.
- Stahl, D. O. and Haruvy, E. (2002). Aspiration-based and reciprocity-based rules in learning dynamics for symmetric normal-form games. *Journal of Mathematical Psychology* 46(5): 531–53.
- Stahl, D. O. and Wilson, P. (1995). On players models of other players: theory and experimental evidence. *Games and Economic Behavior* 10: 218–54.
- Sterelny, K. (2011). From hominins to humans: how sapiens became behaviourally modern. *Philosophical Transactions of the Royal Society B* 366: 809–22.
- Sugden, R. (1986). *The Economics of Rights, Cooperation and Welfare*. Oxford: Basil Blackwell.
- Suppes, P. and Atkinson, R. (1960). *Markov Learning Models for Multiperson Interactions*. Stanford: Stanford University Press.
- Taylor, P. and Jonker, L. (1978). Evolutionary stable strategies and game dynamics. *Mathematical Biosciences* 40: 145–56.
- Thorndike, E. L. (1911). *Animal Intelligence*. New York: Macmillan.
- Tomasello, M. (1999). *The Cultural Origins of Human Cognition*. Cambridge, MA: Harvard University Press.
- Tooby, J. and Cosmides, L. (1992). The psychological foundations of culture. In J. Barkow, L. Cosmides, and J. Tooby (eds.), *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. New York: Oxford University Press.
- Trivers, R. L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology* 46: 35–57.
- Van Damme, E. (1987). *Stability and Perfection of Nash Equilibria*. Berlin: Springer-Verlag.
- Van Damme, E. (1991). *Stability and Perfection of Nash Equilibria*. Berlin: Springer-Verlag.
- Van Huyck, J., Battalio, R., and Beil, R. (1990). Tacit cooperation games, strategic uncertainty, and coordination failure. *American Economic Review* 80: 234–48.
- Van Huyck, J., Battalio, R., and Beil, R. (1991). Strategic uncertainty, equilibrium selection principles, and coordination failure in average opinion games. *Quarterly Journal of Economics* 106: 885–909.
- Van Huyck, J., Battalio, R., and Rankin, F. (2007). Selection dynamics and adaptive behavior without much information. *Economic Theory* 33(1): 53–65.
- Van Huyck, J., Cook, J., and Battalio, R. (1997). Adaptive behavior and coordination failure. *Journal of Economic Behavior and Organization* 32: 483–503.
- Vega-Redondo, F. (1996). *Evolution, Games, and Economic Behavior*. Oxford: Oxford University Press.
- Vega-Redondo, F. (1997). The evolution of Walrasian behaviour. *Econometrica* 65: 375–84.
- Vega-Redondo, F. (2003). *Economics and the Theory of Games*. Cambridge: Cambridge University Press.
- Vickers, G. T. and Cannings, C. (1987). On the definition of an evolutionarily stable strategy. *Journal of Theoretical Biology* 129: 349–53.
- Wallace, C. and Young, P. H. (2014). Stochastic evolutionary game dynamics. In H. P. Young and S. Zamir (eds.), *The Handbook of Game Theory, Volume IV*. Amsterdam: Elsevier. Forthcoming.
- Watson, J. (1993). A reputation refinement without equilibrium. *Econometrica* 61: 199–205.
- Weibull, J. W. (1995). *Evolutionary Game Theory*. Cambridge, MA: MIT Press.
- Wilcox, N. T. (2006). Theories of learning in games and heterogeneity bias. *Econometrica* 74(5): 1271–92.
- Wilson, E. O. (2012). *The Social Conquest of Earth*. New York: Liveright Publishing Corporation.
- Wynne-Edwards, V. C. (1962). *Animal Dispersion in Relation to Social Behavior*. Edinburgh: Oliver & Boyd.
- Young, H. P. (1993a). The evolution of conventions. *Econometrica* 61(1): 57–84.

- Young, H. P. (1993b). An evolutionary model of bargaining. *Journal of Economic Theory* 59(1): 145–68.
- Young, H. P. (1998). *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton, NJ: Princeton University Press.
- Young, H. P. (2004). *Strategic Learning and Its Limits*. Oxford: Oxford University Press.
- Young, H. P. (2006). Social dynamics: theory and applications. In K. Judd and L. Tesfatsion (eds.), *Handbook of Computational Economics, Volume II*. Amsterdam: North Holland.
- Young, H. P. (2008). Social norms. In S. N. Durlauf and L. E. Blume (eds.), *The New Palgrave Dictionary of Economics*. 2nd edition. London: Macmillan.
- Young, H. P. and Burke, M. A. (2001). Competition and custom in economic contracts: a case study of Illinois agriculture. *American Economic Review* 91: 559–73.
- Zeeman, E. (1980). Population dynamics from game theory. In Z. Nitecki and C. Robinson (eds.), *Global Theory of Dynamical Systems* (Lecture Notes in Math 819). Berlin: Springer-Verlag.

PART 6

Emotions

INTRODUCTION TO PART 6

In economics, the standard practice is to focus on *emotionless deliberation* and assume that economic agents make cool, calculated, decisions in constructing trade-offs between alternative choices. In contrast, evidence suggests that humans operate along the entire range of the emotional spectrum, rather than at the cool end only. Emotions are central to human behavior and affect economic outcomes in a fundamental way. Until fairly recently, economic theorists had little to say about emotions, and researchers working in the area of emotions, little to say about the effect of emotions on economic variables. However, in recent years, there has been an impressive spurt of research activity in this area.

One may give the following useful taxonomy of emotions (Elster, 1998). *Social emotions* include anger, hatred, guilt, shame, pride, admiration, and liking. *Counterfactual emotions* are generated by thinking about unrealized possibilities, e.g., regret, rejoicing, disappointment, and elation. *Anticipatory emotions* are caused by anxiety about what may happen, e.g., fear and hope. *Realized emotions* are caused by events that have already been realized in the past, e.g., grief and joy. *Material emotions* are caused by the possession of material or non-material objects by others, e.g., envy, malice, indignation, and jealousy. There are other kinds of emotions that may not fall into any of these categories, e.g., contempt, disgust, and romantic love.

In this part, we are mainly interested in the study of a limited number of emotions and their impact on economic behavior.¹ We also consider the interaction between emotions and cognition in shaping economic decisions. Indeed, as Elster (1998, p. 73) writes: “The more urgent task is to understand how emotions interact with other motivations to produce behavior.”

There is a range of interesting questions that we either omit, or which offer a limited discussion; however, these questions are well discussed in other sources.² For instance, are emotions culture-specific or universal? What hormonal or physiological changes are associated with emotions? What is the valence, or location on a pleasure–pain scale, of different emotional states? Can emotions be called upon/avoided by individuals in a deliberate and rational manner, or are they passively undergone? Can we learn to avoid cues that trigger different emotional states? Can individuals successfully induce emotions in others?

We have already considered, elsewhere in the book, material that could have also been well situated in this chapter. This includes (i) formal models of emotions such as guilt, anger, and disappointment using *psychological game theory* (Part 4 of the book),³ (ii) models of *disappointment aversion* and *regret* (Part 1 of the book), and, (iii) issues of *self-control* that arise from *present-*

¹ There is a corresponding realization among emotion theorists about the importance of economic models. For instance, Rolls (2014) devotes a major part of his Chapter 9 to economic models, albeit the main focus is on neuroeconomics.

² The interested reader may consult Frank (1988), Elster (1998), Loewenstein (2000), Goldie (2009), Winter (2014), and Rolls (2014).

³ For a discussion of guilt inculcation in children by parents that does not rely on the machinery of psychological game theory, see Becker (1996) and Frank (1988), although they invoke very different mechanisms to each other.

biased preferences (Part 3 of the book). In order to avoid repetition, we omit a discussion of these issues here, and focus on the remaining issues in the role of emotions in economics.⁴

Chapter 17 gives an introduction to some components of emotions and physiological states (moods, physical pain, hunger, thirst) that are of interest in economic models; collectively these are known as *visceral factors*. We briefly consider the question of the rationality of emotions. In Part 1, we considered the role of *anticipated emotions* such as regret and disappointment that might be experienced at some future time, on account of decisions made today. Here, in Chapter 17, we introduce *anticipatory emotions* experienced at the time of making decisions, such as anxiety and dread. Part 1 also considered various theories of risk. However, risk may also bring about feelings of dread, anxiety, and apprehension. Thus, it would appear that emotions ought to be inextricably linked to explanations of human behavior under risk. We explain the *risk as feelings hypothesis*, that ought to be taken more seriously by economists.

Visceral factors are typically activated by external cues. We develop a model of *habit formation* in the presence of *stochastic cues* and discuss the role of active cue management. Anticipation of future pain or rewards can create anxiety or elation, thus, we also incorporate the role of *anticipal utility* into a simple life-cycle model. We develop a *preference representation of anxiety* and apply it to a problem of portfolio choice in a life-cycle model of consumption. It turns out that risky assets, whose uncertain returns create anxiety, may command an *anxiety premium* that could also potentially explain the equity-premium puzzle.

When predicted future tastes differ from their actual realization, individuals are said to suffer from *projection bias*. In this case, individuals might make time-inconsistent choices, and may end up making decisions in a hot state that they regret later. This explains the role of public policy and legal provisions that provide for various cooling-off periods, be they for the purchase of consumer durables or applying for marriage/divorce.

We discuss the Gul–Pesendorfer framework of *temptation preferences* that explains temptation and self-control problems in the classical revealed preference framework. In this framework, one response to temptation is to remove a tempting option from the menu of choices. However, in the presence of uncertainty about the future, there are costs and benefits from this option. For instance, removing a currently tempting option from the menu may improve current savings, yet an adverse shock to marginal utility in the future may reduce the future enjoyment of these savings.

In the final section of this chapter we discuss the emerging literature on *happiness economics*. Unlike many other areas in economics, the evidence relies on self-reported measures of well-being, health, and life satisfaction. However, there is a strong correlation between various measures of well-being and these, in turn, successfully predict job quits, suicides, productivity, and probability of a divorce among couples. We examine a celebrated result in this literature, the *Easterlin paradox*. The cross-sectional evidence in any country shows that richer people are happier. However, the time series evidence shows that growth in income is not associated with higher self-reported measures of well-being. A leading explanation is that as incomes increase over time, the reference point of individuals, as captured by their aspiration levels, moves upwards. Thus, an increase in income over time is not associated with greater happiness. In recent years, the data behind the Easterlin paradox has been increasingly questioned based on a longer time series and a more representative sample of poor countries.

⁴ Instructors who wish to teach a specialized course in emotions in economics or a behavioral course with special emphasis on emotions may wish to combine these topics from different parts of the book.

Chapter 18 considers the interaction between emotions and cognition. This interaction has evolutionary origins. The cognitive system is associated with the development of the prefrontal cortex, 150,000 years ago. However, the cognitive system was built on top of an already existing emotional system that we share with other mammals. The result is that the objectives of the emotional and the cognitive systems can sometimes be in conflict, and at other times, in harmony. We also employ a System 1 (quick, reactive emotional system) and System 2 (slow, deliberative cognitive system) distinction in Part 7. However, here we are interested in examining a different set of issues.

We begin by presenting a promising model of the interaction between the two systems that can serve as a template for the development of future models. This model has the advantage of endogenously deriving several behavioral features of interest to economists such as reference dependence, loss aversion, social preferences, and hyperbolic discounting.

In a formal model of *dual selves*, we show how the cognitive system can control the emotional system through the exercise of costly willpower. We derive the optimal solution in the presence and in the absence of a commitment technology.

People often try to discover themselves through their actions. For instance, charitable giving may be a self-signal that the individual has empathy and compassion for others. Withstanding extreme stress and successfully completing a timed task may self-signal willpower to an individual. These issues can be modeled in several ways. A radical proposal is to model the interaction between the dual systems using classical game theoretic asymmetric information models. However, this leaves open the question of how contracts are enforced between the emotional and the cognitive systems, and what does a participation constraint imply in this special case?

In the final two sections, we consider the *strategic transmission of information* in a multiple-selves model. In these models, if the current self chooses to acquire some information, it immediately becomes available to all future selves, who can use this information to the detriment of the current self. In this case, the current self may choose to be strategically uninformed. For instance, individuals vastly overestimate the effects of smoking on lung cancer, yet typically are not proactive in accessing easily available information. There can be other equally plausible explanations for this phenomenon. For instance, individuals may procrastinate in the presence of hyperbolic discounting. However, here we study a different set of explanations that are based on *imperfect willpower*, *imperfect recall*, and *motivated cognition* (i.e., managing the degree to which subsequent selves can recall past events).

CHAPTER 17

Emotions and Human Behavior

17.1 Introduction

In this chapter, we motivate the use of emotions to explain human behavior. In contrast to the cold, emotionless, deliberations undertaken by Homo economicus, the introduction of emotions provides a more satisfactory explanation of many economic phenomena. A wealth of empirical evidence provides a compelling case for taking emotions more seriously in all areas of economics.

In Section 17.2, we focus on intense physiological states such as hunger, moods, thirst, physical pain that are collectively referred to as *visceral influences* (Loewenstein, 1996). Visceral influences can induce individuals to take actions that are often contrary to their own self-interests. Yet individuals often underestimate the effect of visceral influences on their past, current, and future behavior. The usual practice in economics is to consider emotions arising in the future on account of current decisions, as in regret theory and disappointment aversion; these are *anticipated emotions*. A strong case has been made for considering the role of emotions at the time of making decisions, i.e., *anticipatory emotions*, such as anxiety and dread (Loewenstein, 2000).

The issue of the rationality of emotions is an interesting one (Damasio, 1994). Individuals who have suffered damage to the emotional centres of their brains are unable to experience emotional states associated with fear. When faced with risky situations that entail occasional large losses and overall negative expected values, these individuals make extremely risky decisions. Damasio and several others have argued that the absence of emotions is detrimental to sound cognitive decisions, hence, on net, emotions must be rational. Emotions need not always enable us to take socially optimal actions although emotions such as compassion and empathy often do lead to prosocial actions at an individual level. For instance, our insensitivity to many global problems such as global warming, and poverty and hunger in poor countries may be caused by a problem of insufficient emotions (Loewenstein, 2010).

In the final part of Section 17.2, we consider the role of anticipatory emotions in decisions under risk based on the idea of *risk as feelings* (Loewenstein et al., 2001). In many risky situations, individuals do not behave in an entirely consequentialist manner (i.e., they do not pay sufficient attention to probabilities and outcomes). We consider how emotional states may influence the types of decisions made under risky environments. This gives an alternative, affective, account of decision making under risk that deserves to be taken more seriously by economists of all persuasions.

Visceral influences are often activated by external stimuli, which we refer to as *cues*. Traditional habit formation models do not consider the role of cues (Becker and Murphy, 1988). In Section 17.3, we consider the role of stochastic cues in models of habit formation (Laibson, 2001;

Bernheim and Rangel, 2004). The nature of the solution under the cues and the no-cues models is very similar. If the stock of habits exceeds a threshold, then the individual optimally engages in the harmful activity, otherwise not. In traditional habit formation models, the main upshot of this feature is that once rationally addicted, individuals stay addicted. However, in the cues model, the harmful activity is dictated by the observation of a cue. In order to wean themselves off addiction, individuals may then engage in active cue-management, while self-interested firms may try to counter this by making consumption-inducing cues more salient. Models of addiction are also considered in subsequent Sections 17.6, 17.7, and in applications of present-biased preferences in Part 3 of the book that readers interested in a fuller treatment should consult.

We often feel anxious about an uncertain future. Will my pension contributions be sufficient for my retired life? Will early retirement turn out to be a good decision? Should I buy or sell stocks in the current economic environment? Anticipation of a stressful event can create anxiety that in some cases is worse than the event itself, e.g., an impending visit to a dentist. Such feelings of anxiety are caused by *anticipation* of future events, and they are *aversive*. Future events need not, however, always be unpleasant. We might savor a pleasurable experience such as waiting for a dessert at the end of a meal. Or, we might be excited about a forthcoming pleasurable vacation.

For both pleasant and unpleasant events, individuals often exhibit a negative discount rate. This typically takes the form of postponing a pleasant experience (dessert at the end of a meal) and preponing an unpleasant one (taking out an aching and doomed tooth sooner rather than later). This evidence is contrary to the predictions of the discounted utility model, which does not take account of anticipatory utility and anticipatory disutility from future events. Section 17.4 shows how such anticipatory feelings can be incorporated into a formal model (Loewenstein, 1987).

It is natural to suppose that anxiety about future events may be associated with uncertainty about the future. In Section 17.5, we consider a utility representation of preferences under anxiety (Caplin and Leahy, 2001). We use this preference representation to derive equilibrium prices of assets whose returns are ex-ante random. The randomness of returns creates feelings of anxiety. When investors exhibit anxiety, risky stocks are priced even higher relative to the case of no-anxiety. This provides one possible explanation for the *equity premium puzzle* that does not rely on loss-aversion. Furthermore, the price of riskless assets in the presence of anxiety is relatively higher because they provide a “peace of mind function” for the investor that stocks with fluctuating returns cannot.

In the absence of any random shocks, the typical assumption in economics is that individuals are fully aware of their future tastes. When the predictions of future tastes differ from their actual realization, despite no inherent randomness in the model, then the individual is said to make a *projection bias*. Section 17.6 considers issues of projection bias (Loewenstein et al., 2003). For instance, among patients who await kidney transplants, those who are successful (unsuccessful) in getting a transplant overestimate (underestimate) their future quality of life (Smith et al., 2008). A distinguishing feature of projection bias is that individuals use their current state disproportionately to predict their future states. Hungry shoppers choose far more groceries for future consumption, relative to satiated shoppers, even when they have the same future needs (Read and van Leeuwen, 1998).

Individuals who have projection bias make *time inconsistent choices*, without being aware of it. Under the influence of projection bias, in a *hot state*, individuals might, for instance, (i) end up buying goods that they regret later, and (ii) end up marrying/divorcing without sufficient deliberation. In response to these concerns, (i) consumer regulation provides a 21 day cooling-off period during which goods can be returned, provided they are in a satisfactory state, and (ii)

the law typically requires a cooling-off period before which applications for marriage and divorce cannot be filed. Projection bias also provides an alternative explanation for why individuals may get addicted. First, they might underestimate the (negative) effect of current consumption on future utility. Second, they might under-appreciate the formation of habit, i.e., addiction to the harmful product. We give two formal applications of projection bias. The first is to a two-period model of habit formation, and the second is to purchase of consumer durables.

What does it mean to say that one alternative is more tempting relative to another? Can one provide an axiomatization of preferences that incorporates some plausible notion of *temptation*? Consider a two-period model. In the first period, the individual chooses between a set of *menus*. In the second period, the individual chooses an element from the menu that was chosen in the first period. Preferences exhibit *temptation* if the individual prefers to exclude certain menus from first period choice; under the typical preferences specified in neoclassical economics, more options never hurt. If exclusion of the temptation is not possible, then the individual may either be able to exercise self-control in the second period or succumb to the temptation. We give a utility representation of temptation preferences in Section 17.7 (Gul and Pesendorfer, 2001; Dekel et al., 2001). We also show that temptation preferences can also arise in a model of competing dual selves, when there is some uncertainty about the eventual winner in a contest between them (Bénabou and Pycia, 2002).

When individuals have self-control problems, such as those arising from temptations, the standard prescription in a range of models is to commit to eliminate options that lead to such temptations. Section 17.8 considers an important trade-off associated with the removal of temptations in a model of multiple-selves (Amador et al., 2006). Eliminating a subset of the options has the benefit of reducing/eliminating temptations, but at the cost of reduced flexibility, especially when new information is likely to arrive in the future. For instance, in the context of intertemporal choice, a reduction in temptations might improve the savings rate of an individual. However, if the future marginal utility turns out to be unexpectedly low, then the improvement in savings might not have been as worthwhile. In the context of this framework, we derive the justification for a minimum savings rule when current selves have asymmetric information about the subsequent selves. Within the two type case, we also derive conditions under which the first best can be implemented.

Modern economics typically relies on the notion of *decision utility*—the utility experienced at the time of the decision. An older tradition in economics emphasizes *experienced utility*—the hedonic experience associated with an outcome (Bentham, 1798). The modern approach is based on checking if preferences are consistent with the axioms of rational choice—this is encapsulated in the revealed preference approach. However, in many cases, the well-being of individuals is influenced by outcomes that they cannot directly choose. Examples include inflation, unemployment, income inequality, environmental degradation, crime, health shocks, accidents, and exogenous income shocks. In these cases, revealed preference arguments cannot be used to infer the well-being of individuals. It would be bizarre, if one is willing to take experienced utility seriously, to argue that the subject matter of economics should restrict itself to those economic outcomes that can be handled via the revealed preference approach only.

Section 17.9 considers an alternative approach to well-being, based on *experienced utility*. In this approach, individuals are asked directly to report their well-being, life satisfaction, and health status. Clearly, self-reported survey questions that report experienced utility are subjected to biases, such as *duration neglect*, and *peak-end evaluation* (Kahneman et al., 1997). However, subjective measures of well-being are consistent with the actual choices for the vast majority of subjects and these measures are the best predictors of choices (Benjamin et al., 2012). Such

subjective measures successfully predict a range of phenomena such as job quits, probability of suicides, productivity and probability of a future divorce among couples. Despite the drawbacks of self-reported measures of well-being, it may be still preferable to use them rather than have experts guess the well-being of individuals. Furthermore, survey questions may catch people in a deliberative mood rather than an emotional state, hence, self-reported measures may have normative value (Benjamin et al., 2014).

Evidence from many countries showed that at any point in time, richer individuals are happier. However, over time, while GDP growth takes place, mean self-reported happiness stays almost unchanged. This has come to be known as the *Easterlin paradox* (Easterlin, 1974). The main explanation offered for the paradox is that individuals adapt to the higher level of income, and simply raise their aspiration levels. In this view, self-reported happiness is a measure of the difference between actual incomes and aspiration incomes (Easterlin 1974, 1995).

This evidence has been contested in recent years in a series of papers, which argue that mean happiness is increasing in the levels of income over time, hence, there is no Easterlin paradox (Deaton, 2008; Stevenson and Wolfers, 2008, 2013). The main reasons for the difference in conclusions is newer datasets (Gallup World Poll rather than the World Values Survey, which formed the basis of earlier results) that provide data on a more inclusive set of poorer countries, and use a more representative sample. Easterlin has responded with the argument that while the paradox might not hold in the short run, it does hold in the long run (Easterlin et al., 2010). We end Section 17.9 with a brief note on the correlates of self-reported well-being, such as inflation, unemployment, health, income inequality, and democracy.

17.2 Visceral influences and the rationality of emotions

17.2.1 *Visceral influences*

One manifestation of emotions that has an important influence on human decision making is intense physiological states that are lumped under the term *visceral factors/influences*. Examples include hunger, thirst, moods, physical pain, and craving for a drug. So, for instance, consider the visceral factor of sleep deprivation in its extreme form. Individuals do not make a rational decision to fall asleep behind the wheel of their car, but some people do, with disastrous consequences. Overeating, compulsive shopping, and drug addiction are other examples where strong visceral influences can lead to suboptimal behavior.

Loewenstein (1996) has drawn attention to the role of visceral influences in influencing choices. Two important features of visceral influences are the following.

1. *Current visceral factors can have a “disproportionate” influence on current behavior:* When visceral influences are weak, individuals can control them better. However, when they are “intense,” the individual gives almost exclusive attention to these influences at the expense of other possible objectives. Under extreme visceral influences, the individual may behave as if the discount factor were zero. Hence, such factors might drive a wedge between the self-interested optimum of an individual and actual behavior.
Under the influence of visceral factors, individuals can often be aware of the wedge between their actions and best interests. However, they nevertheless persist in taking actions that are contrary to their self-interest. For instance, in a visceral state of anger, an individual, in the heat of the moment, could say something that is deeply regretted later, and lead to a poor outcome.

An extreme example of visceral influences is the behavior of individuals who have various *phobias* that they fully know are non-threatening but are still unable to use rational judgment to counter them. Police and military interrogators use extreme visceral influences such as pain, hunger, or sleep deprivation to elicit confessions. Or, as a strike deadline approaches, labor negotiations often go “round the clock,” evoking the extreme visceral influence of sleep deprivation that is more likely to elicit concessions. Both sides, perhaps due to mutual awareness of a potential bargaining impasse, agree to this bargaining structure.

In conclusion, unlike the assumptions of neoclassical economic theory, on many occasions individuals do not operate at the *cool end* of the spectrum of human emotions. Despite a degree of self-awareness of the problem, they can operate at the *hot end* of the spectrum, with detrimental consequences to their self-interest.

2. *Individuals typically underestimate the importance of visceral factors in determining their past, current, and future behavior:* Loewenstein (1996) reports the following empirical result. Prior to the onset of labor, pregnant women were asked to make a non-binding decision on whether they will require an anaesthetic during labor. A majority were optimistic that they would not need an anaesthetic. However, acute pain is often an extreme visceral influence. Once labor set in, a high proportion of the women reversed their earlier decision. Hence, individuals can underestimate the effect of visceral influences on future behavior.

The “endowment effect” can also be thought of as evoking the visceral factor of attachment. Individuals underpredict the future endowment effect. Individuals also often forget or underreport the visceral influence of pain in the past; memories of pain seem qualitatively different from other sorts of memories.

These concerns are contrary to the behavior of individuals in models where they *rationally decide* to be addicts, in full knowledge of all future consequences, as in Becker and Murphy (1988). By contrast, visceral influences explain drug addiction by a serious error of judgment on the part of individuals, who underestimate the pain of withdrawal and craving for the drug, caused by the visceral activity of addiction.

Visceral factors are not *tastes* or *preferences*, as understood in conventional economic theory. Visceral influences arise from external stimulation or deprivation, hence, they may be more volatile than tastes. As Loewenstein (1996, p. 273) puts it: “In contrast, changes in preferences are caused by slow experience and reflection, are typically not anticipated, and do not imply a permanent change in behavior.” Furthermore, Loewenstein cites evidence to show that tastes and visceral factors draw on different neurophysiological mechanisms. Visceral factors are likely to have a common neurophysiological mechanism. Evidence for this assertion comes from the observation that behavior contrary to one’s self-interest, arising from a range of visceral influences, can be mitigated by using a single drug, Fluoxetine.

The origin of visceral factors are most likely evolutionary. For instance, the fitness of animals who respond to hunger by focusing exclusively on the search for food, might have been higher. However, such behavior could also be a double edged sword. For instance, moderate fear may lead to actions that are conducive for survival, but extreme fear can petrify an individual and hinder rational action.

Under the influence of visceral factors, individuals focus more on their needs and less on others. Interrogators have long known that under appropriate visceral influences, e.g., extreme torture, otherwise sincere individuals are likely to betray collaborators, friends, and family.

17.2.2 *Rationality and emotions*

Is it rational to have emotions? One possibility is that our judgment may be clouded in the presence of emotions, hence, we may make suboptimal decisions (Baumeister et al., 1994; Baumeister et al., 2007; Frijda and Mesquita, 1994). For instance, one may make a reckless purchase in a hot state of mind and regret it later, once the initial emotion subsides. The presence of cooling-off periods, enshrined in law in many countries, appear to be a regulatory response to this possibility.

On the other hand, using data from neurobiological studies, Damasio (1994) argues that emotions aid us in taking better decisions. For supporting findings and the evidence, see Bechara et al., (1994, 1996, 1997, 1999) and Damasio (1996). Using data from individuals who had suffered damage to the emotional regions of their brains, e.g., damage to the ventromedial prefrontal region,¹ Damasio draws several conclusions from the behavior of individuals who suffer such kind of brain damage; we summarize some of these conclusions, below.

1. These individuals are less likely to conform to social norms, which might lead them to suffer social sanctions.
2. These individuals may procrastinate more in social situations because they are less likely to suffer from the guilt and embarrassment of letting others down.
3. These individuals may do worse in decision tasks. For instance, in one experiment, individuals had to choose from four possible decks of cards, A, B, C, and D. Decks A and B contained high payments and occasional large losses, so that the expected temporal pay-offs were negative. Decks C and D contained low payments but positive expected temporal payoffs. Observing the data over time, it was possible to learn that it is a dominated strategy to choose decks A and B. Yet, those with damage to the emotional regions of the brain continued to choose decks A and B, while those who did not suffer from the damage learned to pick decks C and D. A possible interpretation of these findings is that individuals with brain damage did know that decks A and B are associated with occasional large losses, but they lack a fear response that would have led to the avoidance of decks A and B. Indeed, over time, such individuals could go bankrupt. Similar results are expected from psychopathic criminals, who are insensitive to the future consequences of their actions.

These results, however, cut both ways. For instance, one could design an experiment such that, while decks A and B have large occasional losses, the expected temporal value of choosing these decks is higher as compared to choosing decks C and D. In this case, those with a damaged ventromedial prefrontal region may earn a higher long-run return.

Smith and Dickhaut (2005) find that they can make better predictions of prices in English and Dutch auctions when they take account of the heart rates of participants, just prior to making choices. An interesting finding is that emotions play a role in price determination in Dutch auctions but not in English auctions. Hence, there might be important complementarities between emotions and the determination of economic outcomes under some institutions, but not others.

When choices are cognitively challenging, individuals may often make choices based on physiological signals, or emotional states, associated with outcomes. This is also sometimes known as the *somatic markers* hypothesis. For instance, a coiled object on the side of the street may initiate immediate fear, and the fight-or-flight response, bypassing the cognitive

¹ This region of the brain is found to be active when individuals make choices; see Glimcher et al. (2009). Individuals with damage to the ventromedial prefrontal region otherwise appear to be normal in respect of memory, intelligence, and language comprehension.

system. Avoidance of decks A and B after a period of learning may also be based on the fear of losing a large amount, while patients with ventromedial prefrontal damage exhibit no such fear.

The debate on the rationality of emotions, despite the influential work of Damasio and his colleagues, is not a settled one. For instance, anger may be a suitable emotion to get a larger share of the cake when aggression is required to determine the division of the cake. However, more irascible individuals may be shunned by others in subsequent divisions of the cake and may lose out in the long run. The economic models that we consider in subsequent sections, allow for a more explicit consideration of these opposing forces.

Are emotional states induced voluntarily or involuntarily? The somatic markers hypothesis is based on the assumption that emotional states are induced involuntarily. Sinaceur and Tiedens (2006) provide evidence to show that individuals may also make use of emotional states in a strategic manner to gain advantage in bargaining situations. Meshulam et al. (2012) provide evidence that suggests many individuals can induce emotional states at will. In their experiments, players play a modified dictator game in which they play the role of recipients who receive low, medium, and high offers. Their emotional response, in terms of skin conductance response (SCR) and heart rates, is measured, while they receive the different offers. These measurements can determine if their emotional state is angry, calm, or happy.

Subjects are given incentives to exhibit various emotional states, depending on the offers received. For instance, when they receive a low offer, they are given a bonus if their emotional response is angry. Likewise they are given a bonus to demonstrate a happy state in the case of a high offer. In addition, at the end of the experiment, subjects fill in a questionnaire in which they are given two vignettes in which an individual suffers harm from another's negligence. Subjects were then asked to rate their punitiveness towards the negligent individual on a 1–7 scale. This is used to classify individuals into those who are naturally angry, calm, and happy.

Figure 17.1 shows the main results. More angry subjects (as measured by scores on the 1–7 response to the two vignettes) are also angrier when they receive low offers (as measured by SCR




Figure 17.1 Skin conductance response (SCR) amplitude and heart rate (HR), 0–12 s after offer presentation onset. Bars represent mean amplitude of SCR for high and low offers of participants in each condition. Lines represent mean HR during the presentation of high and low offers. Error bars: $\pm$ SEM.

Source: "Rational emotions," Meir Meshulam, Eyal Winter, Gershon Ben-Shakhar, and Itzhak Aharon, *Social Neuroscience*, 7(1), (2012). Reprinted by permission of Taylor & Francis Ltd.

and the heart rate). Since individuals are responding to the incentive effect of bonuses, if they demonstrate the required emotional state, the authors conclude that emotions are voluntary and rational. There is an asymmetry between angry and happy individuals; the latter are not able to fully muster the necessary emotions to take advantage of the bonus on offer. This finding could be a useful candidate for further exploration; the authors speculate that prospect theory may be a possible explanation.

Formal theoretical models of the rationality of emotions are in their infancy. Winter et al. (2014) propose that players choose their strategies and their mental states simultaneously. They define the notion of a *mental equilibrium* in which roughly, (i) the chosen strategies of the players are a Nash equilibrium for a given mental state, and (ii) the mental state of each player is a best response to the mental states of other players. The epistemic foundations of a Nash equilibrium show just how strong the underlying conditions are and the outcomes of many games are not consistent with equilibrium concepts (see Part 4). The cognitive requirements for a mental equilibrium are likely to be even stronger. Ultimately, the usefulness of this line of research, particularly for experimental results, is an open question. Perhaps the framework of a mental equilibrium may have greater usefulness in understanding the development of emotions in certain contexts over a longer, evolutionary, time period. Winter (2014) makes a case for the rationality of emotions that follows in large part from the notion of a mental equilibrium. In the mental equilibrium framework, being angry is necessarily a rational response that is codified as a part of the equilibrium. For an alternative approach that is able to distinguish between anger caused by the intentionality (or lack of it) of the behavior of others using a psychological game theory approach, see the interesting paper by Battigalli et al. (2015).

Loewenstein (2010) proposes that many global problems may be caused by the problem of insufficient emotion. For instance, we may not feel sufficiently emotionally distressed about global warming, or hunger, poverty, and disease in the poorer countries, hence, we react inadequately. Even during wars and genocides, humans appear to feel inadequate emotions for the victims who are often dehumanized through propaganda that paints them as outsiders or dangerous to one's own world view. Statistics of lives lost through wars and genocides often move us insufficiently. Insufficient emotions could be an evolutionary adaptation. For instance, they prevent us from being immobilized by the large number of emotionally disturbing events around us. Perhaps for this reason, we are emotionally more sensitive to next of kin or close members of our own social group, a phenomenon known as *parochial altruism*.

Vividness of events plays an important role in our emotional arousal. For instance, we are less moved by statistics of accidents, poverty, child labor, and disease, but more moved by graphic stories of individuals who have suffered, particularly women and children. Good storytellers often weave personal, individual narratives within complex social situations, for greater effect. For similar reasons, simple cost–benefit analyses, say, associated with global warming, or the cost of saving a life in a poor country, can sometimes be an inadequate impetus for many of us to be proactive.

There is a contrast between the work of economists and psychologists in terms of the role played by emotions in guiding human behavior. Economists have traditionally incorporated *anticipated emotions*, such as regret and disappointment (see Part I of the book). These are emotions that are not experienced at the time of making decisions, but are experienced in the future when the outcomes of current decisions materialize. In contrast, psychologists have stressed the role of *anticipatory emotions*. These are emotions such as fear, dread, and anxiety that are experienced at the time of decision making. Loewenstein (2000) and Loewenstein et al. (2001)

make a compelling case for incorporating anticipatory emotions into economics. While the ideas apply quite generally, the thrust of our discussion below is to consider the effects on decisions under risk.

We now outline some of the main ideas from the seminal work of Loewenstein et al. (2001). Each of the assertions below are backed by good evidence; the interested reader should consult Loewenstein et al. (2001) for an exhaustive list of about 200 references that we omit.

Classical decision theories such as expected utility, rank dependent utility, and prospect theory are *consequentialist* in the sense that decision makers are assumed to be solely influenced by the probabilities and outcomes. In a series of papers, Paul Slovic and co-authors (see, e.g., Slovic et al., 2002) have highlighted anticipatory emotions such as dread and risk of the unknown in determining the outcomes of decisions under risk. By contrast, consequentialist factors such as probabilities and outcomes play little role in the outcomes. For instance, some individuals may experience dread and anxiety at the mere thought of air travel, while being perfectly comfortable with travel by car. However, the probability of car accidents, per mile travelled, is much higher. Phobias are an extreme form of these kinds of anxieties. High-risk activities, even if associated with a positive expected value, may be associated with negative emotions, hence, they may be avoided, which is consistent with the somatic markers hypothesis.

Empirical evidence suggests that jurors often make decisions based on their gut feelings. The difficulty in making decisions can evoke negative emotions even when all possible consequences are positive. People in good moods often make optimistic choices, while people in bad moods make pessimistic choices. When people are exposed to “sad” newspaper articles, their estimates of risk for a range of potential causes of death are higher, relative to people who read “happy” newspaper articles. Under induced anxiety, people choose low risk and low reward options.

The emotional and cognitive pathways in the brain are different, hence, it is not surprising that decisions made in the presence of emotional factors are different from the decisions made under purely cognitive considerations. Loewenstein et al. (2001) highlight the transmission channels behind these differences as well as the differences in outcomes under the two. Their approach to decision making under risk, which they term as *risk as feelings theory*, is a mixture of consequentialist and non-consequentialist reasoning. The outcomes under this theory can differ substantially from the predictions of standard decision theory in economics, particularly when anticipatory emotional reactions are strong, relative to the cognitive evaluations.

A schematic outline of the risk as feelings thesis is shown in Figure 17.2. Following the bottom row of arrows from left to right in the figure, behavior under risk can manifest independently of the associated probabilities and outcomes. Indeed the top two channels in the figure show that behavior can arise under a combination of outcomes and emotions (ignoring probabilities) or probabilities and emotions (ignoring outcomes). Each of these channels makes a prediction that cannot be made under traditional models of risk in economics.

The ability of individuals to form vivid images associated with risky outcomes, an ability that differs among individuals, can affect risky choices. For instance, individuals could underinsure themselves from risks that cannot be vividly imagined. Vividness may be enhanced by knowing someone who has suffered from the negative consequences of a risky situation. Thus, as noted above, individuals are less affected by statistics associated with risk and more affected by personal human stories of suffering that evoke greater emotion. However, inducing too much anxiety by, say, public information, may lead to avoidance of the relevant public information, forgoing the necessary risk mitigation activities, by some individuals. For instance, public information on the importance of self-examination to allow for early detection of breast cancer has led many women




Figure 17.2 Schematic description of the “risk as feelings” thesis.

Source: Copyright © 2001 by the American Psychological Association. Reproduced with permission. “Risk as feelings.” Loewenstein, George F., Weber, Eike U., Hsee, Christopher K., and Welch, Ned. *Psychological Bulletin* 127(2), Mar 2001, 267–86. The use of APA information does not imply endorsement by APA.

to avoid the necessary self-examination, on account of the presence of anticipatory anxiety. Many empirical studies show that women are more risk averse than men. Women are also known to engage in better and more vivid imagery as compared to men. Perhaps women can feel more negative emotions associated with losses, hence, they appear to be relatively more risk averse.

Decision makers are often not very sensitive to variations in probabilities within a broad set of values. Hence, contrary to consequentialist models of risk, they typically do not react as much to changes in probability from, say, 1 in 100,000, to 1 in 1,000,000. In a large number of experiments in psychology, the probability of receiving an impending shock had little effect on anticipatory anxiety. By contrast, in the presence of pallid outcomes, say, losing \$20, the degree of sensitivity to probabilities is relatively higher. These sorts of considerations might explain why the *probability weighting function* (see Part 1 of the book) is relatively flat for mid-range probabilities and why the degree of flatness may depend on anticipatory anxiety.

One explanation for the high take-up rates of expensive warranties is that people do not pay adequate attention to the probability of breakdowns. They appear to think immediately of the “peace of mind” associated with the warranty, which is an anticipatory emotion. Indeed, salesmen rarely highlight probabilities of breakdown, and emphasize the convenience and peace of mind aspects of owning the warranty. Lottery sellers similarly appeal to anticipatory emotions when they use slogans such as “buy a dream.” We did not consider any of these issues in Part 1, neither has the economics profession given adequate attention to this line of research; clearly, it deserves more attention and respect.

17.3 Cue-conditioned behavior and habit formation

Visceral influences are typically activated by external stimuli or *cues* that heighten certain desires. The desire for consumption, in particular, can be enhanced by a number of cues; we will refer

to this as *cue-conditioned consumption*. For instance, the smell of freshly made cookies might create a craving for them. Cue-conditioned consumption, like visceral influences, is not specific to humans. In a famous experiment, conducted more than a century ago, Pavlov (1904) found that if bells are rung at the same time as food is given to a dog, after some period of conditioning, the ringing of bells alone is sufficient to produce salivation in the dog.

Cues can be individual-specific, for instance, the desire for a particular food might be conditioned on the cue being a rainy day. But some cues apply to most of us, perhaps because these have been acquired through a common evolutionary history. For instance, the desire for immediate consumption is heightened by the smell and sight of food. This *current-bias* for consumption does not arise through time discounting (e.g., hyperbolic discounting) but through the proximity of a cue that acts as an external stimulus.

Furthermore, cues often increase the marginal utility from consumption. So, for instance, the smell of baking cookies whets the appetite for them. Cues might also trigger physiological compensatory mechanisms in the body. So, for instance, cues might offset the harmful effects of a drug. For that reason, addicts often need increasingly larger amounts of a drug to get the same “kick.”

Laibson (2001) formalizes the effect of cues on habit forming behavior. This model is also useful to illustrate the rational addiction model of Becker and Murphy (1988, 1993). Laibson describes the central premise of his theory as follows. “The cues model in the current paper captures these patterns by assuming that habit formation effects are turned on and off by the presence of cues (i.e. sensory inputs) that have been associated with past consumption of habit-forming goods.” Unlike the second premise of the visceral influences model in Section 17.2.1, it is assumed here (as also in Becker and Murphy, 1988) that individuals make a rational decision to be addicted to, or abstain from, a particular habit forming activity. Thus, the cues model can be thought of as one formalization of the first premise of the visceral influences model in Section 17.2.1.

Suppose that time is discrete, so $t = 0, 1, 2, \dots$. The individual lives infinitely and uses exponential discounting. In each period, the individual chooses one from two possible activities. A *primary activity*, or activity 1, that is potentially habit forming, and a *secondary activity*, or activity 0, that is not habit forming; the activity set is denoted by $a = \{0, 1\}$.

Before choosing the activity, the individual observes a *cue* that has no economic significance; i.e., the cues are neutral. Ex-ante, the cue is random and takes on one of two discrete values, Red or Green. The probability distribution of the cues is given by:

$$\Pr(\text{Red}) = \mu^R \geq 0, \Pr(\text{Green}) = \mu^G \geq 0, \text{ and } \mu^G + \mu^R = 1. \quad (17.1)$$

In any period t , the individual observes the cue and then makes a choice from the set a . Denote the cue-conditioned choice by $a_t^j = \{0, 1\}$ when the individual observes the cue $j = \text{Red, Green}$. So, for instance, the pair of choices $(a_t^R = 1, a_t^G = 0)$ means that in period t , the individual undertakes the primary activity only when the cue is Red. We wish to determine the optimal temporal sequence of actions $(a_0^R, a_0^G), (a_1^R, a_1^G), (a_2^R, a_2^G), \dots$

The utility function is influenced by two *stock variables*, or *compensatory processes*, each activated in any given period by the realization of the cue. Denote these compensatory processes in period t by x_t^R, x_t^G ; these variables are of the nature of stocks. The Red cue only activates the x_t^R process, while the green cue only activates the x_t^G process. An unactivated process, in any period, experiences no change in value. In particular, the evolution of the two processes is given by the following two equations:

$$\begin{cases} x_{t+1}^R = \theta x_t^R + (1 - \theta) a_t^R; & x_{t+1}^G = x_t^G, & \text{if period } t \text{ cue is Red} \\ x_{t+1}^G = \theta x_t^G + (1 - \theta) a_t^G; & x_{t+1}^R = x_t^R, & \text{if period } t \text{ cue is Green} \end{cases}, 0 < \theta < 1. \quad (17.2)$$

Hence, the activated process at time $t + 1$ is a convex combination of the time t stock and the decision. The initial values of the stock variables are denoted by $x_0^R \in [0, 1]$, $x_0^G \in [0, 1]$, and are given exogenously. Since $a_t^j = \{0, 1\}$, $x_{t+1}^j \in [0, 1]$ for all t .

Like the stock variables, the instantaneous utility is also cue-contingent. The utility of the alternative activity, activity 0, is normalized at $\bar{u}$. The stochastic cues trigger a change in preferences. The instantaneous utility in period t , following the two cues is given by

$$U_t(a_t^j, x_t^j) = u(a_t^j - \gamma x_t^j) + (1 - a_t^j) \bar{u}; j = R, G, 0 < \gamma < 1. \quad (17.3)$$

Since $a_t^j = \{0, 1\}$, the utility from the alternative activity, $\bar{u}$, kicks in only when the individual chooses the alternative activity, i.e., $a_t^j = 0$. The function $u(a_t^j - \gamma x_t^j)$ is increasing, $u' > 0$, and strictly concave, $u'' < 0$; it reflects the instantaneous utility from the current decision, a_t^j , and from the stock of the primary activity, x_t^j .

To understand the nature of instantaneous utility, suppose that the primary activity is drug addiction.

1. A decision to consume drugs in the current period, $a_t^j = 1$, increases current utility, but at an opportunity cost of $\bar{u}$. If the individual has been consuming high levels of drugs in the past, then the stock variable x_t^j is high. Thus, current utility from consuming drugs is low, i.e.,

$$\frac{d}{dx_t^j} u(a_t^j - \gamma x_t^j) = -\gamma u' < 0. \quad (17.4)$$

2. Higher drug consumption in the past, reflected in an increase in the stock variable, x_t^j , increases the current marginal utility of consumption, i.e.,

$$\frac{d^2}{dx_t^j da_t^j} u(a_t^j - \gamma x_t^j) = -\gamma u'' > 0. \quad (17.5)$$

These features are standard in habit formation models. The individual chooses the time path of actions $(a_0^R, a_0^G), (a_1^R, a_1^G), (a_2^R, a_2^G), \dots$, where $a_t^j \in \{0, 1\}$, to maximize

$$\sum_{t=0}^{\infty} \delta^t [\mu^R [u(a_t^R - \gamma x_t^R) + (1 - a_t^R) \bar{u}] + \mu^G [u(a_t^G - \gamma x_t^G) + (1 - a_t^G) \bar{u}]],$$

subject to the evolution of the stock variables in (17.2); $\delta \in (0, 1]$ is the discount factor.

An early discussion of habit formation can be found in Duesenberry (1949) and a formalization is given in Pollack (1970). Becker and Murphy (1988) provided an additional impetus to this class of models. A discrete time version of the Becker–Murphy model can be recovered from the cues model in which the cue takes only one possible value, i.e., one of μ^R, μ^G is zero. This *no-cues*

model, and its solution, also helps gain insights into the solution to the cues model. For this reason, we turn to the no-cues model, next.

17.3.1 *Solution to the no-cues model*

Suppose that the cue takes only one possible value. Thus, we can omit the superscript $j = R, G$ that we used to distinguish between different values of the cues. The problem for an individual in the no-cues model is to choose the time path of cue-independent actions $(a_1, a_2, \dots)$, where $a_i \in \{0, 1\}$, in order to maximize

$$\sum_{t=0}^{\infty} \delta^t [u(a_t - \gamma x_t) + (1 - a_t) \bar{u}], \quad (17.6)$$

subject to the constraint

$$x_{t+1} = \theta x_t + (1 - \theta) a_t.$$

This is a standard problem in dynamic programming. The Bellman equation is written as

$$V(x | \delta) = \underset{a \in \{0,1\}}{\text{Max}} [u(a - \gamma x) + (1 - a) \bar{u} + \delta V(\theta x + (1 - \theta) a) | \delta]. \quad (17.7)$$

Let us motivate the solution through the following intuition. As the stock of habits increases, it has two opposing effects. (i) From (17.4), it reduces current and future utility. (ii) From (17.5), the marginal utility from undertaking the primary activity increases. From (i) and (ii) it follows intuitively that if the stock of habits is large enough, say, greater than some threshold value $\tilde{x}$, the second effect dominates in a marginal utility based calculation, so it is optimal to choose the primary activity $a = 1$. Otherwise, the secondary activity is chosen. We summarize this result in Proposition 17.1, below. The proof is not difficult but it is lengthy; an exercise asks for a proof of the problem and gives hints to solve it.

Proposition 17.1 (*Laibson; 2001*): *In the no-cues problem, there exists some threshold value of the stock variable, $\tilde{x}$, such that the optimal decision rule of the individual is*

$$a = \begin{cases} 1 & \text{if } x > \tilde{x} \\ 0 & \text{if } x < \tilde{x} \end{cases}.$$

When $x = \tilde{x}$, some sort of tie breaking rule can be used. Several comparative static results can be derived. The threshold value, $\tilde{x}$, is increasing in $\bar{u}$ and δ . An increase in the primary activity increases the future stock of habits, hence, it reduces future utility. An increase in δ makes this reduction of future utility more salient. Hence, the threshold $\tilde{x}$ must be larger for the primary activity to be undertaken. A higher value of $\bar{u}$ increases the relative utility of the alternative good, hence, the marginal utility of the primary good must be correspondingly higher. From (17.5) this occurs if the stock of habits is higher, hence, we need a higher threshold, $\tilde{x}$.

Example 17.1 (*The dynamic programming solution in a two-period model*): *Given the importance of the no-cues model for forming intuition about the solution of the cues model, consider a two-period example. Using (17.6), the individual wishes to maximize the intertemporal utility $U_1(a_1, x_1) + \delta U_2(a_2, x_2)$, or*

$$[u(a_1 - \gamma x_1) + (1 - a_1)\bar{u}] + \delta [u(a_2 - \gamma x_2) + (1 - a_2)\bar{u}] \quad (17.8)$$

subject to the constraint

$$x_2 = \theta x_1 + (1 - \theta)a_1, \text{ and } x_1 \text{ given.} \quad (17.9)$$

The problem of the decision maker is to choose $a_1 \in \{0, 1\}$ and $a_2 \in \{0, 1\}$. Let the utility function take the CARA form, so $u(c) = \frac{-b}{\alpha} e^{-\alpha c}$, $\alpha, b > 0$; since we are not fundamentally interested in the degree of curvature of the utility function, suppose $\alpha = 1$. Check that $u' > 0$ and $u'' < 0$ are satisfied. For pedagogical simplicity, let $\delta = 1$ so the individual does not discount future utility. Denote the optimal value of the action in any period by a_t^* , $t = 1, 2$. In the second period, the decision maker chooses a_2 , given x_2 . It is optimal to choose $a_2^* = 1$ if

$$-be^{-(1-\gamma x_2)} > -be^{\gamma x_2} + \bar{u}$$

or

$$x_2 > \tilde{x} = \frac{1}{\gamma} \ln \left(\frac{\bar{u}/b}{1 - e^{-1}} \right). \quad (17.10)$$

If the inequality in (17.10) is reversed, then $a_2^* = 0$. Clearly the threshold value $\tilde{x}$ is increasing in the outside option, $\bar{u}$, as claimed earlier. Suppose that the parameter values are

$$\gamma = 0.7, \bar{u} = 1, b = 1.25, \text{ and } \theta = 0.5. \quad (17.11)$$

Then (to three decimal places) $\tilde{x} = 0.423$, which gives the decision rule

$$a_2^* = \begin{cases} 1 & \text{if } x_2 > 0.423 \\ 0 & \text{if } x_2 < 0.423 \end{cases}. \quad (17.12)$$

The second period utility is $U_2(a_2^*, x_2) = -be^{-(a_2^* - \gamma x_2)} + (1 - a_2^*)\bar{u}$. Substituting x_2 from (17.9) in U_2 we get

$$U_2(a_2^*, \theta x_1 + (1 - \theta)a_1) = \begin{cases} -be^{-(1-\gamma(\theta x_1 + (1-\theta)a_1))} & \text{if } x_2 > 0.423 \\ -be^{\gamma(\theta x_1 + (1-\theta)a_1)} + \bar{u} & \text{if } x_2 < 0.423 \end{cases}. \quad (17.13)$$

The unconstrained first period optimization problem of the individual is

$$a_1^* \in \arg \max_{a_1 \in \{0, 1\}} -be^{-(a_1 - \gamma x_1)} + (1 - a_1)\bar{u} + U_2(a_2^*, \theta x_1 + (1 - \theta)a_1), \quad (17.14)$$

where x_1 is given exogenously, U_2 is given in (17.13), and the parameter values are given in (17.11).

I. The simplest case arises when the initial stock of habits is $x_1 = 0$. Using (17.9),

$$x_2 = \begin{cases} 0.5 & \text{if } a_1 = 1 \\ 0 & \text{if } a_1 = 0 \end{cases}. \quad (17.15)$$

From (17.12) and (17.15), if the individual chooses $a_1 = 1$, then it is optimal to choose $a_2^* = 1$, otherwise if $a_1 = 0$, then the optimal second period choice is $a_2^* = 0$.

If the individual chooses $a_1 = 1$, then, using the parameter values in (17.11), his utility is $-(1.25)e^{-1} - (1.25)e^{-(1-0.7(0.5))} = -1.1124$. If he chooses $a_1 = 0$, his utility is $-(1.25)e^{-0} +$

$1 - (1.25)e^0 + 1 = -0.5$. Hence, $a_1^* = 0$. In summary, when $x_1 = 0$ and given the parameters values in (17.11), the optimal dynamic programming solution is $(a_1^*, a_2^*) = (0, 0)$.

II. Suppose at the other extreme that the initial stock of habits is high, $x_1 = 1$. In this case, for any of the two choices, $a_1 \in \{0, 1\}$, we have $x_2 \geq 0.5$, so, from (17.12), $a_2^* = 1$. Thus, irrespective of the first period choice of the individual, the second period action is identical. By choosing $a_1 = 1$, the decision maker's utility is

$$- (1.25)e^{-(1-0.7)} - (1.25)e^{-(1-0.7)} = -1.852, \quad (17.16)$$

while if he chooses $a_1 = 0$, the utility is

$$- (1.25)e^{-(0.7)} + 1 - (1.25)e^{-(1-0.7(0.5))} = -2.1697. \quad (17.17)$$

Hence, $a_1^* = 1$.

Thus, starting from a high enough stock of habits, the individual is more likely to undertake the action, say, drug consumption, in each period. However, starting from a low stock of habits, the individual is likely to refrain from drug consumption. In the case of smoking, the initial stock of drugs may well be determined by factors such as smoking in the peer group, or exposure to smoking at home. Hence, addiction to smoking may become an accident of history rather than a rational choice. ■

17.3.2 Solution to the cues model

The solutions to the model with and without cues are very similar because a given cue only activates one of the two compensatory processes or stock variables. The Bellman equation for the cues model is given by

$$W(x^R, x^G) = \frac{\mu^R}{1 - \delta\mu^G} V\left(x^R \mid \tilde{\delta} = \frac{\delta\mu^R}{1 - \delta\mu^G}\right) + \frac{\mu^G}{1 - \delta\mu^R} V\left(x^G \mid \tilde{\delta} = \frac{\delta\mu^G}{1 - \delta\mu^R}\right),$$

where $V(x \mid \tilde{\delta} = z)$ is the solution to the no-cues Bellman equation when the discount factor $\tilde{\delta}$ takes the value z and μ^R, μ^G are defined in (17.1). The solution in the presence of cues turns out to be a simple generalization of Proposition 17.1. We state this result in Proposition 17.2.

Proposition 17.2 (Laibson, 2001): *In the presence of cues, there exist threshold values $\tilde{x}^R$ and $\tilde{x}^G$ such that the optimal decision of the individual is given by*

$$a^R = \begin{cases} 1 & \text{if } x^R > \tilde{x}^R \\ 0 & \text{if } x^R < \tilde{x}^R \end{cases}$$

$$a^G = \begin{cases} 1 & \text{if } x^G > \tilde{x}^G \\ 0 & \text{if } x^G < \tilde{x}^G \end{cases}$$

As in Proposition 17.1, when $x^R = \tilde{x}^R$ or $x^G = \tilde{x}^G$, then some tie breaking rule can be used to determine the optimal action. As in the no-cues model, the threshold values are increasing in δ and $\bar{u}$. Figure 17.3, drawn in (x^R, x^G) space, shows the direction of evolution of the state variables for each possible value of the states.² The figure shows the threshold values $\tilde{x}^R, \tilde{x}^G$ that partition

² Despite the presence of discrete time, we have shown continuously differentiable phaselines for a clearer diagram.




Figure 17.3 Dynamics in the cue-dependent model.

the diagram into four quadrants, labeled, I–IV; arrows indicate the direction of motion in each quadrant. We also show the four steady states of the model, $(0, 0)$, $(1, 1)$, $(0, 1)$, $(1, 0)$, and one trajectory in each quadrant that converges to a steady state. The adjustment processes in (17.2) imply a gradual adjustment path, unless θ is close to its extreme values of zero or one.

The steady states $(0, 0)$ and $(1, 1)$ can be used to illustrate the dynamics in the no-cues Becker–Murphy model. At the steady state $(0, 0)$, the individual always undertakes the secondary activity ($a = 0$, no-addiction) because $x^R < \tilde{x}^R$ and $x^G < \tilde{x}^G$. At the steady state $(1, 1)$, the individual always undertakes the primary activity ($a = 1$, addiction), irrespective of the realization of the cue. These are the cue-independent history-dependent states, as shown in Example 17.1.

The other two steady states, $(0, 1)$, $(1, 0)$, are the cue-dependent states and illustrate the additional steady states that can arise once we allow for the introduction of cues. Consider quadrant I where $x^G < \tilde{x}^G$ and $x^R > \tilde{x}^R$; here the individual chooses the secondary activity ($a = 0$, no-addiction) if the cue is Green, and the primary activity ($a = 1$, addiction) when the cue is Red (see Proposition 17.2). The intuition follows from Example 17.1. If the stock of the primary activity is high enough, then the marginal utility from the primary activity is greater than the secondary activity, so the primary activity is preferred. At the steady state $(0, 1)$ the individual's action is also cue-dependent; choose the alternative activity if the cue is Red and the primary activity when the cue is Green.

The cues, however, are completely neutral and have no economic significance in the model. The theory, therefore, explains how neutral cues can come to form the basis of behavior, provided that the individual believes in them. One prediction of the model is that since the stochastic cues trigger a change in preferences (see (17.3)) individuals may check themselves into rehab centres in order to avoid the cravings associated with the cues. For this, one needs to invoke models of multiple selves or planner–doer models, in which the current self is more patient about future choices relative to the future selves.

In the Becker–Murphy model, once addicted, the individual rationally stays addicted forever. However, in the model with cues, the individual's consumption of a harmful good depends on the observed cues. In actual practice, we do observe that people engage in cue-management that

sometimes successfully controls consumption of the harmful good. At the same time, firms may wish to sell more by highlighting the salience of certain cues that induce individuals to consume their products, even if they are harmful.

Laibson (2001, pp. 107–8) gives examples of *cue-management behaviors* that include: hiding cigarettes, avoiding parties where alcohol will be served, and using the candy-free checkout lane. However, counter practices by firms include: installing hotel minibars in which snacks and alcohol are visible, even if the minibar has not been opened; generating artificial appetite-arousing food smells in supermarkets; packaging snack food in see-through containers; visually displaying dessert options in restaurants; and placing candy and gum in all checkout lanes to take advantage of impulse buying. Thus, an emotions-augmented account of the Becker–Murphy model gives a better explanation of several important aspects of human behavior.

This is not the only place in the book where we treat issues of addiction. Addiction can be explained through several channels. We explain it using present-biased preferences in Part 3 of the book (e.g., Gruber and Köszegi, 2001); using projection bias in Section 17.6; and using models of temptation utility in Section 17.7. Other than models of projection bias, individuals do not make mistakes; they simply make rational choices at each instant in time.

The possibility of mistakes in the presence of stochastic cues is considered by Bernheim and Rangel (2004). Stochastic cues are also a feature of the Laibson model, but the possibility of mistakes is not. An important difference in predictions between the Bernheim–Rangel model and the Laibson model is that in the former, individuals who wish to check into rehab centres fully appreciate that they will experience cravings for harmful goods. However, they hope that by checking into the rehab centres they will be able to control their actions in response to the cravings. In the Laibson model, individuals check in because they think they will be able to control their cravings by, say, avoiding the cues that give rise to the cravings.

None of the other models of addiction that we consider in this book, except for projection bias, allow for mistakes. In the presence of a projection bias, individuals might currently underappreciate the difficulty of kicking a future habit. Indeed, models of present-biased preferences and temptation utility assign no role to the empirically important phenomenon of cues either. Present-biased preferences, in addition, do not differentiate between addictive and non-addictive goods.

17.4 Anticipation and delay under certainty

The idea that anticipation about future outcomes gives rise to current utility, or *anticipal utility*, is not new in economics. It goes back to Bentham (1789) who placed anticipation of consumption on a footing similar to physical consumption in that they potentially impart pleasure and pain. These Benthamite ideas were applied further by Jevons (1905) who gave particular emphasis to the “memory of past events” and “anticipations of future events” in causing pleasure and pain. Anticipations may take the form of *anticipal pleasure* and *anticipal pain*, depending on whether the events cause *savoring* or *dread*. Anticipations may lead us to postpone consumption, as for instance, a dessert at the end of a meal. Indeed, we may exhibit negative discount rates for such cases.

Loewenstein (1987) provides a theoretical framework for anticipal utility under certainty that we now outline. Issues arising from uncertainty are considered in Section 17.5 below. Loewenstein reports the following experiment. Among other options (not reported here) students were asked how much they would pay now to (1) obtain a kiss from a movie star of their choice, and (2) avoid




Figure 17.4 Maximum payment as a proportion of current value to obtain/avoid outcomes at selected time delays.

Source: Loewenstein, G. (1987) "Anticipation and the valuation of delayed consumption," *Economic Journal* 97, 666–84. © 1987 Royal Economic Society.

receiving a non-lethal 110 volt shock. These events occurred at various time delays: 0, 3 hours, 24 hours, 3 days, 1 year, and 10 years. The results are shown in Figure 17.4.

The outcomes in Figure 17.4 are at odds with the prediction of the discounted utility (DU) model. Under positive discounting, the DU model predicts that the immediate occurrence of the pleasant event (kiss) would be most valued because the discounted value of the event occurring in the future is lower. However, subjects were willing to pay the most for the kiss to be delayed by three days, in line with the prediction of anticipatory utility. Anticipatory utility for a kiss exhibits a hump shaped pattern, peaking at three days.

Similarly, DU predicts that an unpleasant event (shock) should be delayed because the discounted value of a negative event is lower. In contrast, subjects found the unpleasant event more painful the longer the delay, hence, they were willing to pay more now to avoid shocks the later they were supposed to occur. The analogy is with the common preference of getting an aching tooth out earlier rather than later. Hence, individuals appear to prefer the unpleasant event to occur immediately to avoid the anticipatory pain of the event.

We now incorporate Jevon's *anticipatory pleasure* and *anticipatory pain* in a standard DU model. Consider an individual who, at some initial time t_0 , must decide to engage in a single act of consumption of one unit of a good at time $T > 0$. There are two kinds of utilities. The act of consumption at time T confers a utility v per period over periods $[T, T + L]$, $L > 0$, which occurs, for instance, for consumer durables. Thus, denoting the consumption utility experienced at time t by u_t^c we have

$$u_t^c = \begin{cases} v & T \leq t \leq T + L \\ 0 & \text{otherwise} \end{cases}. \quad (17.18)$$

We consider here the case of *anticipatory pleasure* and indicate later the implications arising from *anticipatory pain*. The consumer derives anticipatory utility, u_t^a , at time $t \in [t_0, T]$ from the *entire stream of consumption between time periods $T \leq t \leq T + L$* . Anticipatory utility arises prior to

the act of consumption, while consumption utility arises during the act of consumption. Hence, for times $T \leq t \leq T+L$ when the individual derives consumption utility, anticipal utility is zero. Conversely, during times $t_0 \leq t \leq T$ the individual derives anticipal utility but no consumption utility. Let $\alpha > 0$ be a measure of the anticipation or vividness of the consumption experience to follow, then

$$\begin{aligned} u_t^a &= \alpha \int_T^{T+L} e^{-\sigma(\tau-t)} v d\tau, \quad t_0 \leq t \leq T, \\ &= \frac{\alpha v}{\delta} e^{-\sigma(T-t)} (1 - e^{-\sigma L}). \end{aligned} \quad (17.19)$$

In (17.19), σ is the discount rate applicable to anticipated consumption; it need not equal the traditional discount rate r that applies to actual consumption (see (17.20)). From (17.19), we can make the following inferences that were noted by Jevons (1905).

$$\frac{\partial u_t^a}{\partial L} > 0, \quad \frac{\partial u_t^a}{\partial v} > 0, \quad \frac{\partial u_t^a}{\partial t} > 0, \quad \frac{\partial^2 u_t^a}{\partial t^2} > 0; \quad 0 \leq t \leq T.$$

Hence, instantaneous anticipal utility over $t \in [t_0, T]$ increases as one approaches the date of actual consumption. Moreover, it increases at an increasing rate. Think, for instance, about the approaching date of a long anticipated holiday. The future per-period enjoyment from the holiday, v , and the length of the holiday, L , also increase anticipal utility.

We can now combine the instantaneous consumption and anticipal utilities over the time interval $[t_0, T+L]$ into total intertemporal utility from the one-off consumption decision:

$$U = \int_{t_0}^T u_t^a e^{-rt} dt + \int_T^{T+L} u_t^c e^{-rt} dt, \quad (17.20)$$

where $r > 0$ is the conventional discount rate used in exponential discounting functions. Substituting (17.18) and (17.19) in (17.20) we get

$$U = \int_{t_0}^T \frac{\alpha v}{\delta} e^{-\sigma(T-t)} (1 - e^{-\sigma L}) e^{-rt} dt + \int_T^{T+L} v e^{-rt} dt. \quad (17.21)$$

Simplifying (17.21), we get

$$U = v \left[\frac{\alpha}{\sigma(\sigma - r)} (e^{-rT} - e^{-\sigma T}) (1 - e^{-\sigma L}) + \frac{1}{r} e^{-rT} (1 - e^{-rL}) \right]. \quad (17.22)$$

Will the consumer ever want to delay consumption, i.e., is the optimal starting date of consumption, T^* , greater than t_0 ? To answer this question, differentiate (17.22) with respect to T

$$\frac{\partial U}{\partial T} = v \left[\frac{\alpha (1 - e^{-\sigma L})}{\sigma(\sigma - r)} (\sigma e^{-\sigma T} - r e^{-rT}) - (1 - e^{-rL}) e^{-rT} \right]. \quad (17.23)$$

The first term in (17.23) gives the marginal benefit of anticipal utility (or savoring), while the second term captures the delay costs. In the absence of the first term, we get the classical outcome

under discounted utility, namely, that the individual should engage in immediate consumption of goods that give anticipal pleasure.

Suppose that $t_0 = 0$. If $\frac{\partial U}{\partial T}\big|_{T=0} > 0$, then the individual will like to defer consumption from the present date. Using (17.23), we have

$$\frac{\partial U}{\partial T}\bigg|_{T=0} = v \left[\frac{\alpha}{\sigma} (1 - e^{-\sigma L}) - (1 - e^{-rL}) \right],$$

thus,

$$\frac{\partial U}{\partial T}\bigg|_{T=0} > 0 \Leftrightarrow \frac{\alpha}{\sigma} (1 - e^{-\sigma L}) > (1 - e^{-rL}). \quad (17.24)$$

From (17.24), factors conducive to delaying are: high α (greater salience of anticipal utility), and low σ (lower discounting of future anticipated outcomes). If condition (17.24) is met, then it can be shown that there is an optimal time for consumption $T^* > 0$ that maximizes the individual's utility. At this point, $\frac{\partial U}{\partial T}\big|_{T^*} = 0$ and $\frac{\partial^2 U}{\partial T^2}\big|_{T^*} < 0$. The outcome $T^* > 0$ immediately contradicts the predictions of the discounted utility model without anticipal considerations.

We leave the reader to derive the corresponding conditions in the case of unpleasant outcomes. Discounted utility theory predicts delayed consumption in this case, while the incorporation of anticipal utility opens up the possibility that the individual would like to speed up consumption.

17.5 Fear and anxiety under uncertainty

Several models consider ex-post disappointment in static models (Bell, 1985; Loomes and Sugden, 1986). Ex-post disappointment presupposes, however, that there must have been ex-ante hopefulness or anxiety about an event. Hence, the appropriate scope of such models would appear to be dynamic. The model in Section 17.4 considered feelings of anxiety under certainty. However, since anxiety is often caused by uncertainty about future events, we need a model that considers uncertainty explicitly. Caplin and Leahy (2001) provide one such model, within an expected utility framework, which they call *psychological expected utility*.

17.5.1 An outline of the framework

Consider a two-period model. We first present a brief outline of Caplin and Leahy's general framework and then deal with a concrete life-cycle example. Suppose that the set of psychological states in the two periods are vectors of real numbers $X_t \subseteq R^n$, $t = 1, 2$. Let $X = X_1 \times X_2$ be the Cartesian product of the two psychological states, and let $P(X)$ be the space of all probability distributions over X .

Example 17.2 Suppose that the emotional states are happiness, elation, and disappointment, denoted respectively by h , e , d . Let $X_1 = \{h\}$ and $X_2 = \{e, d\}$. The model requires us to put numerical values on these psychological states that are real numbers, so $h, e, d \in R$. Then we have $X = X_1 \times X_2 = \{(h, e), (h, d)\}$. Consider two illustrative probability distributions over X , $p, q \in P(X)$ such that $p = (p_1, p_2) = (0.3, 0.7)$ and $q = (q_1, q_2) = (0.5, 0.5)$.

Definition 17.1 We assume that $P(X)$ is a mixture space. Thus, for any $\lambda \in [0, 1]$ and any $p, q \in P(X)$ there exists an element $p\lambda q \in P(X)$ such that for any $x \in X$ we have $p\lambda q(x) = \lambda p(x) + (1 - \lambda)q(x)$.

The individual has preferences $\succeq$ that are defined directly over $P(X)$. We make the following standard assumptions on preferences that readers familiar with expected utility theory will easily relate to.

A1: $\succeq$ is complete and transitive.

A2: (Substitution axiom) For all $p, q, r \in P(X)$ and $\lambda \in [0, 1]$, $p \succeq q$ implies $p\lambda r \succeq q\lambda r$.

A3: Let $p, q, r \in P(X)$ such that $p \succeq q$ and $q \succeq r$. Then there exists numbers $\lambda_1, \lambda_2 \in [0, 1]$ such that $p\lambda_1 r \succeq q$ and $q \succeq p\lambda_2 r$.

A4: (Indifference) Let $p, q \in P(X)$ such that for any $x = (x_1, x_2)$, $p(x) = (p_1(x_1), p_2(x_1))$, and $q(x) = (q_1(x_1), q_2(x_1))$, where $p_t(x_t)$ and $q_t(x_t)$ are respective probabilities of the vector of psychological state x at time t (see Example 17.2). If $(p_1(x_1), p_2(x_1)) = (q_1(x_1), q_2(x_1))$ then $p \sim q$.

A5: (Continuity) Let $\lambda \in [0, 1]$ and $p, q \in P(X)$. Then, the sets $\{z : p\lambda q \succeq z\}$ and $\{z : z \succeq p\lambda q\}$ are closed.

Since preferences are defined directly on probability distributions over X , assumption A1 alone ensures that a utility representation exists.

Proposition 17.3 (Expected utility representation) Suppose that only assumption A1 holds. Then there exists a utility function $U: X \rightarrow \mathbb{R}$ such that for any $p, q \in P(X)$, if $p \succeq q$, then $E_p U \geq E_q U$, where $E_r U$ is the expected utility under the probability density $r \in P(X)$.

Example 17.3 In Example 17.2 suppose that $p \succeq q$. Then Proposition 17.3 implies that

$$E_p U \geq E_q U \Leftrightarrow 0.3u(h, e) + 0.7u(h, d) \geq 0.5u(h, e) + 0.5u(h, d). \quad (17.25)$$

Proposition 17.4 (Additive separability of the representation; Caplin and Leahy, 2001) Suppose that assumptions A1–A5 hold. Let $p, q \in P(X)$. Then the utility function in Proposition 17.3, U , takes an additive form, such that $U(x) = u_1(x_1) + u_2(x_2)$, where $u_t: X_t \rightarrow \mathbb{R}$ for $t = 1, 2$.

Example 17.4 Continuing with Example 17.3, the function u in (17.25) can be specialized to the form, $u(x_1, x_2) = u_1(x_1) + u_2(x_2)$.

The set of possible consumption levels (or prizes) in the two periods are monetary and given by $Z_t \subseteq \mathbb{R}^{n_t}$, $t = 1, 2$. In any period $t = 1, 2$ the psychological state X_t is realized immediately after a choice is made from the consumption set Z_t . Hence, there is no within-period anticipatory utility. Anxiety is caused by anticipation of future events, hence, there is no anxiety in the second period of a two-period model. To simplify the exposition, let the psychological and physical states coincide in period 2, i.e., $Z_2 = X_2$. However, in the first period, the realization of the psychological state $x_1 \in X_1$ depends on two factors.

1. The draw from the consumption set $z_1 \in Z_1$.
2. Uncertainty about the outcomes in the second period, Z_2 . In order to formalize the second effect, let l_2 be a probability distribution over Z_2 .

We can now use the framework of Kreps and Porteus (1978) to deal with *evolving uncertainty* that is implicit in this formulation. Let $Y_1 = Z_1 \times l_2$ and denote a typical element of Y_1 by

$y_1 = (z_1, l_2)$. Let l_1 be the space of all probability distributions over Y_1 . Following Kreps–Porteus, elements of l_1 are called *temporal lotteries*. Define a function $\phi : Y_1 \rightarrow Z_1 \times X_1$ that specifies how elements of the set Y_1 translate into first period consumption levels and psychological states at time $t = 1$.

Recalling that $Z_2 = X_2$, Caplin and Leahy then define the following utility function over temporal lotteries

$$V(y_1) = u_1[\phi(y_1)] + E_{l_2} u_2[x_2], y_1 \in Y_1, x_2 \in X_2. \quad (17.26)$$

Equation (17.26) looks like the standard time separable functional form. However, the innovation here is that the distribution, l_2 , over second period outcomes is included in the first period utility function. Hence, the anxiety created by anticipation of future uncertainty affects current utility. We now look at a specific example using (17.26).

17.5.2 *Life-cycle consumption, portfolio choice, and asset prices under anxiety*

As an application of the Caplin–Leahy framework, we now consider a two-period model of consumption-portfolio choice. In the first period, a representative consumer is endowed with e_1 units of a non-storable consumption good, and 1 unit each of N productive assets. A unit of each asset yields a random output equal to $s_i \in S_i$, $i = 1, \dots, N$ in terms of the second period consumption good. There are competitive markets and the competitively determined price of the i th asset in period 1, in terms of the consumption good, is p_i . The representative-agent nature of the model implies that the aggregate supply of each asset in the economy is unity as well.

The portfolio allocation problem in the first period is to choose the amount θ_i of the i th asset to buy, $i = 1, \dots, N$. The consumer's first period budget constraint is given by

$$c_1 + \sum_{i=1}^N \theta_i p_i = e_1 + \sum_{i=1}^N p_i. \quad (17.27)$$

The RHS of (17.27) is the representative consumer's first period income. It is the sum of income from the non-storable consumption good and income from ownership of the assets. The LHS is the total expenditure, which can be split into first period consumption c_1 and investment in a portfolio of N assets. Given the definition of p_i , the units on both sides of (17.27) are in terms of the first period consumption good.

Second period income is the sum of incomes from investments in each of the N assets made in period 1. Income from the i th asset in the second period is $\theta_i s_i$, which is a random variable because s_i is random. Hence, second period consumption, c_2 , is bounded above by the random second period income $\sum_{i=1}^N \theta_i s_i$.

In terms of the notation in Section 17.5.1, we proceed as follows. Define the physical prizes in each period to be identical to the consumption in that period, so $Z_1, Z_2 \subseteq R$ and $c_t \in Z_t$, $t = 1, 2$. Thus, a typical element of Z_2 , depending on the realization of the random terms $s_i \in S_i$, is given by

$$c_2 \equiv z_2 = \sum_{i=1}^N \theta_i s_i. \quad (17.28)$$

On account of the randomness in returns, s_i , from the perspective of period 1, c_2 is random. Furthermore, each portfolio choice, $\theta = (\theta_1, \theta_2, \dots, \theta_N)$, assumed to belong to a compact set, Θ , alters the moments of this random variable. Insofar as uncertainty about the second period creates anxiety in the first, each portfolio choice is associated with a possibly different level of anxiety.

Define l_2 as the probability distribution over Z_2 . From (17.28), second period outcomes depend on θ , so it is more illustrative to write $l_2 = l_2(\theta)$. The space of *temporal lotteries* is the set of pairs $y_1 = (c_1, l_2)$. The psychological state in period 1 can now be found by using the function $\phi : Y_1 \rightarrow Z_1 \times X_1$. Let us assume a particular function, $\phi(y_1) = (c_1, a(l_2))$, where $a(l_2)$ is assumed to be a differentiable function of l_2 that captures the anxiety associated with a distribution l_2 over second period outcomes and $c_1 \in Z_1$. Noting the dependence of l_2 on θ we may write $\tilde{a}(\theta) = a(l_2(\theta))$. The temporal expected utility in (17.26) now takes the form

$$V(c_1, l_2) = u_1(c_1, \tilde{a}(\theta)) + \beta E_{l_2} u_2(c_2). \quad (17.29)$$

Substituting (17.27), (17.28) in (17.29), the unconstrained problem of the consumer is

$$\theta^* \in \arg \max_{\theta \in \Theta} u_1 \left(e_1 + \sum_{i=1}^N p_i - \sum_{i=1}^N \theta_i p_i, \tilde{a}(\theta) \right) + \beta E_{l_2} u_2 \left(\sum_{i=1}^N \theta_i s_i \right). \quad (17.30)$$

The first order condition for an interior solution is

$$-p_i \frac{\partial u_1}{\partial c_1} + \frac{\partial u_1}{\partial \tilde{a}} \frac{\partial \tilde{a}}{\partial \theta_i} + \beta E_{l_2} s_i \frac{\partial u_2}{\partial c_2} = 0; i = 1, 2, \dots, N. \quad (17.31)$$

The condition (17.31) is a standard Euler equation with the addition of an extra term (the second term) to account for anxiety. This is an arbitrage condition that illustrates the trade-off from reducing first period consumption by a unit (first term) and purchasing an extra unit of the i th asset that gives a random return s_i in the second period (third term). The second term captures the marginal effect of anxiety on first period utility that arises from the purchase of an extra unit of the i th asset.

Each asset is in fixed supply of one unit. Competitive market clearing ensures that the price of the i th asset, p_i , adjusts so as to clear the market for the asset. Given that we have a representative agent model, market clearing of the one unit implies

$$\theta_i^* = 1, i = 1, 2, \dots, N, \quad (17.32)$$

or $\theta^* = \mathbf{1}$, where $\mathbf{1}$ is a vector of ones, so $\tilde{a}(\theta^*) = \tilde{a}(\mathbf{1})$. Substituting $\theta_i^* = 1$ in (17.27) we get the optimal consumption as

$$c_1^* = e_1. \quad (17.33)$$

Substituting $\theta_i^* = 1$ in (17.28), we get

$$c_2 = \sum_{i=1}^N s_i. \quad (17.34)$$

Substituting (17.32), (17.33), (17.34) in (17.31), the price of the i th asset can be solved out as

$$p_i^* = \frac{1}{\frac{\partial u_1(e_1, \tilde{a}(1))}{\partial c_1}} \left[\frac{\partial u_1(e_1, \tilde{a}(1))}{\partial \tilde{a}} \frac{\partial \tilde{a}(1)}{\partial \theta_i} + \beta E_{l_2} s_i \frac{\partial u_2\left(\sum_{i=1}^N s_i\right)}{\partial c_2} \right]. \quad (17.35)$$

Anxiety is *aversive*, so $\frac{\partial u_1(e_1, \tilde{a}(1))}{\partial \tilde{a}} < 0$. Note that in the absence of anxiety we get $\frac{\partial u_1(e_1, \tilde{a}(1))}{\partial \tilde{a}} = 0$. Since $\frac{\partial u_1}{\partial c_1} > 0$, we immediately get our first result.

Proposition 17.5 *Anxiety alters the price of assets. If (i) $\frac{\partial \tilde{a}(1)}{\partial \theta_i} > 0$, then the price of the asset is relatively lower as compared to the no-anxiety case and the expected rate of return is higher, and if (ii) $\frac{\partial \tilde{a}(1)}{\partial \theta_i} < 0$, then the price of the asset is relatively higher as compared to the no-anxiety case and the expected rate of return is lower.*

The formula in (17.35) depends on $\frac{\partial \tilde{a}(1)}{\partial \theta_i}$. Consider the following mean-variance functional form for anxiety,

$$a(l_2) = -\alpha E_{l_2} c_2 + \gamma \text{Var}_{l_2}(c_2), \alpha > 0, \gamma > 0. \quad (17.36)$$

Hence, anxiety from investment in any vector of assets is increasing in the variance of second period consumption and decreasing in the average level of second period consumption. Using (17.28) we get

$$\text{Var}_{l_2}(c_2) = \sum_{i=1}^N \theta_i^2 \text{Var}(s_i) + \sum_{i=1}^N \sum_{j=1, j \neq i}^N \theta_i \theta_j \text{Cov}(s_i, s_j). \quad (17.37)$$

$$E_{l_2}(c_2) = \sum_{i=1}^N \theta_i E_{l_2}(s_i).$$

Hence, it follows that

$$\frac{\partial \text{Var}_{l_2}(c_2)}{\partial \theta_i} = 2 \left(\theta_i \text{Var}(s_i) + \sum_{j=1, j \neq i}^N \theta_j \text{Cov}(s_i, s_j) \right) = 2 \text{Cov}(c_2, s_i). \quad (17.38)$$

$$\frac{\partial E_{l_2}(c_2)}{\partial \theta_i} = E_{l_2}(s_i). \quad (17.39)$$

Using (17.36), (17.38), (17.39) we get

$$\frac{\partial \tilde{a}(1)}{\partial \theta_i} = -\alpha E_{l_2}(s_i) + 2\gamma \text{Cov}(c_2, s_i). \quad (17.40)$$

Substituting (17.40) in (17.35), we can now derive the prices of the N assets in terms of the fundamentals. For $i = 1, \dots, N$, the optimal asset price is

$$p_i^* = \frac{1}{\frac{\partial u_1(e_1, \tilde{a}(1))}{\partial c_1}} \left[\frac{\partial u_1(e_1, \tilde{a}(1))}{\partial \tilde{a}} \left[-\alpha E_{l_2}(s_i) + 2\gamma \text{Cov}(c_2, s_i) \right] + \beta E_{l_2} s_i \frac{\partial u_2\left(\sum_{i=1}^N s_i\right)}{\partial c_2} \right]. \quad (17.41)$$

Suppose that the i th asset is riskless, so $\text{Cov}(c_2, s_i) = 0$, hence, $\frac{\partial \tilde{a}(1)}{\partial \theta_i} = -\alpha E_{l_2}(s_i) < 0$ because a riskless asset would be held in the portfolio only if its expected return, $E_{l_2}(s_i)$, is positive. The next proposition directly follows from Proposition 17.5.

Proposition 17.6 *For a riskless asset, the price of the asset under anxiety is relatively higher and so the expected rate of return is relatively lower.*

In the presence of anxiety, by investing in a riskless asset, the investor is purchasing “peace of mind” so he is prepared to pay a relatively higher price as compared to the no-anxiety case. On the other hand, stocks are risky and the fluctuations in their returns causes anxiety. So their purchase is likely to (i) increase the expected value, which is anxiety reducing, and (ii) increase the variance of second period consumption, which is anxiety enhancing. From (17.40), the trade-off between risk and return on the level of anxiety depends on the relative magnitudes of α and γ and the underlying probabilistic structure. The next result directly follows from these observations and Proposition 17.5.

Proposition 17.7 *Suppose that γ is relatively large as compared to α .*

- (i) *If $\text{Cov}(c_2, s_i) < 0$, then $\frac{\partial \tilde{a}(1)}{\partial \theta_i} < 0$, so the price of the asset under anxiety is relatively higher and the expected rate of return is lower*
- (ii) *If $\text{Cov}(c_2, s_i) > 0$, then $\frac{\partial \tilde{a}(1)}{\partial \theta_i} > 0$. In this case, the price of the risky asset under anxiety is relatively lower and so the expected rate of return is relatively higher.*

Proposition 17.7(ii), which is the psychologically more interesting case, shows that the ownership of a risky stock creates anxiety, which reduces its price and increases its return relative to the no-anxiety case. This provides one possible explanation for the *equity premium puzzle* that we consider in more detail in Parts 1 and 7 of the book. The equity premium puzzle refers to the gap between returns on equities and bonds that cannot be explained by differences in the risk characteristics of the two financial assets.

One may distinguish between anxiety and risk aversion. Anxiety is an anticipatory feeling that arises before the resolution of an event, while under expected utility, risk aversion is a property of the curvature of the utility function. By considering risk aversion in the absence of anxiety, traditional measures of risk aversion might be partially capturing anxiety.

There are several other applications of the Caplin–Leahy framework. For an application to the revelation of information by a concerned party to an uninformed party, see Caplin and Leahy (2004). An example of this setting is a doctor who considers revealing information about a life threatening condition to two kinds of patients. An information-avoiding patient who avoids anxiety by shunning new information and an information-seeking patient who seeks information to lower anxiety (e.g., to rule out potentially worse possibilities). The natural vehicle of analysis for such a model is psychological game theory that we develop in Part 4 of the book. Caplin and Eliaz (2003) consider the possibility that individuals may be put off by the anxiety of a positive test result and, therefore, avoid AIDS testing. They embed the Kreps–Porteus preferences used above, into a mechanism design approach to propose an optimal mechanism that encourages testing and reducing the incidence of the disease.

17.6 Projection bias

Projection bias arises when individuals make predictions of their future tastes that are disproportionately influenced by their current tastes and are typically not realized in equilibrium. Consider the following examples.

Example 17.5 *Smith et al. (2008) asked patients awaiting a kidney transplant to predict their quality of life in a year’s time if they did or did not receive the transplant. The subjects*

were then followed up a year later and asked to report on their actual quality of life. The predictions failed systematically. The quality of life predicted by patients who received (did not receive) the transplant was lower (higher) relative to the actual quality of life.

Example 17.6 When ordering food at the beginning of a meal, we must predict how hungry we will be at the end of the meal. We often overestimate our future hunger and order more than we can eat. Furthermore, shopping on an empty stomach leads people to buy more (Gilbert et al., 2002). When making the decision to smoke, an individual must predict how the future desire for smoking will be affected. An underappreciation of the future desire by smokers often leads to greater current overconsumption.

Example 17.7 Read and van Leeuwen (1998) asked office workers to choose among healthy and unhealthy snacks to be received in a week's time. When asked to make this decision, office workers were in one of two possible states. In the hungry state, the question was posed just before lunchtime, and subjects were more likely to choose the unhealthy snack. However, in the satiated state, the question was posed just after lunchtime, and subjects were more likely to choose the healthy snack. The authors conclude that subjects in a currently hungry state overestimate their future hunger; in contrast, satiated subjects underestimate future hunger. In other words, people project their current tastes onto future tastes.

What accounts for projection bias? Loewenstein et al. (2003) argue that the reason is twofold. First, the ability of humans to adapt to changes is remarkably good. Second, individuals underestimate the extent of their adaptation. The two main stylized facts about projection bias that we need for the moment are:

1. Individuals are able to predict the *direction* of taste changes.
2. Individuals under-appreciate the true *magnitude* (or absolute value) of the change.

Suppose that we measure the current hunger-state of an individual on a scale of 1 (very satiated) to 10 (very hungry). Assume that normally the individual experiences a hunger-state of about 5. Suppose that we consider individuals who are currently at the two extremes, 1 and 10, and ask them to predict their hunger-state in the future. The first stylized fact implies that both individuals will report future hunger-states in the direction of 5. However, the second stylized fact implies that they will underappreciate the full extent of the movement to 5. For instance, the very satiated individual may predict a future hunger-state of 3, while the very hungry individual may predict a hunger-state of 7.

The two stylized facts also provide a plausible explanation of the endowment effect. Before obtaining a mug, individuals are not adapted to owning the mug and underappreciate the satisfaction derived from its ownership, hence, they undervalue it. After obtaining the mug, individuals adapt to ownership but underappreciate how well they will adapt if they were to be deprived of it, hence, they overvalue it.

17.6.1 Preferences under projection bias

Suppose that an individual lives for T periods, $t = 1, 2, \dots, T$. The instantaneous utility in period t is given by $u(c_t, s_t)$, where c_t is consumption at time t and s_t is the state of the world at time t that reflects tastes of the individual. The set of possible states is $S \subset \mathbb{R}^n$. Given the current state, s_t , the individual's prediction of utility from consumption c and state $s_{t'}$ is given by $\tilde{u}(c, s_{t'} \mid s_t)$.

We now give a definition of projection bias that incorporates stylized facts 1 and 2.

Definition 17.2 (Loewenstein et al., 2003): An individual exhibits projection bias if for all c, s_t and $s_{t'}$ there exists $\alpha \in [0, 1]$ such that

$$\tilde{u}(c, s_{t'} | s_t) = \alpha u(c, s_t) + (1 - \alpha)u(c, s_{t'}); \alpha \in [0, 1]. \quad (17.42)$$

In Definition 17.2, the prediction of future utility is a linear combination of the current instantaneous utility and the actual instantaneous utility in the relevant state. The central feature is that current tastes, as reflected in s_t , affect one's prediction of future tastes. If $\alpha = 0$, then we have the classical case of no projection bias. When $\alpha = 1$, the individual has extreme projection bias and equates current tastes fully with future tastes. Intermediate values of α capture varying degrees of projection bias. In particular, for $0 < \alpha < 1$, the individual accurately predicts the direction of change in tastes but underpredicts their magnitude.

Example 17.8 Consider two states of the world, hungry (state h), and not hungry (state n). Suppose that we have only one consumption good c and an identical consumption level in each period. Then we have the following four cases.

$$\tilde{u}(c, h | h) = \alpha u(c, h) + (1 - \alpha)u(c, h) = u(c, h). \quad (17.43)$$

$$\tilde{u}(c, n | h) = \alpha u(c, h) + (1 - \alpha)u(c, n). \quad (17.44)$$

$$\tilde{u}(c, n | n) = \alpha u(c, n) + (1 - \alpha)u(c, n) = u(c, n). \quad (17.45)$$

$$\tilde{u}(c, h | n) = \alpha u(c, n) + (1 - \alpha)u(c, h). \quad (17.46)$$

From (17.43), (17.45), there is no projection bias if the future and current states are identical. But from (17.44), (17.46) projection bias exists if $\alpha \neq 0$ and the two states differ. Suppose that

$$u(c, h) = 20, u(c, n) = 10, \alpha = 0.5,$$

then we have $\tilde{u}(c, n | h) = \tilde{u}(c, h | n) = 15$. If the current state is n and the future state is h , then the change in predicted utility is 5. Reversing the states, the change in predicted utility is -5 . In each case, the direction of change is correct (stylized fact 1), but the absolute value of the change is $5 < 10$. However, in the absence of projection bias, the absolute value of the change is 10. Hence, in each case, the magnitude of the change is underappreciated (stylized fact 2).

In writing down the individual's intertemporal preferences in the presence of projection bias, there are two polar cases to consider.

1. The individual has complete foresight about the entire future temporal path of states $s_1, s_2, \dots, s_T$: In this case, the intertemporal utility at date $t = 1$ is given by

$$U(c_1, c_2, \dots, c_T | s_1) = u(c_1, s_1) + \sum_{\tau=2}^{\tau=T} \delta^{\tau-1} \tilde{u}(c_\tau, s_\tau | s_1), \quad (17.47)$$

where δ is the discount factor and $\tilde{u}$ is defined in (17.42).

2. There is uncertainty over the future states: Denote by E_1 , the expectation operator taken at date 1, over the relevant probability distribution over future states. Then, the expected intertemporal utility is given by

$$E_1 U(c_1, c_2, \dots, c_T | s_1) = u(c_1, s_1) + E_1 \sum_{\tau=2}^{\tau=T} \delta^{\tau-1} \tilde{u}(c_\tau, s_\tau | s_1). \quad (17.48)$$

Time inconsistency problems can arise when individuals have projection bias because the prediction of individuals about their future tastes will often not be realized. For instance, in (17.44), (17.46) $\tilde{u}(c, n | h) \neq u(c, n)$ and $\tilde{u}(c, h | n) \neq u(c, h)$. Since the projection bias arises from a mis-prediction, individuals would not know it (for otherwise they could easily correct it by setting $\alpha = 0$). Hence, individuals perceive preferences to be time consistent but they are not. The model does not address possible learning by individuals when predictions do not come out to be true. In other words, individuals are unaware that they might be making systematic forecasting errors. This does typically arise when humans are hardwired to engage in certain behaviors; considerable cognitive effort and practice are needed to shed such behaviors.

17.6.2 *Habit formation, life-cycle consumption, and projection bias*

Consider a consumer with projection bias who, at time $t = 1$, allocates a given initial income Y over T periods. For pedagogical simplicity, there is no uncertainty, the discount factor $\delta = 1$, and the interest rate is zero. The actual instantaneous utility of the consumer at time t is given by $u(c_t, s_t)$, where c_t is consumption and s_t is the state of the world at time t . Preferences are of the habit formation type, i.e.,

$$u(c_t, s_t) = v(c_t - s_t), v' > 0, v'' < 0. \quad (17.49)$$

s_t is a stock of habits that evolves in the following manner

$$s_{t+1} = (1 - \lambda)s_t + \lambda c_t; \lambda \in (0, 1]. \quad (17.50)$$

Current consumption adds to the habit stock, and the initial state, $s_1 > 0$, is given.

Solving (17.50) recursively, we get

$$s_{t+1} = (1 - \lambda)^t s_1 + \lambda c_t + \lambda (1 - \lambda)^1 c_{t-1} + \lambda (1 - \lambda)^2 c_{t-2} + \dots + \lambda (1 - \lambda)^{t-1} c_1. \quad (17.51)$$

This setup satisfies the following two standard features of the habit formation model.

1. Current consumption, by adding to the stock of habit, directly reduces utility in all future time periods. Using (17.49), (17.51) we get

$$\frac{\partial u(c_t, s_t)}{\partial c_{t-k}} = -\lambda (1 - \lambda)^{k-1} v'(c_t - s_t) < 0; k = 1, 2, \dots, t - 1. \quad (17.52)$$

Hence, one needs to consume ever greater amounts to get the same “kick” from a habit.

2. The marginal utility from current consumption is increasing in the current stock of habits because

$$\frac{\partial}{\partial s_t} \left(\frac{\partial u(c_t, s_t)}{\partial c_t} \right) = -v''(c_t - s_t) > 0. \quad (17.53)$$

Given (17.42), (17.47), (17.49), the consumer's optimization problem at time $t = 1$ is to choose $c_1, c_2, \dots, c_T$ in order to maximize

$$U(c_1, c_2, \dots, c_T | s_1) = v(c_1 - s_1) + \sum_{\tau=2}^{\tau=T} [\alpha v(c_\tau - s_1) + (1 - \alpha)v(c_\tau - s_\tau)], \quad (17.54)$$

subject to the equation of motion for the stock variable (17.50), which we rewrite as

$$s_{t+1} - s_t = \lambda(c_t - s_t); \text{ given } s_1 > 0, \quad (17.55)$$

and the feasibility constraint that takes account of the assumption of zero interest rate,

$$\sum_{\tau=1}^{\tau=T} c_\tau \leq Y. \quad (17.56)$$

Denote the solution under no-projection bias ($\alpha = 0$), or the unbiased solution, by $c_1^u, c_2^u, \dots, c_T^u$, and the solution under projection bias ($\alpha > 0$), or the biased solution, by $c_1^b, c_2^b, \dots, c_T^b$. The qualitative features of the solution for the general case can be found in Loewenstein et al. (2003).³

We derive below the closed-form solution for an illustrative two-period model, $T = 2$, and assuming log preferences

$$v(c_t - s_t) = \ln(\theta + c_t - s_t); \theta = \max\{s_1, s_2\};$$

this ensures that $\theta + c_t - s_t > 0$, so we have a well-defined utility function with $v' > 0, v'' < 0$.

For $T = 2$, substituting s_2 from (17.55) and c_2 from (17.56) into (17.54), the consumer chooses c_1 to maximize the following concave utility function

$$\begin{aligned} U(c_1, c_2 | s_1) &= \ln(\theta + c_1 - s_1) + \alpha \ln[\theta + Y - c_1 - s_1] \\ &\quad + (1 - \alpha) \ln[\theta + Y - (1 + \lambda)c_1 - (1 - \lambda)s_1], \end{aligned}$$

given the initial state variable s_1 . We assume an interior solution throughout. The first order condition with respect to c_1 is

$$\frac{1}{\theta + c_1 - s_1} - \frac{\alpha}{\theta + Y - c_1 - s_1} - \frac{(1 - \alpha)(1 + \lambda)}{\theta + Y - (1 + \lambda)c_1 - (1 - \lambda)s_1} = 0. \quad (17.57)$$

Consider first, the absence of projection bias, so $\alpha = 0$. Evaluating (17.57) at $\alpha = 0$, the unbiased solution, c_1^u , is given by

$$c_1^u = \frac{1}{1 + \lambda} \left(\frac{\tilde{Y}}{2} + \lambda s_1 \right), \quad (17.58)$$

where

$$\tilde{Y} = Y - \lambda\theta.$$

³ The solution to the general case cannot be obtained by the usual optimal control method. The reason is that the presence of a habit stock variable does not allow one to reduce the system to a system of first order difference equations. What one would get, for instance, by eliminating the shadow prices, at any point t , is a system of two $t - 1$ order difference equations in c, s . Hence, the qualitative features of the solution can only be derived by using the method of induction.

We can rewrite (17.58) as $c_1^u = \frac{1}{1+\lambda} \left(\frac{\tilde{Y}}{2} - s_1 \right) + s_1$, so a sufficient condition for $c_1^u > s_1$ is that $\frac{\tilde{Y}}{2} > s_1$. We next show that when an individual consumes more than her stock of habits, then she consumes an increasing profile of consumption. This outcome occurs if the initial stock of habits is sufficiently small.

Proposition 17.8 Suppose that there is no projection bias. If $\frac{\tilde{Y}}{2} > s_1$ (or $c_1^u > s_1$), then $c_1^u < c_2^u$.

Proof: (a) Substituting (17.58) in (17.56). Then, for $T = 2$, the optimal solution for c_2^u is

$$c_2^u = c_1^u + \left(\frac{2\lambda}{1+\lambda} \right) (\tilde{Y} - 2s_1) + \lambda\theta. \quad (17.59)$$

From (17.58) we get

$$c_1^u - s_1 = \frac{1}{1+\lambda} \left(\frac{\tilde{Y}}{2} - s_1 \right). \quad (17.60)$$

From (17.60), $c_1^u > s_1 \Leftrightarrow \frac{\tilde{Y}}{2} > s_1$, which implies, from (17.59), that $c_2^u > c_1^u$. ■

Proposition 17.9 (Loewenstein et al., 2003): If $c_t^u > s_t$ for some t , then $c_t^u < c_{t+1}^u < c_{t+2}^u < \dots < c_T^u$.

From Proposition 17.8, we get the standard result in habit formation models that $s_1 = 0$ is a sufficient condition for an increasing profile of optimal consumption. By frontloading consumption, in this case, the stock of habits is raised early on, reducing the marginal utility of consumption in the subsequent periods. Hence, the opposite, planning an increasing profile of consumption, is optimal.

Let us now introduce the possibility of projection bias. The next proposition states the main result.

Proposition 17.10 Suppose $T = 2$ and there is projection bias ($\alpha > 0$). If $\frac{\tilde{Y}}{2} > s_1$ (so $c_1^u > s_1$), then $c_1^u < c_1^b$.

Proof: Implicitly differentiate (17.57) with respect to α

$$\frac{\partial c_1}{\partial \alpha} = \left(-\frac{\partial^2 U}{\partial c_1^2} \right)^{-1} \left(\frac{(1+\lambda)}{\theta + Y - (1+\lambda)c_1 - (1-\lambda)s_1} - \frac{1}{\theta + Y - c_1 - s_1} \right). \quad (17.61)$$

For $\frac{\partial c_1}{\partial \alpha} > 0$, we require that $\frac{(1+\lambda)}{\theta + Y - (1+\lambda)c_1 - (1-\lambda)s_1} > \frac{1}{\theta + Y - c_1 - s_1} \Leftrightarrow Y > 2s_1 - \theta$. But we know from Proposition 17.8 that the condition $\frac{\tilde{Y}}{2} > s_1$ ensures that $c_1^u > s_1$, and $\frac{\tilde{Y}}{2} > s_1 \Leftrightarrow Y > 2s_1 + \lambda\theta$. Clearly, in this case, $Y > 2s_1 - \theta$. Hence, if $\frac{\tilde{Y}}{2} > s_1$ (so $c_1^u > s_1$), then $\frac{\partial c_1}{\partial \alpha} > 0$. Since c_1^u corresponds to the optimal solution when $\alpha = 0$ and c_1^b corresponds to the optimal solution when $\alpha > 0$ it follows that $c_1^u < c_1^b$, as claimed. ■

The analogous result when there are T time periods is given in Proposition 17.11 below.

Proposition 17.11 (Loewenstein et al., 2003): If $c_1^u > s_1$ and $\alpha > 0$, then $\sum_{\tau=1}^t c_\tau^u < \sum_{\tau=1}^t c_\tau^b$ for $t < T$.

The intuition behind Proposition 17.10 is as follows. Projection bias implies an underappreciation of the magnitude of change. Here, the consumer underappreciates the effect of habit formation as well as the negative effect of the current stock on all future instantaneous utilities. Hence, relative to the unbiased case, the consumer with projection bias consumes too much. Therefore, a person with projection bias will consume her income “too quickly” relative to the optimal unbiased consumption profile. An important implication of this idea is that consumers might not save enough. This provides an alternative explanation of the savings inadequacy problem that does not rely on arguments of present-biased preferences (see Part 3 of the book).

17.6.3 *An application to consumer durables*

Suppose that a consumer can purchase one unit of a consumer durable at time $t = 1$ at some price P that provides a stream of utility $d_1, d_2, \dots, d_T$. The consumer’s “state” in any period is the utility in that period, i.e., $s_t = d_t$, $t = 1, \dots, T$. There is no discounting, so the discount factor $\delta = 0$. Suppose that the consumer buys the good in period 1. Then her true or unbiased utility (corresponding to $\alpha = 0$), U^u , is given by

$$U^u = \sum_{\tau=1}^{\tau=T} d_{\tau} - P = T\bar{d} - P, \quad (17.62)$$

where $\bar{d} = \frac{1}{T} \sum_{\tau=1}^{\tau=T} d_{\tau}$ is the mean, across time, of instantaneous utility.

The intertemporal utility of a consumer who has projection bias, U^b , at time $t = 1$, is given by

$$U^b = \sum_{\tau=1}^{\tau=T} [\alpha d_1 + (1 - \alpha) d_{\tau}] - P = T\bar{d} + \alpha T (d_1 - \bar{d}) - P. \quad (17.63)$$

Using (17.62), (17.63) it follows that

$$U^b \geq U^u \text{ if } d_1 \geq \bar{d}. \quad (17.64)$$

From (17.64), if $d_1 > \bar{d}$, then a consumer with projection bias may buy the consumer durable even if in the absence of projection bias he does not buy (e.g., if $T\bar{d} - P < 0$ and $T\bar{d} - P + \alpha T (d_1 - \bar{d}) > 0$). Hence, a person with projection bias may be too sensitive to her valuation at the time of purchase. Insofar as the valuation may be influenced by an emotional *hot state*, it plays a disproportionate role in the presence of projection bias.

Businesses might wish to exploit the tendency of consumers to buy in a *hot state*. They may induce such a state through sales hype, attractive displays, in-store music, and one-click shopping on the internet. Often salesmen will offer a discount, conditional on buying today, but the discount is not available if the purchase is delayed. Sellers might also try to turn non-durable goods into durable goods. So, for instance, they might sell an annual or bi-annual membership to a health club rather than a daily one.

DellaVigna and Malmendier (2002) provide evidence that people overpay for health club membership, by buying on an annual basis, relative to a “pay as you use” basis. One possible explanation is that people have self-control problems; once people sign up to a health club membership,

they procrastinate and do not fully use it (see Part 3 of the book). However, the explanation based on projection bias suggests that people overestimate their use of the membership because at the time of signing, they are in a hot state. In the future, they cool down as their enthusiasm wanes, and they underuse their membership. However, due to projection bias, individuals underpredict the extent to which they will cool down.

This has interesting implications for government policy. In the US and UK, for instance, “cooling-off” laws allow consumers to cancel their purchase within a certain number of days subsequent to the purchase. One interpretation of these laws is that they are designed to counter the effect of projection bias. Such laws can also alter the incentives of salespersons to mis-sell because they too wish to reduce costly returns of their products. Thus, salespersons are more likely to highlight longer-run implications of a purchase decision to the consumer that are more consistent with the cooler rather than the hotter end of the spectrum of human emotions.

17.6.4 *Other applications of projection bias*

Habit formation is a plausible hypothesis that is based on the psychological foundations of adaptation. However, evidence of habit formation is not readily available; but see Dynan (2000) for the experimental evidence. Projection bias suggests one explanation for why it might be hard to find evidence of such behavior, even if individuals exhibit it. Projection bias might lead people to underreact to habit formation. For instance, using (17.57), if the individual has extreme projection bias, say, $\alpha = 1$, then the individual’s optimal consumption profile is $c_1 = Y/2$, $c_2 = Y/2$ whether habit formation is present or absent. Hence, projection bias might mask the presence of habit formation even if it exists.

Projection bias might also explain why many pledges of charitable support are not eventually honored. This occurs particularly when a major disaster occurs that elicits a widespread feeling of empathy. If individuals make their pledges in a hot state, they overestimate their charitable urges in the future. One’s own future-selves might not honor the earlier pledge because they make decisions in a cooler state, once the immediate shock of the disaster lessens.

Habit formation models often find it difficult to explain why people develop harmful habits such as addiction in the first place.⁴ Projection bias provides two potential explanations. First, people underestimate the (negative) effect of current consumption on future utility. Second, individuals might under-appreciate the formation of the addictive habit.

The making of decisions in hot states that characterizes individuals with projection bias might also explain why some people get married without sufficient deliberation. Or why costly proposals to get married are sometimes reversed. Or why people undertake costly actions/commitments in moments of rage or extreme generosity. The existence of cooling-off periods could be one way of mitigating the harmful consequences of such actions. For instance, there are mandatory time intervals between filing for marriage or divorce.

Such decisions are, of course, irreversible for suicides. If individuals in a currently depressed state predict, on account of projection bias, that their future states are also likely to be similar, then their predicted lifetime discounted value of life is relatively lower. Hence, some individuals might, on the spur of the moment, make a disastrous and irreversible decision to end their life. Such individuals, if they are known to the social services, are typically closely monitored for the existence of such hot states.

⁴ See, however, Bernheim and Rangel (2004), above, who explain addiction as a mistake.

17.7 Temptation preferences: a revealed preference approach

Models of dual selves and models that consider the interaction between the emotional and the cognitive system in the brain rely on the premise that human behavior is the outcome of *brain modularity*. The primitives in such a model are the relevant modules in the brain and their interaction.

A more classical approach takes preferences as the relevant primitives and uses the revealed preference approach to recover the underlying preferences. In this approach, individuals face temptations but may either make *commitments* to eliminate them, or they might not be able to eliminate the temptations. There are two possibilities—individuals may either exercise *self-control*, or they may succumb to the temptations. This section considers these issues.

Consider two time periods. In the first period, the individual chooses from a set of *menus* and has preferences over menus given by $\succsim_1$. The items in the menu can be outcomes or, more generally, probability distributions over outcomes. In the second period, conditional on the chosen menu in the first period, the individual chooses items in the menu, based on second period preferences, $\succsim_2$. The underlying framework was introduced by Kreps and Porteus (1978).

Example 17.9 Consider two periods and two possible restaurants: A vegetarian restaurant, V , that sells healthy but bland food and a fast food restaurant, F , that sells unhealthy, but tempting food. In the first period, the individual needs to choose over three possible “menus”: $\{V\}$, $\{F\}$ and $\{V, F\}$. Menus are defined more formally below. For the moment, if menu $\{V\}$ (respectively $\{F\}$) is chosen in the first period, then, in the second period, the individual can only visit restaurant V (respectively restaurant F). If menu $\{V, F\}$ is chosen, then the individual can choose any of the two restaurants V or F in the second period. The first period preferences over menus are denoted by $\succsim_1$. In the second period, the individual can visit any of the restaurants in the chosen menu (if there is a single option on the menu, then the second period choice is trivial). Second period preferences over restaurants in the chosen menu are given by $\succsim_2$.

Option F is said to be a temptation if $\{V\} \succsim_1 \{V, F\}$. In this case, in the first period, the individual would like to commit to just choose option V rather than having to choose among options V and F in the second period. Conditional on the first period preferences $\{V\} \succsim_1 \{V, F\}$, suppose that the individual cannot commit to choosing $\{V\}$. In this case, in the second period, the individual faces a choice from the menu $\{V, F\}$. If the individual chooses F , he gives in to temptation, and if he chooses V , he is said to exercise self-control.

The seminal theoretical models that explore these issues are Gul and Pesendorfer (2001, 2004, 2005, 2007) and Dekel et al. (2001). In Example 17.9, temptation arises on account of second period preferences, $\succsim_2$. First period preferences may dictate a preference for commitment. Noor (2011) introduces the possibility that individuals might also be tempted by menus in the first period. For instance, the very fact that F is a tastier, although unhealthy, option may create a preference for the menu $\{V, F\}$ over the menu $\{V\}$. The subsequent discussion relies on Gul and Pesendorfer (2001), henceforth, GP.

17.7.1 The formal model

Consider a two-period model. Suppose that $C = \{c_1, c_2, \dots, c_n\}$ are the possible levels of consumption in period 2. Let ΔC be the set of probability distributions over C ; we denote typical

elements of ΔC by lowercase alphabets such as p, q, r . Thus, a typical element p in ΔC takes the form $p = (p_1, \dots, p_n)$, such that $\sum_{i=1}^n p_i = 1$, and $p_i \geq 0$ for all $i = 1, \dots, n$; the interpretation of p_i is the probability of consumption c_i .

Let Z be the set of all non-empty compact subsets of ΔC . The set Z comprises all possible menu of choices that, in period 1, the individual can choose from. We will typically denote elements of Z , or *menus*, as M_i, M_j, M_k , and so on. Denote the preferences of the individual over elements in the set Z , or his first period preferences, by $\succsim_1$. For instance, $M_i \succsim_1 M_j$ means that menu M_i is “at least as good as” M_j . The corresponding strict preference and indifference relations are given, respectively, by $\succ_1$ and $\sim_1$.

Suppose that the menu chosen in the first period is M_i . In the second period, the individual makes an optimal choice from M_i . An element $m_i \in M_i$ induces a second period lottery of the form $(c_1, p_1; \dots, c_n, p_n)$. Second stage preferences, $\succsim_2$, are defined over ΔC .

Define the following mixture operation over elements of the set Z . For any $M_i, M_j \in Z$ and $0 < \lambda < 1$ we have

$$\lambda M_i + (1 - \lambda)M_j = \{\lambda q + (1 - \lambda)p, \text{ where } q \in M_i, p \in M_j\}.$$

The binary preference relations $\succsim_1$ and $\succsim_2$ are the primitives of the model and they satisfy the following conditions for $M_i, M_j, M_k \in Z$.

- A1: (Preference relations) $\succsim_1$ and $\succsim_2$ are complete and transitive.
- A2: (Strong continuity) The sets $\{M_i \in Z : M_i \succsim_1 M_j\}$, $\{M_i \in Z : M_j \succsim_1 M_i\}$, and the set $\{p \in \Delta C : p \succsim_2 q\}$ are closed.
- A3: (Independence) If $M_i \succ_1 M_j$, then for all $0 < \lambda < 1$,

$$\lambda M_i + (1 - \lambda)M_k \succ_1 \lambda M_j + (1 - \lambda)M_k.$$

- A4: (Set betweenness) $M_i \succsim_1 M_j \Rightarrow M_i \succsim_1 M_i \cup M_j \succsim_1 M_j$.

Axioms A1–A3 are standard in economics and are used, for instance, in the axiomatic derivation of the expected utility formula. However, Fudenberg and Levine (2006) and Fudenberg (2006) argue that the independence axiom is less plausible when economic agents have self-control problems. They argue that the evidence that self-control is a scarce resource, which gets depleted with cognitive load is not consistent with the independence axiom.

The new axiom is *set betweenness* that we now explain in more detail by considering the following cases.

1. $M_i \sim_1 M_i \cup M_j \succ_1 M_j$: The preference $M_i \sim_1 M_i \cup M_j$ indicates that the preferred choice belongs to the set M_i . In this case, adding the option M_j to M_i is irrelevant. Hence, M_j is not a temptation.
2. $M_i \succ_1 M_i \cup M_j \sim_1 M_j$: From the first preference, the individual prefers M_i to the case where both M_i and M_j are available. Such an alternative, M_j , is termed as a *temptation*. From the second preference, when faced with a choice between elements of the sets $M_i \cup M_j$ and M_j the individual chooses an element from the set M_j . Hence, ex-ante the individual prefers to exclude the set M_j altogether from the available options. But, from the preference $M_i \cup M_j \sim_1 M_j$, if the option M_j is available, then the individual succumbs to the temptation.
3. $M_i \succ_1 M_i \cup M_j \succ_1 M_j$: The preference $M_i \succ_1 M_i \cup M_j$ reveals that M_j is a temptation. However, the preference $M_i \cup M_j \succ_1 M_j$ differentiates this from the previous case. In

this case, even when the option M_i and the temptation M_j are jointly available at the time of making the choice, the individual successfully resists the temptation, possibly by using some form of costly *self-control*.

To summarize, there are three possibilities. Either there is no temptation (case 1) or there is a temptation that cannot be resisted (case 2) or there is a temptation that can be successfully resisted (case 3).

We have made sufficient assumptions to ensure that utility functions exist that represent the preferences, $\succsim_1$ and $\succsim_2$.

Definition 17.3

- (i) A utility function $U : Z \rightarrow \mathbb{R}$ represents the first period preferences $\succsim_1$ if for $M_i, M_j \in Z$, $M_i \succsim_1 M_j \Leftrightarrow U(M_i) \geq U(M_j)$.
- (ii) A utility function $u : \Delta C \rightarrow \mathbb{R}$, sometimes called lottery utility, represents the second period preferences $\succsim_2$ if for $p, q \in \Delta C$, $p \succsim_2 q \Leftrightarrow u(p) \geq u(q)$. Furthermore, u is said to be linear if

$$u(\lambda p + (1 - \lambda)q) = \lambda u(p) + (1 - \lambda)u(q). \quad (17.65)$$

There is no difference between first stage and second stage preferences when the role of temptation preferences is ignored. In this case, the individual evaluates each menu $M_i \in Z$ by computing the expected utility, as in (17.65), for each element, $m_i \in M_i$ and then picking the element with the highest expected utility. This motivates the next definition.

Definition 17.4 (Classical preferences): Suppose that an individual satisfies the conditions on the utility function in Definition 17.3(ii). The individual has classical preferences if he proceeds as follows. (i) For each menu $M_j \in Z$, let $p_j^* \in \arg \max_{p \in M_j} u(p)$. (ii) Pick the menu, $M_k \in \arg \max_{M_j \in Z} u(p_j^*)$. In this case, the pair of preferences $(\succsim_1, \succsim_2)$ are called classical preferences.

It is obvious from Definition 17.4 that an individual who has classical preferences cannot be worse off by expanding the size of the menu. The reason is that at worse, the expected utility of all new items added to the menu will have an expected utility lower than the maximal element in the old menu. But then, the decision maker will simply not chose from any of the new elements in the menu. So, for classical preferences, we have the following condition

$$M_i \succ_1 M_j \Rightarrow M_i \cup M_j \sim_1 M_i. \quad (17.66)$$

We now introduce an axiom that connects first and second period preferences.

A5: If $M_i \cup \{p\} \succ_1 M_i$, then $p \succ_2 q$ for all $q \in M_i$.

Axiom A5, sometimes known as the *sophistication axiom* (see Lipman and Pesendorfer, 2013), requires that if the individual prefers to add an option p to a menu in M_i , then that option must be chosen in period 2. Axiom A5 is not always innocuous. For instance, firms may offer new contracts to existing customers (adding to their menus) in order to make existing contracts look more attractive, even when they know that such contracts will not be chosen.

We now state the main result in the absence of temptation preferences.

Proposition 17.12 (Classical preferences; Kreps, 1979; Gul and Pesendorfer, 2001): Suppose that the following conditions hold. (i) For all $M_i, M_j \in Z$, condition (17.66) holds. (ii)

Axioms 1–5 hold. Then $(\succsim_1, \succsim_2)$ are classical preferences. Furthermore, if $(\succsim_1, \succsim_2)$ are classical preferences, then conditions (i) and (ii) hold.

Once we depart from classical preferences and introduce the role of temptation, then the condition in (17.66) need not hold any longer. Indeed, we may have $M_i \succ_1 M_i \cup M_j$ so that individuals prefer to have a smaller set of choices in their menus in order to avoid temptation. The main result of GP can be stated in the following theorem.

Proposition 17.13 (Self-control preferences; Gul and Pesendorfer, 2001): *The preference relation $\succsim_1$ satisfies Axioms A1–A4, if and only if, there exists a utility function U satisfying Definition 17.3(i), and utility functions u and v satisfying Definition 17.3(ii) such that*

$$U(M) = \max_{p \in M} [u(p) + v(p)] - \max_{p \in M} v(p), \quad M \in Z. \quad (17.67)$$

We may call such preferences self-control preferences.

If preferences can be represented as in (17.67), we say that they have a (u, v) representation. Given below is a proof of Proposition 17.13, using the planner–doer framework in Bénabou and Pycia (2002). While the planner–doer framework is not in the tradition of the revealed preference framework, the proof shows nicely the potential relation between the two frameworks.

One of the problems at the end of the chapter asks you to show that the following proposition is true.

Proposition 17.14 *Classical preferences (Definition 17.4) are a special case of self-control preferences (see Proposition 17.13) if either (i) v is an affine transformation of u , i.e., $v = \alpha u + \beta$ for some $\alpha \geq 0$ and $\beta \in \mathbb{R}$, or (ii) u is a constant function. Furthermore, when v is a non-affine transformation of u , i.e., $v = \alpha u + \beta$ but $\alpha < 0$ and $\beta \in \mathbb{R}$, then the individual has a sole desire for commitment.*

Three main features underlie the GP framework.

B1: Suppose that the individual chooses the menu $M \in Z$ in period 1. The individual's preferences in the second period depend on tuples of the form (M, p) , where $p \in M$. In the second period, the individual's optimization problem is

$$\max_{p \in M} (u(p) + v(p)). \quad (17.68)$$

B2: Conditional on the chosen menu $M \in Z$ in period 1, suppose that the individual chooses some element $p^* \in M$ in the second period; this is the solution to (17.68). Then the first period utility of the individual is given by

$$U = u(p^*) + v(p^*) - \max_{p \in M} v(p). \quad (17.69)$$

From (17.68), the last term $\max_{p \in M} v(p)$ in (17.69) does not affect the actual choice made in the second period. It does, however, affect the individual's level of welfare. Temptation lowers the individual's utility. The second period preferences directly induce first period preferences over menus given in (17.67). In other words, in the first period, for each menu, $M \in Z$, the individual computes (17.69), and assigns a level of utility to it.

B3: In period 1, $M_i \succsim_1 M_j$ if and only if $U(M_i) \geq U(M_j)$.

Suppose that the individual chooses the menu $M \in Z$ in period 1. Denote by $p^* \in M$, the solution to the problem stated in (17.68). Suppose that p^* is unique and suppose that the individual could commit to a singleton menu $\{p^*\}$ in the first period. Using (17.69)

$$U(p^*) = u(p^*) + v(p^*) - v(p^*) = u(p^*). \quad (17.70)$$

For this reason, GP refer to u as the *commitment ordering over outcomes*. It represents the utility from singleton menus.

Now suppose that the individual were to choose a larger (non-singleton) set M of alternatives in period 1 that contains p^* . Suppose also that the individual is able to resist any potential temptations so that the chosen alternative is still p^* in the second period. Then, using (17.68), (17.69), (17.70), welfare loss to the individual is given by

$$U(p^*) - U(M) = u(p^*) - \left(u(p^*) + v(p^*) - \max_{p \in M} v(p) \right) = \max_{p \in M} v(p) - v(p^*). \quad (17.71)$$

Hence, the term $\max_{p \in M} v(p) - v(p^*)$ has the interpretation of *cost of self-control*. The more tempting the alternatives in the larger set M , the greater is the cost of self-control. For this reason GP, refer to $v(p)$ as the individual's *temptation ordering*.

Example 17.10 Consider two possible menus, $M_1 = \{p\}$ and $M_2 = \{p, q\}$. Suppose that

$$u(p) > u(q) \text{ and } v(q) > v(p). \quad (17.72)$$

We now show that the individual responds to temptation by choosing the smaller menu, M_1 , if commitment is possible. If commitment is not possible, and menu M_2 is chosen, then the individual must choose, in period 2, whether to give in to temptation or to exercise self-control. Using (17.70), we have $U(M_1) = u(p)$, while using (17.67), (17.72),

$$U(M_2) = \max \{ u(p) + v(p), u(q) + v(q) \} - v(q).$$

Consider now the following two cases.

I. Let $\max \{ u(p) + v(p), u(q) + v(q) \} = u(p) + v(p)$. In this case,

$$U(M_2) = u(p) - (v(q) - v(p)) < u(p) = U(M_1). \quad (17.73)$$

II. Let $\max \{ u(p) + v(p), u(q) + v(q) \} = u(q) + v(q)$. In this case,

$$U(M_2) = u(q) < u(p) = U(M_1). \quad (17.74)$$

Using (17.73), (17.74) we get $\{p\} \succ_1 \{p, q\}$, hence, q is a temptation. In (17.73), the individual can exercise self-control in period 2 and does not succumb to temptation. However, there is a cost of self-control, given by the term $v(q) - v(p)$. In (17.73), the individual succumbs to temptation. Indeed when preferences, $\succ_1$, have a (u, v) representation it can be shown that q tempts p in the sense $\{p\} \succ_1 \{p, q\}$ if and only if (17.72) holds.

Suppose that, in addition, a third menu $M_3 = \{q\}$ is available. In this case,

$$U(M_3) = u(q) < u(p) = U(M_1). \quad (17.75)$$

Of the two cases, I and II, considered above, let us focus only on case I:

Using (17.73) we have

$$U(M_2) \geq U(M_3) \Leftrightarrow u(p) - u(q) \geq v(q) - v(p).$$

The interesting case is the one where $u(p) - u(q) > v(q) - v(p)$, so that $U(M_2) > U(M_3)$. Hence, the overall picture in this case is

$$\{p\} \succ_1 \{p, q\} \succ_1 \{q\}. \quad (17.76)$$

Thus, q is a temptation and the individual successfully exercises self-control if presented with the menu M_2 . From (17.76), we get the main empirical prediction of this class of models. Namely, that an individual prefers to commit even when the individual will not engage in choice reversals. However, Lipman and Pesendorfer (2013) write: “We are unaware of any experimental or empirical study that seeks to document self-control.” Clearly this is an area that is in urgent need for empirical testing.

We now consider the invariance of GP preferences to affine transformations of commitment utility, u , and temptation utility, v .

Proposition 17.15 (Gul and Pesendorfer, 2001) Suppose that preferences, $\succ_1$, have a (u, v) representation (see Proposition 17.13). Suppose also that u, v are non-constant and v is not an affine transformation of u (see Proposition 17.14). Then $(\hat{u}, \hat{v})$ also represent the preferences $\succ_1$ if and only if there exist $\alpha > 0$ and $\beta_u, \beta_v \in \mathbb{R}$ such that $\hat{u} = \alpha u + \beta_u$ and $\hat{v} = \alpha v + \beta_v$.

A special case of temptation preferences arises in the context of time inconsistency (Strotz, 1955). These preferences, known as Strotz preferences, ensure that the long-run plans of the individual are time consistent. In a two-period model, the preferences prescribe the following lexicographic ordering. Begin with period 2, where some menu $M \in Z$ has already been chosen. The individual then maximizes temptation utility, v , over all elements in the menu M ; suppose that the maximal element is $p \in M$. Rolling back to the first period, the individual then assigns to the menu M , the number $u(p)$, where u is commitment utility. Finally, the individual chooses the menu with the maximal commitment utility. This is stated in the next definition.

Definition 17.5 (Strotz, 1955): Suppose that a preference pair $(\succsim_1, \succsim_2)$ can be represented by the following utility function:

1. For any first stage choice of $M \in Z$, and $p, q \in M$, $p \succsim_2 q$ if and only if the following lexicographic ordering holds

$$v(p) > v(q), \text{ or } v(p) = v(q) \text{ and } u(p) > u(q). \quad (17.77)$$

2. Looking ahead at the second period choice from any menu $M \in Z$, as given in (17.77), first period preferences $\succsim_1$, satisfy

$$U(M) := \max_{p \in M} u(p) \text{ such that } v(p) \geq v(q) \text{ for all } q \in M.$$

The preference pair $(\succsim_1, \succsim_2)$ so defined is said to be a Strotz pair or that $\succsim_1$ can be represented by a Strotz utility function.

Proposition 17.13 also holds for Strotz preferences. Thus, the binary relation $\succsim_1$ satisfies Axioms 1–4 if and only if it is a Strotz preference. Furthermore, a modified form of Proposition 17.15 also holds for Strotz preferences.⁵

⁵ The only difference is that we need separate values of α , say, $\alpha_u, \alpha_v > 0$ for, respectively, the utility functions $\hat{u}$ and $\hat{v}$; see Lipman and Pesendorfer (2013).

Example 17.11 (Lipman and Pesendorfer, 2013) Consider the following example of time inconsistent choices that is explained using quasi-hyperbolic discounting, or (β, δ) preferences, in Strotz (1955), Phelps and Pollak (1968), and Laibson (1997).⁶ Suppose that there are two time periods and $C = \{(c_1, c_2) \in [0, \bar{c}]^2\}$ where $\bar{c}$ is an upper bound on consumption in each period. Consider the following Strotz preferences (Definition 17.5):

$$\begin{cases} u(c_1, c_2) = w(c_1) + \delta w(c_2) \\ v(c_1, c_2) = w(c_1) + \beta \delta w(c_2) \end{cases}, \quad 0 < \beta < 1,$$

where u is commitment utility that applies to ex-ante choices in period 1, and v is temptation utility that applies to choices in period 2. Hence, in period 2, the individual will choose the consumption pair $(0, l)$ over $(s, 0)$ if $\beta \delta w(l) > w(s)$. In period 1, the individual must choose the optimal menu using commitment utility, u . He chooses the menu $M_1 = \{(0, l)\}$ over the menu $M_2 = \{(s, 0)\}$ if $\delta w(l) > w(s)$. Choice reversals occur if

$$\delta w(l) > w(s) > \beta \delta w(l).$$

Here commitment utility chooses $(0, l)$ over $(s, 0)$, and temptation utility chooses $(s, 0)$ over $(0, l)$, if presented with the menu $M_3 = \{(s, 0), (0, l)\}$. Faced with this possibility, in period 1 the individual might exercise self-control if it is not too costly, by choosing only the menu M_1 .

We refer the reader to the excellent survey by Lipman and Pesendorfer (2013) to pursue other developments in this literature. These include the idea that temptations might take the form of deliberate distortion in beliefs; see Epstein (2006), Epstein and Kopylov (2007), and Epstein et al. (2008). For instance, individuals might have a temptation to believe that a flood in the neighboring county will not cause them any flood risk, so that they can avoid buying insurance. Relative to the Bayesian benchmark, this model can lead to over reaction and under reaction to new information. For extensions to multiple periods, see Gul and Pesendorfer (2004) and Krussel et al. (2010).

17.7.2 Interpreting temptation preferences in a planner–doer framework

It is possible to derive GP preferences in the planner–doer framework in the tradition of Thaler and Shefrin (1981). This can help to provide a useful reinterpretation of the GP model. Bénabou and Pycia (2002) pursue this approach that we describe below.

Suppose there are two time periods and there is a long-lived planner (cognitive system) and a short-lived myopic doer (emotional system) in each period. The doer only cares about temptation utility, v , while the planner cares about the sum of commitment and temptation utility, i.e., $u + v$. We assume that u, v are strictly concave. As in the GP model, suppose that the individual chooses from some opportunity set, ΔC , which is a set of lotteries, in the first period. Conditional on a choice of M in the first period, there are two possibilities for the second period choice of an element in M , given by

$$\begin{cases} x \in \arg \max_{x' \in M} \{u(x') + v(x')\} & \text{if choice is made by planner.} \\ y \in \arg \max_{y' \in M} v(y') & \text{if choice is made by doer.} \end{cases} \quad (17.78)$$

⁶ This case is discussed at length in Part 3 of the book.

As in the Loewenstein and O'Donoghue (2004, 2007) model (see Section 18.2.2 below), we assume that the cognitive and the emotional systems simultaneously incur respective costs $r_c > 0$ and $r_e > 0$ to probabilistically influence the final outcome. The authors write (p. 423): "One can also think of the two subelves as lobbying the brain's motor control areas in the same way as interest groups lobby the government." The respective probabilities, p_c, p_e , that either of the cognitive and the emotional systems is decisive in making the final choice, are directly proportional to the resources expended, in the following manner.

$$p_c = \frac{r_c}{r_e + r_c}; p_e = \frac{r_e}{r_c + r_e}, \text{ and } p_c + p_e = 1. \quad (17.79)$$

We suppose that the costs r_c, r_e are drawn, respectively, from some compact sets $S_c \subseteq R_+$ and $S_e \subseteq R_+$. We are interested in a static Nash equilibrium in the actions of the emotional and cognitive systems, (r_c^*, r_e^*) . In light of (17.78), (17.79), the respective objective functions of the emotional and the cognitive system are given by

$$U^e = p_c(r_c, r_e) v(x) + p_e(r_c, r_e) v(y) - r_e. \quad (17.80)$$

$$U^c = p_c(r_c, r_e) (u(x) + v(x)) + p_e(r_c, r_e) (u(y) + v(y)) - r_c. \quad (17.81)$$

Solving the two first order conditions $\frac{\partial U^e}{\partial r_e} = 0$ and $\frac{\partial U^c}{\partial r_c} = 0$, simultaneously, for an interior solution, we get the unique optimal solution

$$r_e^* = \frac{(r_e + r_c)^2}{(u(x) + v(x)) - (u(y) + v(y))}; r_c^* = \frac{(r_e + r_c)^2}{v(y) - v(x)}. \quad (17.82)$$

Substituting (17.82), in (17.79), the unique respective probabilities are

$$p_e = \frac{v(y) - v(x)}{u(x) - u(y)}; p_c = \frac{u(x) + v(x) - u(y) - v(y)}{u(x) - u(y)}. \quad (17.83)$$

Having determined period 2 choices, let us now revert back to period 1. Suppose that in period 1, the individual evaluates date 2 lotteries using expected utility, and individual outcomes using her *commitment preferences* u . Then, ex-ante, the choice set M gives a level of expected utility

$$EU(M) = p_c u(x) + p_e u(y). \quad (17.84)$$

Substituting (17.83), in (17.84) we get

$$EU(M) = u(x) + v(x) - v(y). \quad (17.85)$$

Using (17.78), we can rewrite (17.85) as

$$EU(M) = \max_{x' \in M} \{u(x') + v(x')\} - \max_{y' \in M} v(y'). \quad (17.86)$$

Equations (17.67) and (17.86) are identical, as claimed. Thus, the planner–doer model in which emotions or cognition are probabilistically decisive in the second period, generates the same first period preferences as in GP.

17.8 Temptation and conflicts between commitment and flexibility

When individuals have self-control problems, say, arising through temptation, the standard prescription, in a range of models, is to commit to eliminate options that lead to such temptations. Amador et al. (2006) point out an important trade-off in this context. Eliminating a subset of the options has the benefit of reducing/eliminating temptations, but at the cost of reduced flexibility, especially when new information is likely to arrive in the future. For instance, in the context of intertemporal choice, a reduction in temptations might improve the savings rate of an individual, however, if the future marginal utility turns out to be too low, then the improvement in savings might not have been as worthwhile.

Consider a simple two-period example. In the first period, the individual can choose to buy one of two prepaid skiing holiday in the Alps: a nonrefundable holiday and a refundable holiday. The nonrefundable holiday is relatively cheaper. The individual can commit to eliminating temptations in the first period that lead him to buy the nonrefundable option. However, in the second period, due to an unfortunate accident, the individual is unable to enjoy the fruits of his first period efforts. Ex-post, he would have been better off buying the more expensive, but refundable option. If the individual had foreseen the possibility of such an event, his decision would have been the outcome of a *preference for commitment*, as well as a *preference for flexibility*.

We explore these considerations in the model of Amador et al. (2006) below, using a *multiple-selves* framework. This class of models is considered in more detail in Part 3 of the book. Here we simply note that the individual has multiple-selves, one for each time period. Each of the multiple-selves has its own objective function and none of the selves can directly control the other, although indirect control is possible.

Consider a three-period model of a single, representative, individual. Time periods are subscripted by $t = 0, 1, 2$. The individual has initial endowment y at time $t = 0$. There is no consumption in period 0 ($t = 0$) and consumption in period 1 ($t = 1$) is given by $c_1 = c$. The individual saves an amount k at time $t = 1$, to be carried into the second period ($t = 2$). There is no other source of income, and the interest rate on savings is normalized to zero, so $c_2 \equiv k$. Thus, the budget set of the individual is given by

$$B = \{(c, k) \in \mathbb{R}_+^2 : c + k \leq y\}. \quad (17.87)$$

The individual has multiple-selves and we shall refer to the self in period t as self- t . Self-2 has no important economic decisions to make, so we ignore it for the moment. Self-0 is the *principal*, while self-1 is the *agent*. The descriptions of the problems for self-0 and self-1 are given below.

1. In period 1, the preferences of self-1 are defined over first period consumption, c , and second period consumption, k . These preferences are given by

$$\theta u(c) + \beta v(k), \quad (17.88)$$

- where $\theta > 0$, $0 < \beta \leq 1$ are taste parameters, and u, v are increasing, concave and continuously differentiable functions. The parameter θ is private information to self-1 and takes on two possible values, θ_L, θ_H , with respective probabilities $p, 1-p$, and $0 < \theta_L < \theta_H$.
2. Self-0 does not know the value of θ but knows its distribution. The preferences of self-0 are given by

$$E[\theta u(c) + v(k)], \quad (17.89)$$

where E denotes the expectation operator with respect to θ . Notice that self-1 has a bias toward first period consumption ($\because 0 < \beta \leq 1$) but self-0 has no such bias. The *strength of the bias towards current consumption* is given by the term $1 - \beta$, so the greater is β , the lower is the bias towards current consumption.

17.8.1 First best outcome

Suppose that the principal (self-0) has full information on θ . The first best solution maximizes the utility of the principal, given in (17.89), subject to the budget constraint, given in (17.87). Thus, under the first best, there is no bias towards current consumption. Denoting the solution by $(c^*(\theta), k^*(\theta))$, $\theta = \theta_L, \theta_H$, we have

$$c^*(\theta) \in \arg \max_{c \in [0, y]} [\theta u(c) + v(y - c)].$$

The first order condition is given by

$$\theta u'(c^*(\theta)) = v'(y - c^*(\theta)). \quad (17.90)$$

Substituting $c^*(\theta)$ from (17.90) into (17.87) we get $k^*(\theta)$. Since u, v are strictly concave, so the first best is unique. Implicitly differentiating (17.90)

$$\frac{dc^*(\theta)}{d\theta} = \frac{-u'}{\theta u'' + v''} > 0. \quad (17.91)$$

Using (17.87), (17.91), we get the following monotonicity condition.

Lemma 17.1 *At the first best allocation $c^*(\theta_H) > c^*(\theta_L)$ and $k^*(\theta_H) < k^*(\theta_L)$.*

Thus, under full information, consumption is higher when the marginal utility of consumption at $t = 1$ is higher. We shall sometimes denote $c^*(\theta_i)$ and $k^*(\theta_i)$ by, respectively, c_i^* and k_i^* , $i = H, L$.

17.8.2 The optimal commitment solution under asymmetric information

Now suppose that self-0 does not know θ . Self-1 cannot control (even by exercising cost control) the temptation towards current consumption arising from $0 < \beta \leq 1$. Self-0 now optimally chooses an appropriate budget set $B = \bar{B}$ in the full knowledge that self-1 is decisive in making the consumption-savings choice. The main trade-off in the model is to tailor consumption to the

taste shock θ that is unobserved at $t = 0$ (flexibility) and restrict consumption opportunities for self-1 due to present-biased preferences (commitment).

Hence, the problem of self-0 is to choose $\bar{B}$ in order to maximize (17.89), subject to the constraint that the choice of self-1, $\bar{c}(\theta), \bar{k}(\theta)$, will be made according to

$$(\bar{c}(\theta), \bar{k}(\theta)) \in \arg \max_{c, k \in \bar{B}} \theta u(c) + \beta v(k).$$

This problem is equivalent to the following problem in which self-0 directly chooses the allocations $(c_i, k_i) = (c(\theta_i), k(\theta_i))$, $i = L, H$ in order to maximize her objective

$$Z = (1 - p) [\theta_H u(c_H) + v(k_H)] + p [\theta_L u(c_L) + v(k_L)], \quad (17.92)$$

subject to the following constraints.

$$IC_H : \theta_H u(c_H) + \beta v(k_H) \geq \theta_H u(c_L) + \beta v(k_L). \quad (17.93)$$

$$IC_L : \theta_L u(c_L) + \beta v(k_L) \geq \theta_L u(c_H) + \beta v(k_H). \quad (17.94)$$

$$c_H + k_H \leq y; \quad c_L + k_L \leq y. \quad (17.95)$$

Effectively, we have that $\bar{B} = \{(c_H, k_H), (c_L, k_L) : (17.95) \text{ holds}\}$, so one allocation is chosen for each type. (17.92) is just a restatement of (17.89) to take account of the uncertainty of self-0 about the values of θ , while (17.93), (17.94) are, respectively, the incentive compatibility conditions for types θ_H and θ_L ; we denote these, respectively, by IC_H and IC_L . The incentive compatibility constraint ensures that each type prefers its intended allocation. There are no participation constraints because it is not clear what an outside option might be in a model of multiple selves. Finally, condition (17.95) requires that the type-contingent budget constraint must be respected.

We first show that if the *strength of bias towards current consumption* is low enough (or β is high enough), then there is no problem in implementing the first best allocation; this result does not hold for a continuum of types.

Proposition 17.16 *There exists some $\bar{\beta} \in (\frac{\theta_L}{\theta_H}, 1)$ such that for all $\beta \in [\bar{\beta}, 1]$, the first best allocation $(c_L^*, k_L^*), (c_H^*, k_H^*)$ can be implemented under asymmetric information.*

Proof: Suppose that we evaluate IC_L at the first best allocation (c_i^*, k_i^*) when $\beta = 1$. By definition, IC_L is satisfied, so

$$\theta_L u(c_L^*) + v(k_L^*) > \theta_L u(c_H^*) + v(k_H^*) \Leftrightarrow v(k_L^*) - v(k_H^*) > \theta_L [u(c_H^*) - u(c_L^*)]. \quad (17.96)$$

The strict inequality in (17.96) arises because the first best is unique. From (17.96), it follows that there exists some $\bar{\beta} < 1$ such that

$$\bar{\beta} [v(k_L^*) - v(k_H^*)] = \theta_L [u(c_H^*) - u(c_L^*)]. \quad (17.97)$$

$$\Rightarrow \bar{\beta} = \frac{\theta_L [u(c_H^*) - u(c_L^*)]}{[v(k_L^*) - v(k_H^*)]}. \quad (17.98)$$

$$\Rightarrow \bar{\beta} \geq \theta_L \frac{u'(c_H^*)}{v'(k_H^*)}. \quad (17.99)$$

$$\Rightarrow \bar{\beta} \geq \frac{\theta_L}{\theta_H}. \quad (17.100)$$

Expression (17.99) follows from the property that for an increasing and concave function, w , $w'(c_H^*) \leq \frac{w(c_H^*) - w(c_L^*)}{c_H^* - c_L^*}$ and $k_L^* - k_H^* = c_H^* - c_L^*$, while (17.100) follows from (17.99) by using the first order condition under full information, (17.90). From (17.97), IC_L evaluated at the first best allocation holds for all $\beta \geq \bar{\beta}$. Analogously, evaluating IC_H at the first best allocation, a simple calculation shows that it holds if

$$\beta \leq \frac{\theta_H [u(c_H^*) - u(c_L^*)]}{[v(k_L^*) - v(k_H^*)]}.$$

Using (17.98) this implies that

$$\beta \leq \frac{\theta_H}{\theta_L} \bar{\beta}.$$

Using (17.100), $\frac{\theta_H}{\theta_L} \bar{\beta} > 1$, so IC_H , evaluated at the first best allocation, holds for all $\beta \leq 1$.

Thus, for all $\beta \in [\bar{\beta}, 1]$, IC_L , IC_H hold at the first best allocation, which is therefore, implementable. ■

Denote the asymmetric equilibrium solution by $(\bar{c}_H, \bar{k}_H)$, $(\bar{c}_L, \bar{k}_L)$. We now show that for high enough temptation $(1 - \beta)$, it is optimal to bunch the types, while for intermediate levels of temptation, it is optimal to separate the types. Bunching shares a common feature with *minimum savings rules* in other models, namely, that types above some threshold have the same consumption and saving bundle.

Proposition 17.17 (Amador et al., 2006)

- (a) If $\beta \geq \frac{\theta_L}{\theta_H}$, then a separating equilibrium exists such that $\bar{c}_H > \bar{c}_L$ and $\bar{k}_H < \bar{k}_L$. Both inequalities in (17.95) bind at the optimal solution.
- (b) If $\beta \leq \frac{\theta_L}{\theta_H}$, then bunching is optimal, i.e., $\bar{c}_H = \bar{c}_L$ and $\bar{k}_H = \bar{k}_L$. Both inequalities in (17.95) bind at the optimal solution.

In Part 3 we considered evidence that individuals do not save enough towards their retirement. Suppose that the government imposed a minimum savings requirement on all individuals, independent of their types. Is there a justification for such a *minimum savings rule* in a model of multiple-selves where selves have asymmetric information about the subsequent selves? The answer, it turns out from Proposition 17.17(b) is in the affirmative. A second contribution of the analysis in this section is to show that the removal of all tempting options, is suboptimal. Rather, individual allocations, tailored to the flexible needs of different types (see Proposition 17.17(a)) are optimal. Only when β is low enough as in Proposition 17.17(b) (self-control problems of self-1 are high enough), does self-0 remove the flexibility associated with a separating equilibrium and forces each type to save the same amount (commitment).

17.8.3 *Some empirical evidence on self-control*

There is surprisingly little formal evidence on the direct measurement of self-control, although one observes potential implications of self-control problems such as inadequate savings for retirement. There is now a fair bit of evidence on hyperbolic discounting that leads to self-control problems; see Part 3 of the book for details that we omit here.

Ameriks et al. (2007) is a relatively unusual study that tries to directly measure self-control problems. The authors assume that individuals correctly predict their future self-control problems, while anecdotal evidence suggests that individuals may be imperfectly aware of these problems. Survey data is used to pose hypothetical choices, so there is no guarantee that their procedure is incentive compatible; this would be more of a concern for experimental economists, although the view taken in this book is that all forms of evidence, including survey data, is admissible.

The study proceeds as follows. Individuals are told that they have won 10 certificates and each certificate entitles them to eat one free meal at a good restaurant. The certificates must be used entirely within the first two years, otherwise they are not valid. One measure of the contribution of self-control problems is the self-reported expected-ideal (EI) gap. This is the gap between the stated *expected* number of gift certificates that people use and the *ideal* number of gift certificates they would like to use in year 1 as compared to year 2. For 60% of the sample, the ideal split of the certificates over the two years is 50–50. Of the remaining sample, 15% wished to use all their certificates in the first year.

The EI gap is found to be quite small, suggesting low self-control problems. For 95% of the sample, the EI gap belongs to the set $\{-2, -1, 0, 1, 2\}$. Approximately two-thirds of the subjects choose the modal value of the EI gap, which is 0. In a regression analysis, individuals with higher self-reported value of the degree of conscientiousness exhibit smaller absolute values of the EI gap. Older individuals exhibit a smaller absolute value of the EI gap, suggesting that temptations fall with age. This may be helpful in understanding why we typically do not observe consumption splurges from lump-sum liquid retirement incomes.

In order to further explore self-control problems, subjects in the sample were given a choice of precommitting to the number of certificates that they could use in the first year. This gives rise to another measure, the revealed preference (RP) gap. This is the difference between the expected number of certificates that subjects hope to consume in the first year and the number they would like to commit to. The RP gap turns out to be quite low and is indicative of low self-control problems. Only 10% of the subjects have a desire to impose strictly binding constraints on themselves. Perhaps, as in the work of Amador et al. (2006), they have a preference for flexibility.

17.9 Happiness economics

The dominant practice in economics is to infer the utility function of individuals from their actual choices and to use the utility function to predict the choices that people make. We may refer to such utility as *decision utility* or *wantability*. Bentham (1789) had in mind a different concept, namely, *experienced utility*, which is not normally considered in modern economics. Experienced utility is the hedonic experience associated with an outcome. As such, one may obtain it by integrating the instantaneous experienced utilities associated with an outcome.

17.9.1 *Why should we be interested in subjective well-being?*

Kahneman and Thaler (2006, p. 222) write: “In the older interpretation of utility, the question of whether choices maximize utility has a simple meaning: do people choose the options that they will most enjoy? In modern decision theory, which ignores the distinction, the question is quite different: are preferences consistent with each other and with the axioms of rational choice?” Insofar as the stream of instantaneous utilities has accrued in the past, the individual has to contend with issues of imperfect memory; in this case, experienced utility is sometimes known as *remembered utility*. If instantaneous utilities accrue in the future (say, a purchase of a consumer durable), the individual might be subjected to projection bias (see Section 17.6 above). In this case, experienced utility is sometimes known as *predicted utility*. Hence, an important issue is whether experienced utility can be successfully measured.

Kahneman et al. (1997) review the evidence on remembered utility, drawing on Varey and Kahneman (1992), Fredrickson and Kahneman (1993), Kahneman et al. (1993), Redelmeier and Kahneman (1996). The typical procedure in these studies is to measure the discomfort of the subjects, say, during a colonoscopy, every 60 seconds (instantaneous utilities) and then ask them to report the disutility of the entire procedure at the end (remembered utility). The remembered utility can then be compared to the reports of instantaneous utilities. For a preference representation that takes account of these features, see Kahneman et al. (1997). The two main findings are as follows.

1. *Peak-end evaluation*. Peak-end utility is measured as a combination of the (i) average of the remembered peaks of past instantaneous utilities, and (ii) the instant utility towards the end of the experience. Peak-end evaluation has the following two implications. (a) *Duration neglect*: The duration of the experience has little effect if the average of the remembered peaks does not change much. (b) *Violations of temporal monotonicity*: Since peak-end evaluation weights the latest episode of utility, if there is a particularly low (respectively, high) recent instantaneous utility experience, overall remembered utility can be reduced (respectively, increased).
2. *Decisions by remembered utility*: When comparing two streams of past utilities and evaluating them by peak-end evaluation, the decision maker chooses the one with the highest remembered utility.

Subjective well-being relies on *expressed preferences* rather than on *revealed choice*. Revealed preference cannot be relied upon to judge the utility implications for individuals in cases where they have little or no effect on the outcome. Examples include environmental degradation, income inequality, fairness of the tax system, social capital, quality of governance, inflation, and unemployment. Indeed, in these cases, self-reported measures of *life satisfaction*, *well-being*, or *happiness* (the literature often uses these terms interchangeably), may provide an insight into the utility effects on individuals.

Utility is ordinal in the revealed preference approach. If a utility function, u , represents preferences, so does the function $\phi(u)$, where ϕ is any monotonically increasing transformation of u . Hence, based on actual choices made by individuals, we cannot know which utility function in the class of all monotonically increasing transformations of u we have actually measured. Yet, we routinely hear individuals making statements of the form: Peter is happier than John. The ability to make interpersonal comparisons of utility suggests that individuals may have some underlying cardinal measure of preferences.

Many of the measures of well-being are non-pecuniary. Economic policy that focuses purely on pecuniary measures of well-being will, therefore, not optimize overall well-being. *Subjective well-being* encompasses both *experienced utility* and *procedural utility* (Frey and Stutzer, 2002). Procedural utility refers to the actual procedure that leads to an outcome; see, for instance, Part 2 for a discussion of procedural fairness.

Can subjective well-being be measured? Traditionally the economics profession has been suspicious of self-reported measures of well-being that are based on hypothetical questions and the absence of incentives. Furthermore, the memories of subjects could be imperfect, they may be subject to projection bias, to framing, and to emotional and contextual effects. However, the implications of these potential errors depend on the intended usage of the data. Senik (2014) offers the following interesting quote attributed to Alan Blinder: “If molecules could talk, would physicists refuse to listen?”

The data are typically intended to provide an ordinal measure of well-being, and identify the determinants of well-being and happiness. For instance, we might be interested to see if reported subjective well-being is higher or lower when the quality of governance increases. The numerical increase in the reported subjective well-being is of lesser importance for this question. There is typically a high degree of consistency in the data, both in time series and cross-sectional data. Reassuringly, there is a high correlation between various measures of happiness. These considerations lead many to believe that individuals are perfectly capable of answering questions about subjective well-being in a consistent manner; see Frey and Stutzer (2002, Section 2.3) for the supporting evidence and the details.

Benjamin et al. (2012) evaluate the claim that *subjective well-being* (SWB) is a good proxy for utility in the sense that individuals make choices that would maximize their subjective well-being. Using three different samples, they give subjects a series of hypothetical choices in a variety of scenarios. Consider one of the scenarios, Scenario 1, in which individuals are asked to consider the trade-off between sleep and income.

“Say you have to decide between two new jobs. The jobs are exactly the same in almost every way but have different work hours and pay different amounts.

Option 1: A job paying \$80,000 per year. The hours for this job are reasonable, and you would be able to get about 7.5 hours of sleep on the average work night.

Option 2: A job paying \$140,000 per year. However, this job requires you to go to work at unusual hours, and you would only be able to sleep around 6 hours on the average work night.

Between these two options, taking all things together, which do you think would give you a happier life as a whole?”

At this point, subjects indicate a preference for one of the options and also indicate the intensity of the choice, based on a three point scale. Having answered this subjective well-being question, subjects are asked the following *choice question*.

“If you were limited to these two options, which do you think you would choose?”

For the choice question as well, subjects make one of the choices and indicate the intensity of the preference on a three point scale.

A total of 13 different scenarios were used across the subject pool. Subjects are given choices for jobs or housing options that trade-off financial rewards versus more sleep (Scenario 1), a shorter commute (Scenario 12), being around friends (Scenario 13), making more money relative to others (Scenario 3), and so on. In other scenarios, subjective well-being is not the only consideration. For instance, in Scenario 4, subjects choose between two career paths. One involves an easier life with fewer sacrifices, while another promises posthumous impact and fame. Scenario 2 involved an ethical choice: A more convenient option versus one that is the “right thing to do.”

The main results are as follows. Across all samples, the *SWB ranking* of the two options and the *actual choices* match for 83% of all respondents. For one sample, the Denver sample, only 13% of the rankings do not match, and for another, the Cornell sample, only 24% of the rankings do not match. Thus, subjective well-being is a good predictor of actual choices. However, for both samples, a test of proportions rejects the null hypothesis that the proportion of subjects making choice 1 is the same in the subjective well-being question and the choice question. When one of the options involves more money, respondents are more likely to choose the option with more money in the choice question, relative to predicting that this option will yield greater subjective well-being.

In order to explain the difference between the two sets of questions, the authors introduce several new explanatory variables for the Cornell sample. These include the following variables, each measured on a seven point scale: Your family's happiness, social life, control over life, level of spirituality, social status, and sense of purpose. When actual choices are regressed on subjective well-being and 11 other variables, it turns out that subjective well-being predicts choices better than all the other variables combined. Further, on average, the increase in explanatory power that is achieved by adding the 11 extra variables is modest, although in some scenarios the increase is substantial.

These results establish the importance of measures of social well-being as well as offer caveats about its use. It is no longer tenable to dismiss this line of research. The evidence supports plausible links of self-reported measures of well-being with meaningful economic variables. These include, for instance, the ability of these measures to predict job quits (Clark, 2001), probability of suicides (Helliwell, 2007), productivity (Oswald et al., 2015), and probability of a future divorce among couples (Guyen et al., 2012).

A dominant tradition in the literature, as we shall see below, equates subjective well-being with self-reported measures of happiness, health, and life satisfaction; this is the approach that we shall take in this section. Indeed, in this approach, these self-reported measures of well-being are taken to directly reflect utility. But what actually constitutes subjective well-being? Benjamin et al. (2014) explore this issue. In addition to happiness, their list includes goals, achievements, freedom, morality, and relationships. They find that happiness, health, and life satisfaction do indeed play a major role in well-being, but so do many other variables such as the quality of relationships, financial and physical security, morality, and freedom of choice.

We have considered the importance of character virtues, typically suppressed in neoclassical economics, in Chapter 7 of the book. The empirical results in Benjamin et al. confirm the importance of character virtues for well-being. Indeed, many of the variables that they find important for well-being correspond to the use in Sen (1985, 1992) of *functionings* and *capabilities*; the former refers to material standards of living and the latter to opportunities to achieve the relevant functionings, and to individual freedoms.⁷

There is a recognition in the literature that self-reported measures of well-being may not always give an accurate picture of utility, due to some of the reasons considered above. Yet, there is merit in considering these measures on two counts. First, as Benjamin et al. (2014, p. 2703) write: "Nonetheless, we believe it is more attractive to rely on what people's own stated preferences suggest about what they themselves care about than to paternalistically rely on the opinions and introspections of 'experts' (such as researchers and policymakers) regarding which aspects to track and how to weight them." Second, there might be a normative use of self-reported

⁷ For the social welfare foundations of economics in the presence of these measures of well-being, see Fleurbaey (2009) and Fleurbaey and Blanchet (2013).

measures of well-being. These measures may embody deliberative responses to a greater extent (as compared to emotional responses) that are more reflective of the actual welfare of the individuals.

Carter and McBride (2013) provide evidence that the shape of experienced utility is similar to the S-shaped utility function under prospect theory. Furthermore, they also report some evidence for loss aversion. The similarity between decision utility and experienced utility at a fundamental level, despite individual-level heterogeneity, is clearly an important development that needs to be replicated and refined.

17.9.2 *The relation between life-satisfaction/happiness and GDP*

The relation between life-satisfaction/happiness and GDP lies at the heart of the subject of subjective well-being. Many surveys prefer the word *life satisfaction* to *happiness* as it may be considered to invoke lesser emotional feelings. Recent research has thrown the conventional wisdom in the area into a state of flux. In this section, we shall primarily be interested in the relation between happiness and income. However, there might also be a relation between happiness and wealth; for the emerging evidence, see Senik (2014).

Early work in this area is associated with the work of Easterlin (1974). Van Praag (1968, 1971) had already started analyzing issues of subjective well-being using survey questionnaire techniques, a few years earlier. The results can be illustrated using a more updated dataset for the US, shown in Figures 17.5, 17.6. The data are from the United States General Social Survey. The self-reported measure of happiness takes the form of the following question: “Taken all together, how would you say things are these days—would you say that you are (3) very happy, (2) pretty happy, or (1) not too happy?”

Figure 17.5 shows that for the cross-sectional data from 1994, average US happiness increases with per capita real household income. So the richer appear to be happier. Furthermore, the curve shows diminishing marginal utility of income if we accept that happiness reflects utility. From




Figure 17.5 Mean happiness and real household income for a cross section of Americans in 1994.

Source: Di Tella and MacCulloch (2006) with permission from the American Economic Association.




Figure 17.6 Relation between mean happiness and real per capita GDP for the US between 1975 and 1997. Source: Di Tella and MacCulloch (2006) with permission from the American Economic Association.

this, we might expect that, since per capita real GDP has an upward trend, average happiness should increase over time. However, from Figure 17.6, we see that, while GDP per capita increases over time, average happiness is flat. This has come to be termed as the *Easterlin paradox*.

When Figure 17.5 is drawn for a cross section of countries, using the *World Values Survey*, the relationship between life satisfaction and GDP is relatively flat for incomes above \$10,000. Hence, further increases in income do not lead to appreciable increases in well-being; see for instance Figure 4 in Frey and Stutzer (2002) but also Figure 17.11. The idea is that income adds to happiness until basic needs are met, but not beyond that point (Veenhoven, 1991). Deaton (2008) presents the alternative idea, namely, that only after basic needs have been met can individuals derive happiness for cultural and intellectual pursuits. This view will have a bearing for the emerging results that we review below.

For a replication of the results in Figure 17.6 to other developed countries such as Japan, France, Germany, United Kingdom, Switzerland, see Easterlin (1995), Frey and Stutzer (2000), Blanchflower and Oswald (2004), Myers (2000), and Di Tella et al. (2001).⁸ Regressions of life satisfaction on lagged incomes show that individuals adapt to changes in income over a relatively short span of a few years. Based on German micro-level data, Di Tella et al. (2010) find complete adaptation within four years.

It is perhaps useful to see the results from Japan, over the period 1958–1991, in Figure 17.7. The data is drawn from Penn World Tables and the World Database of Happiness. Self-reported happiness was measured on a four point scale; average happiness was 2.7 in 1958, and despite spectacular growth in per capita real GDP, was 2.7 in 1991 as well. These results are very similar to those in Figure 17.6 and reinforce the Easterlin paradox.

The cross-sectional findings that we report above, only establish a correlation between happiness and income but do not establish any causal direction. One may expect that higher income makes people happier. However, there could also be reverse causation in which happier people earn more, perhaps because they are more productive or are able to deal with job stress better (Kenny, 1999; Graham et al., 2004).

⁸ For a book length treatment of this evidence, see Layard (2005).




Figure 17.7 The Easterlin paradox for Japan, over the period 1958–1991.

Source: Frey and Stutzer (2002) with permission from the American Economic Association.

One natural experiment to establish the direction of causality is to examine the before and after happiness of people who unexpectedly win the lottery. Gardner and Oswald (2007) use the British Household Panel Survey (BHPS) to examine this question. The survey contains a mental well-being measure based on 12 questions. For each question, the lowest well-being score is 3 and the highest is 0. Thus, the aggregate worst well-being score is 36 and best is 0. A simple sum of these scores across the 12 questions for each subject gives the General Health Questionnaire score (GHQ) for the subject. Using data over 1998–2001, the authors cull out individuals who have won differing amounts of money in the lottery. There are 4822 observations for individuals who have won a small amount (less than £999) and 137 observations on medium-sized wins (greater than £1000).

The results are shown in Figure 17.8 for the three groups titled: No win (the control group), small-sized wins (£1–£999), and medium-sized wins (£1000 and above). The mean change in GHQ scores is reported for each group and for various time intervals, assuming that the lottery win takes place in period T. The standard errors are reported in the corresponding figures on the right.

In the first row in Figure 17.8, labeled (a), there is an increase in mental stress in the year of the win, particularly for medium-sized wins. In the second row, labeled (b), there is no difference in the GHQ score changes for non-winners and winners of small amounts in the year immediately after the win. However, there is an increase in mental well-being (fall in GHQ score) for individuals who have medium-sized wins. This is even more pronounced if one considers the two-year interval on either side of the win. The standard errors allow one to reject the null hypothesis at 5% that there is no change in the GHQ score for medium-sized winners. In a nutshell, small wins




Figure 17.8 Before-win and after-win happiness of people who unexpectedly win the lottery.

Source: Reprinted from *Journal of Health Economics* 26, 49-60. Gardner, J. and A. J. Oswald, "Money and mental wellbeing: a longitudinal study of medium-sized lottery wins." Copyright © 2007, with permission from Elsevier.

do not matter but larger wins raise happiness. In this case, the causality runs from income to happiness. For related findings see also Brickman et al. (1978).

Another exogenous variation in income is explored by Frijters et al. (2004). They find that after the fall of the Berlin Wall, the resulting increase in income in East Germany had a positive causal effect on reported happiness.

What accounts for the Easterlin paradox? Drawing on the work of Veblen ([1899] 1973) and Duesenberry (1949), Easterlin (1974, 1995, 2001) gives an explanation in terms of *relative income concerns* and *shifting aspiration levels*. Put simply, people feel happier with higher incomes but if everyone's income increases, then, given sufficient time, happiness does not increase. Aspiration adaptation theories postulate that the happiness of individuals depends on the gap between the actual realization of their income and their aspiration level; the higher the gap, the greater is the happiness. Furthermore, increases in income give temporary increases in happiness as the aspiration level adjusts upwards, eventually swamping the increase in happiness; this is sometimes known as the *hedonic treadmill* (Brickman and Campbell, 1971; Diener et al. 2006).

These adaptation features are similar to models of habit formation that we have studied in this chapter. Another explanation is based on the *set point theory of happiness*. In this theory, individuals have some inherent level of happiness; temporary shocks can lead to short-term departures in happiness but happiness eventually returns to its inherent level (Easterlin, 2003). The finding that approximately 33% of the variation in life satisfaction is on account of genetic factors is consistent with this explanation (De Neve et al., 2012). Furthermore, there have been attempts to isolate a gene for happiness (De Neve, 2011).

We have considered utility functions that depend on relative incomes in Part 2; see, for instance, the model of Bolton and Ockenfels (2000). Clark and Oswald (1996) provide evidence that British workers form an income reference group with other workers who have shared labor market characteristics. Knight et al. (2009) use Chinese household data to show that the main reference group is other households in the same village. For empirical evidence that shows the dependence of well-being on the difference between actual income and income aspirations, see, for instance, Stutzer (2004) and Knight and Guanatilaka (2012). Luttmer (2005) shows that controlling for an individual's income, higher incomes of neighbors reduce self-reported happiness.

Kuhn et al. (2011) exploit the exogenous variation in income caused by the Dutch postcode lottery in which the set of winners are all the buyers of the lottery tickets in a particular postcode. Each winner wins 12,500 euros per ticket and one of the postcode winners is chosen to additionally receive a brand new BMW. The data allow for tracking the consumption patterns of the winners and their neighbors. The neighbors of winners who receive a BMW, are also more likely to buy a BMW. Angelucci et al. (2009) also find positive social spillover externalities of bimonthly conditional grants to rural communities in Mexico, designed to improve their education, health, and nutrition.

The issue of reference groups is of independent interest and is needed to also impart greater empirical content to models of other-regarding preferences in Part 2 of the book. In this regard, the work of Clark and Senik (2010) is particularly important. They use data from 18 European countries in Wave 3 of the European Social Survey in which income comparison questions were asked for the year 2006–2007; they restricted the sample to the working age population, 16–65. We now briefly describe their methods and results.

The two relevant questions asked in the Survey were the following. (1) “How important is it for you to compare your income with other people's income?” The answers could be chosen on a discrete scale of 0 (not at all important) to 6 (very important). (2) “Whose income would you be most likely to compare with your own? Please choose one of the groups on this card: Work colleagues/ Family members/ Friends/ Others/ Don't compare/ Not applicable/ Don't know.” The first question about the *intensity of comparison* was answered by 19,053 respondents; 25% choose 0 (comparisons not important). However, for the remaining 75% of the subjects, income comparisons are important. Each of the numbers 1–4 is chosen by approximately 15% of the subjects; the numbers 5 and 6 are chosen by approximately 10% and 5% of the subjects, respectively.


The responses to the second question are shown in Table 17.1. Here, 35.9% of the 18,936 respondents claim that they do not compare their incomes. In Q1, this percentage was 25, so there is a degree of time inconsistency of choices for some respondents. However, 73% of those who chose 0 in Q1, choose Don't compare in Q2, so they are consistent. Of those who compare, the vast majority compares their income with work colleagues and the next most significant group is friends.

An important finding is that there is a negative correlation between intensity of comparison (higher numbers chosen on Q1) and real per capita GDP across countries as well as within countries; see Figure 17.9. A regression analysis shows that poorer households compare their

Table 17.1 Responses to the income comparison question.

	Observations	%
Work colleagues	6874	36.3
Family members	1103	5.8
Friends	2825	14.9
Others	1344	7.1
Don't compare	6789	35.9
<i>Total</i>	18936	100.0

Source: Clark and Senik (2010).

**Figure 17.9** Comparisons of own incomes with the incomes of others in a cross section of countries with different income levels.

Source: Reprinted from Clark, A. E. and Senik, C. (2010), "Who compares to whom? The anatomy of income comparisons in Europe." *The Economic Journal* 120: 573–94. © Clarke, A. E. and Senik, C. Journal compilation © Royal Economic Society 2010.

incomes more to others. Self-employed compare less, singles compare more, and there are no noticeable gender effects. There are differences in cross-country effects of the intensity of comparison; Eastern European countries compare more relative to Western European countries.

Regarding Q2, self-employed subjects who compare are more likely to compare with family members and friends. The self-employed, on average, also appear happier as compared to the employed, possibly due to the greater job autonomy that they have; see Frey and Benz (2008) for the evidence. The employed are more likely to compare with colleagues. Women are more likely to compare with family members and income comparisons increase among the 25+ age group. Eastern Europeans are less likely to compare to other family members; such a comparison is more

likely in Spain, Poland, Finland, and Ireland. Countries in central continental Europe compare more with colleagues.

Those who compare more are less happy. The correlation coefficient between the intensity of income comparison and subjective happiness is negative and significant. The well-being correlation with income comparisons is smaller for richer as compared to the poorer individuals. The correlation of income comparisons with well-being depends on the reference group. Individuals who compare with colleagues are significantly happier than others who compare with friends. One possibility is that comparison with colleagues signals to individuals that their future incomes might be higher, but that is not true in a comparison with friends. Finally, those who compare more are also more likely to favor greater income redistribution.

Recent years have, however, seen a re-evaluation of the evidence considered above. This re-evaluation is based on alternative sources of information, such as the Gallup World Poll in 2006, and a consideration of the results based on a fuller set of countries and time periods. The results from the Gallup World Poll have a much fuller list of poorer countries as compared to the World Values Survey. These results, as we see now, question the very existence of an Easterlin paradox.

The Gallup World Poll asks questions about *life satisfaction* on an 11 point scale, while the World Values Survey asked questions about *happiness* on a 3 point scale. It is common practice in much of the literature to treat life satisfaction and happiness as synonymous; see, Van Praag et al. (2011) for supporting arguments. However, Deaton (2008) distinguishes between the two. He argues that life satisfaction involves making an overall evaluation of one's life, while happiness reflects an effect that can be gauged from *experiential questions* such as asking the individual if they have been smiling a lot, or if they have been depressed?




Figure 17.10 Relationship between mean life satisfaction and income, measured with Gallup World Poll data. GDP per capita in 2003 is measured in purchasing power parity chained dollars at 2000 prices.

Source: Deaton (2008) with permission from the American Economic Association.




Figure 17.11 The relationship between life satisfaction and income.

Source: Deaton (2008) with permission from the American Economic Association.

Figure 17.10 shows the relation between mean life satisfaction from the Gallup World Poll and the log of per capita GDP for 2003, for three different age groups. The sizes of the circles represent the associated country sizes. The relation is approximately linear. When the absolute level of per capita GDP is used, the relationship between life satisfaction and income is increasing and concave. This is shown in Figure 17.11 in which the data from the World Values Survey (10 point scale) is shown as the shaded circles and the Gallup World Poll data (11 point scale) is shown as hollow circles.

Deaton (2008) draws the following conclusions from Figures 17.10 and 17.11.

1. From Figure 17.11, individuals in richer countries in North America, Japan, Australasia, and Western Europe have higher mean life satisfaction (7.5–8.5 on an 11 point scale) relative to those in poorer countries in sub-Saharan Africa, Haiti, and Cambodia (3.1–4.5). In general the relation between life satisfaction and income is upward sloping.
2. Earlier research based on World Values Survey identified a cutoff income of about \$10,000 (Layard, 2003) or \$20,000 (Layard, 2005) beyond which income is unrelated to happiness. This assertion was widely made; see, for instance, Diener and Seligman (2004), Di Tella et al. (2003), and Frey and Stutzer (2002). Figure 17.10 paints a contrary picture. Indeed there is an almost linear relation between life satisfaction and log of income at all income levels. If anything, the relationship is the steepest for very high income levels. This is supported by the regression results in Deaton (2008).
3. The main reasons for the difference in results between the World Values Survey and the Gallup World Poll are as follows. First, the relative coverage of poor countries in the Gallup Poll is much richer, as reflected in Figure 17.11; there are hardly any filled-in circles for very low levels of income. Second, there is a disproportionate share of Eastern European countries in the World Values Survey who had middling incomes as well as great dissatisfaction

following the events of the breakdown of the former Soviet Union. Third, the World Values Survey in relatively poor countries of that time (Ghana, Nigeria, India, China) sampled the relatively better-off individuals in the poorer countries, in order to ensure comparability with the individuals in the richer countries. Presumably these individuals had much higher levels of happiness relative to the representative individuals in these countries. But this led to a much flatter relation between income and happiness.

4. Life satisfaction is age-dependent. The life satisfaction for the 15–25 year olds is highest, while that of individuals aged 60 or above is the lowest. Life satisfaction declines with age for most countries in the world except for the highest income countries such as the US, Canada, the UK, Australia, and New Zealand, where it is U shaped; see Deaton (2008) for details.

Stevenson and Wolfers (2008, 2013) and Sacks et al. (2013) also provide evidence that supports the analysis of Deaton (2008). Figure 17.12 gives more recent data based on a sample of 155 countries from the first five waves of the Gallup World Poll over the period 2008–2012; it captures 95% of the world's population.

The simple correlation between life satisfaction and log of per capita GDP is 0.79; the regression coefficient in a simple OLS regression is 0.335 with a standard error of 0.018. The regression line is shown as the unbroken line in Figure 17.12, while the dashed line is a non-parametric line. The two lines are closely aligned. In particular, the non-parametric line does not flatten out beyond some income threshold. Choosing a threshold level, say, \$15,000, and fitting separate linear regressions for countries below and above the threshold, the authors are unable to reject the null hypothesis that the regression coefficients are identical. Indeed, the relationship between life satisfaction and income appears quite steep when life satisfaction is plotted against income for the 25 largest countries; see Figure 17.13.




Figure 17.12 The relationship between life satisfaction and per capita GDP for 95% of the world's population.

Source: Stevenson and Wolfers (2013) with permission from the American Economic Association.




Figure 17.13 Relation between life satisfaction and income for the 25 largest countries.

Source: Stevenson and Wolfers (2013) with permission from the American Economic Association.

These results are supported by Inglehart et al. (2008), Diener et al. (2013), and Veenhoven and Vergunst (2013). The response by Easterlin et al. (2010) and Easterlin (2013) is that the positive relation between income and happiness is a short-term one and in the long term (10 years or more) there is no relation between the two. For instance, Easterlin et al. (2010) write in their abstract: “Recent critiques of the paradox, claiming the time series relationship between happiness and income is positive, are the result either of a statistical artifact or a confusion of the short-term relationship with the long-term one.” However, this conclusion has been contested by the proponents of a positive relation between happiness and per capita GDP. Beja (2014) proposes a compromise. He uses very long time series data (1973 to 2012) for a cross section of European countries and finds that, while the relation between happiness and income is a positive one, it is weak and economically not significant enough.

In light of these disparate results, the Easterlin paradox no longer enjoys undisputed acceptance in the profession, which it did a decade ago.

17.9.3 Other important economic correlates of well-being

We consider a range of correlates of well-being in this section. These include, unemployment, inflation, income inequality, social capital, health, and democratic institutions. An alternative approach is to consider *domain satisfaction*. This refers to a more disaggregated measure of satisfaction that studies the determinants of satisfaction in individual domains such as job satisfaction, and financial satisfaction (we have considered health satisfaction above). For a further discussion of these issues and the relevant empirical evidence, we refer the reader to the survey by Van Praag and Carbonell (2011).

UNEMPLOYMENT AND INFLATION

Di Tella et al. (2001) explore the foundations of the standard social welfare function in macroeconomics that typically proposes quadratic losses in unemployment and inflation. They regress measures of happiness on inflation and unemployment, using a very large cross-country dataset

from 12 European countries and the US; as is typical in these studies, they allow for country fixed effects and time trends.⁹

Both unemployment and inflation are found to reduce happiness. A one percentage point increase in unemployment from 9% (the mean unemployment rate at the time of the empirical study) to 10% reduces life satisfaction by 0.028 when happiness is measured on a scale of 1 to 4. On the same scale, a one percentage increase in inflation from the sample mean rate of 8% to 9%, reduces happiness by 0.01. Both effects are substantial. However, unemployment is relatively more deleterious for happiness. On average, individuals are willing to trade a 1 percentage point increase in unemployment against a 1.7 percentage point increase in the inflation rate.

These results have been replicated. Using data from 16 European countries, Wolfers (2003) finds that a one percentage point increase in unemployment causes 4.7 times as much unhappiness as a one percentage point increase in the inflation rate. Reference point effects have been identified in the observed aversion of unemployment. The happiness reducing effect of unemployment is smaller if the fraction of other people in the region who are unemployed is larger, or a partner is unemployed; presumably this reduces the stigma attached with being unemployed (Clark, 2003).

Unemployment is possibly more aversive to happiness than any other single characteristic, even more aversive than divorce or separation, and relatively more aversive for the more educated (Clark and Oswald, 1994). The causality runs from unemployment to unhappiness, rather than the other way around (Winkelmann and Winkelmann, 1998). It appears that the unhappiness of the unemployed cannot be accounted for, purely by their lower income. Thus, there are psychic costs of unemployment.

Psychic costs arising through anxiety, depression, loss of self-esteem, resort to alcoholism, stigma costs, scarring costs of unemployment, and strained relationships all appear to play a role in greater unhappiness.¹⁰ Luechinger et al. (2010) test for anticipatory anxiety among the unemployed by comparing private and public sector employees in Germany. Public sector employees have greater job security. For this reason, private sector workers exhibit relatively greater anxiety towards general economic shocks. For state-level panel data, Ruhm (2000) finds that there is a positive correlation between unemployment and the number of suicides in a state.

The picture emerging from this analysis shows up the inadequacy of macroeconomic explanations, in the tradition of real business cycle models, that solely use the transmission channel of voluntary unemployment as an explanation of unemployment. If people were voluntarily unemployed, it would be difficult to account for their pronounced unhappiness and low self-reported measures of well-being. This not only illustrates the usefulness of happiness economics, but also raises uncomfortable questions about the dominant practice in macroeconomics for so long.

INCOME INEQUALITY AND HAPPINESS

While the relation between happiness and income is well documented, there is comparatively less work on the relation between inequality and happiness. Alesina et al. (2004) find that inequality reduces happiness. However, while the relation is strong and statistically significant in the case

⁹ The European data is taken from the Euro-Barometer Survey Series, the US data is taken from the United States General Social Survey, and the data on unemployment and inflation is taken from OECD statistics. The dataset used in the study contains about a quarter of a million individuals.

¹⁰ For the importance of non-pecuniary costs of unemployment, see Kassenboehmer and Haisken-DeNew (2009). For the relation of non-pecuniary costs with social work norms, see Clark (2003) and Stutzer and Lalive (2004). For limited adaptation to unemployment and its scarring role, see Clark et al. (2001), Clark (2006), and Knabe and Rätzel (2011).

of Europe, it is far weaker in magnitude for the US. The authors conjecture that upward income mobility is greater in the US relative to Europe. So individuals who are poorer, do not see inequality as unduly limiting their future prospects in the US, hence, inequality has a lesser impact on their happiness.

HEALTH SATISFACTION AND HAPPINESS

Figure 17.14 shows the relation between health satisfaction and per capita GDP. The non-parametric regression curves are upward sloping, indicating that higher income is associated with greater health satisfaction. Health satisfaction falls with age. For the richest countries, the 50–59 age group has a lower health satisfaction relative to the two older groups.

Countries from the former Soviet Union exhibit extremely low health satisfaction. Other countries in the low health satisfaction group include Haiti, Rwanda, Burundi, Uganda, Cambodia, and Benin. Deaton (2008) reports that a fall in life expectancy due to HIV/AIDS cannot fully explain the poor levels of health satisfaction in sub-Saharan Africa.

Causality need not just run from good health to happiness. Indeed happier people are likely to be healthier (Diener and Chan, 2011). Other variables may be correlated with both; hence, the relationship between health and happiness may reflect these other factors, particularly genetic factors (De Neve et al., 2012). There is also significant adaptation to health disabilities. Oswald and Powdthavee (2008) report that within three years, individuals adjust to 30% of the effect of severe disability and 50% of the effect of moderate disability.

Oswald and Wu (2011) use state-level data for the US that is based on telephone surveys of the mental health of the respondents and questions about their life satisfaction and incomes in the previous month. They find a high degree of heterogeneity among the US states. Particularly interesting is the result that the high GDP states in the sample are less happy; Louisiana is one




Figure 17.14 Non-parametric regressions of average health satisfaction against per capita GDP for seven age groups. GDP is measured in purchasing power parity chained dollars at 2000 prices.

Source: Deaton (2008) with permission from the American Economic Association.

of the happier states, while California is not happier than the other states. The authors' preferred explanation is based on Adam Smith's idea of compensating wage differential. States with lower life satisfaction have to offer a compensating income differential. However, there is no correlation between mental health and per capita state GDP.

SOCIAL CAPITAL AND HAPPINESS

Social capital refers to the effects of social networks, including the associated norms of trust and reciprocity, on production efficiency and well-being (Coleman, 1993; OECD, 2001). Helliwell and Putnam (2004, p. 1436) write: "Social capital can be embodied in bonds among family, friends and neighbors, in the workplace, at church, in civic associations, perhaps even in Internet based 'virtual communities.'" They use cross-sectional data from the World Values Survey, the European Values Study, and the Canadian Benchmark Survey comprising of 29,000 observations. A range of variables is considered at the national/community level and at the individual level, such as average trust, trust in neighbors and police, membership in associations, governance, the importance of God/religion, and marital status.

The authors report a range of findings that are consistent with the importance of social capital for well-being. Being married improves happiness and life satisfaction, particularly in comparison with those who are divorced or separated; this reflects the effect of family level social capital. Church attendance reflects community level social capital. Those who attend Church and have a strong belief in God, trust others more. Frequent interactions with friends and neighbors are associated with higher well-being. Those who report that they live in a high trust neighborhood report higher life satisfaction and happiness as well as better health.

This literature continues to expand. For the effects of social capital on long-term trends in subjective well-being, see Sarracino (2010); for the positive effects of volunteer work on well-being, see Meier and Stutzer (2008); and for the public goods nature of social capital, see Becchetti et al. (2008).




Figure 17.15 The relationship between economic freedom, democracy, and well-being and happiness.

Source: Rode, M., Knoll, B., and Pitlik, H. (2013). "Economic freedom, democracy, and life satisfaction," in James Gwartney, Robert Lawson, and Joshua Hall (eds.), *Economic Freedom of the World: 2013 Annual Report* (Fraser Institute): 215-33. www.freetheworld.com. Reprinted with kind permission of The Fraser Institute.

DEMOCRACY, FREEDOM, AND LIFE SATISFACTION

Economic freedom and democracy might also directly contribute to well-being and happiness. This is illustrated directly in Figure 17.15; the left panel shows life satisfaction in autocracies and democracies, and the right panel shows life satisfaction under different economic freedom ratings (higher rating shows higher freedom).

Economic freedom is measured by the EFW index published in *Economic Freedom of the World: 2012 Annual Report*. Life satisfaction is measured by the World Values Survey. The classification of countries into autocracies and democracies is based on Cheibub et al. (2010). In a regression analysis, Rode et al. (2013) show that the effects of economic freedom and democracy on life satisfaction are significant at 1%. One possible transmission channel is that in a democracy, politicians respect the choices of people, giving the latter a feeling of well-being (Frey and Stutzer, 2002).

Veenhoven (2000) shows that for 38 developed countries, indices of political freedom, economic freedom, and personal freedom have a statistically significant correlation with happiness. Furthermore, economic freedom enhances happiness relatively more in poorer countries and political freedom relatively more in richer countries. Direct democracy refers to the exercise of direct citizen choices through referenda. Frey and Stutzer (2000) show that Swiss cantons with higher direct democracy indices are also associated with higher self-reported happiness.

CHAPTER 18

Interactions between Emotions and Cognition

18.1 Introduction

Observed human behavior is the outcome of interactions between a *cognitive* and an *emotional* module/system in the brain. The emotional system is thought to be quick, reactive, and automatic; it is typically modeled as operating at the hot end of the spectrum of human emotions. The cognitive system is slow, deliberative, focuses on the larger goals, and operates at the cold end of the spectrum of human emotions. The human emotional system is shared in common with many other mammals, but the cognitive system is unique to humans, although some elementary functions are shared with close primate relatives such as chimps. The cognitive system, developed about 150,000 years ago, and was superimposed on top of the existing emotional system. The objectives of the two modules can be in harmony or, in the more interesting case, in conflict (Tooby and Cosmides, 1992; Wilson, 2012). In this chapter, we consider the economic implications of competition and cooperation between emotions and cognition.

Section 18.2 introduces the reader to models of *dual brain modules*, or simply dual process models. In Section 18.2.1, we introduce the main premise of these models. In Section 18.2.2, we describe a general blueprint for the interaction between the emotional and the cognitive processes (Loewenstein and O'Donoghue, 2004, 2007). The blueprint is speculative in several respects, does sometimes make heroic assumptions, and one cannot always disentangle the source of observed behavior as cognitive or emotional. However, the model is an important first pass, relies on plausible ideas, is very general, and suitable for serving as a template for the development of further ideas. It has the potential to derive simple relations and concepts in behavioral economics, endogenously, and in a simple manner, playing a role similar to that of IS-LM analysis in macroeconomic analysis. We show that the model can give important insights into behavioral models of risk, time preference, and social preferences. This model deserves to be taken more seriously.

In Section 18.3, we model the interaction between a patient long-run self and a myopic short-run self in a simple consumption–savings problem (Fudenberg and Levine, 2006). In the absence of a commitment technology, the long-run self can control the consumption splurges of the short-run self through the exercise of costly willpower. We examine the dependence of optimal consumption plans on willpower costs. In a richer model, we then consider the presence of a commitment technology. Here, we show that for small windfalls in income, an increase in unanticipated income is fully consumed, while the converse holds true for large windfalls. In the absence of any costs of control, the individual never fully splurges on consumption.

It is possible that parts of the brain could have access to private information that other brain areas might not possess (Berns et al., 1997). In other contexts, it is possible that individuals might have restricted access to their own motives and beliefs (Brocas and Carrillo, 2008). People might infer their own emotions or abilities from their actual behavior. They could also take their donation of money as evidence of their altruism, or their ability to jog in extreme heat as evidence of their willpower (Bodner and Prelec, 2003).

In Section 18.4, we consider a somewhat speculative two-period model of interaction between a principal (cognitive system) and an agent (emotional system) when the agent has asymmetric information about its preferences (Bodner and Prelec, 2003; Brocas and Carrillo, 2008). We contrast the allocation of consumption and effort under full information with the case of asymmetric information. It turns out that an endogenous rate of time preference develops, despite the fact that the cognitive system undertakes no time discounting. The model does leave open several questions that are not answered in a fully satisfactory manner. For instance, how are contracts between competing brain modules enforced? How should one treat the participation constraint of a brain module? What does renegotiation of contracts mean in this case?

When a current self interacts with her future selves, the decision to acquire information has considerable strategic importance. By acquiring information, one makes all of one's own future selves aware of that information. If there is a possibility that the future selves will use this information to the detriment of the current self, then the current self might choose to be strategically uninformed. For instance, individuals vastly overestimate the effect of smoking on lung cancer (Viscusi, 1990; Lundborg and Lindgren, 2004). However, such information is freely and easily available. Why then do individuals choose not to acquire such information? One possibility, that we consider in Part 3, is that individuals procrastinate to acquire the relevant information (O'Donoghue and Rabin, 1999).

But in this chapter we consider alternative explanations. Section 18.5 considers formal models that address these issues within the framework of multiple-selves models. We first consider, in Section 18.5.1, a simple three-period model of strategic information transmission (Carrillo and Mariotti, 2000). The current self has to decide whether to learn something of strategic importance that determines the consumption behavior of future selves. We derive the conditions under which such learning is justified, relative to being strategically uninformed.

In models of asymmetric information in economics, information asymmetry typically pertains to the types of other players. Psychological evidence, however, suggests that individuals often have doubts about who they really are. Individuals often engage in introspection and reflection in an attempt to figure out the motivation, justification, and the implications of their current and past actions. Individuals often question their ability to successfully undertake a task, and sometimes struggle to discover their true underlying preferences.

Actions taken by individuals, who have asymmetric information about some relevant characteristics of their own self, are sometimes characterized by self-deception and at other times by repentance. Another relevant feature is the presence of *imperfect willpower* in humans; for instance people might lack the willpower to reduce or to give up smoking, or keep to a deadline. Finally, people might have *imperfect recall* and *motivated cognition*, i.e., they actively seek ways of managing the awareness about the degree to which subsequent selves can recall past events. In Section 18.5.2, we formally consider these issues in a model where individuals have asymmetric information about their own abilities, motivation, or self-confidence and they attempt to strategically manage information available to their future selves (Bénabou and Tirole, 2002, 2004). Among other predictions, this model has the potential to explain the roles of strategic ignorance, selective memory transmission, and self-handicapping.

18.2 Emotions and a two-modules view of the brain

In this section, we explore the implications for human behavior when the brain is viewed as comprising of two modules: an emotional module and a cognitive module. While our focus will be on economic models of more recent vintage, there is also an older literature on these models in psychology (Schneider and Shiffrin, 1977a, 1977b).

18.2.1 *An introduction to the two-modules view*

One may imagine the brain as a group of modules, or a *confederation of mechanisms* in the language of Cohen (2005), that jointly determine human behavior. These modules could have different objectives. The desired behavioral outcomes of the competing modules may sometimes be in agreement, while at other times, these could be in conflict. The premise in these models is that observed human behavior is the outcome of competition among these modules.

A useful abstraction, that has a long history in many diverse literatures, is to consider the brain as having just two main modules. One controls emotions, while the other has a cognitive function; see Loewenstein and O'Donoghue (2004, 2007) for a lucid and informative review that we draw on here. Adam Smith used the terms *passions* versus *impartial spectator* to capture the emotions versus cognitive contrast. *Passions* describe motivational forces and other visceral states such as anger, hunger, and thirst. The *impartial spectator* refers to an ability to stand back, or be emotionally neutral in some sense, and take an impartial view of one's emotional motivations. Freud (1924) referred to a similar conflict between the *id* and the *ego*. Such dual process models are also common in modern psychology; see for instance Chaiken and Trope (1999).

In common usage, we often speak of the *feeling of being in two minds*, or the *conflict between the heart and the mind*; these phrases appear to illustrate the conflict between emotions and cognition. The conflict between emotions and cognition is likely to have an evolutionary explanation (Dawkins, 2006; Tooby and Cosmides, 1992; Massey, 2002; Wilson, 2012). When humans departed from their primate relatives in their evolutionary history, about 6 million years ago, the human brain was not redesigned from scratch. Cognitive abilities, the ability to think about broader goals, and the broader implications of one's actions, were developed by adding brain structures on top of the existing brain structures. The most notable new brain structure in this regard is the prefrontal cortex, which has developed significantly in relatively recent evolutionary history.

Emotionality clearly preceded rationality in the evolutionary sequence. Define emotions for the moment, as in Cohen (2005), as low level automatic psychological processes that are engaged by stimuli (or memories) and are associated with strong positive and negative psychic utility. For instance, witness the emotion of fear that is triggered by observing a coiled object on a dark street.

As rationality developed, it did not replace emotionality as the basis of human interaction. Rather, rational abilities were gradually added to the existing, and simultaneously developing, emotional capacities. Indeed, the neural anatomy essential for full rationality, the *prefrontal cortex*, is a recent evolutionary innovation. It only emerged in the last 150,000 years of a 6-million-year existence, representing only about 2.5% of humanity's total time on earth.

Three features of evolution, namely, that it is *efficient*, *opportunistic*, and *conservative*, are important in this regard; see for instance, Cohen (2005). Evolution is *efficient* in the sense that, conditional on the environment, it optimally meets the needs of the organism. It is *opportunistic*

in the sense that it produces solutions that are closely tailored to the environment.¹ Finally, it is *conservative* in the sense that it is reluctant to discard older solutions (e.g., existing brain structures) as new environmental situations arise, even if the older solutions are not well adapted to the new environment. Evolution develops additional solutions/structures to deal with new environments.

Metcalf and Mischel (1999) also make a distinction between a hot emotional module or process, and a cool cognitive module. The emotional response is rapid, automatic, simple, relatively inflexible, and attuned to particular stimuli. The cognitive response is just the opposite, i.e., it is reflective, deliberate, slower, and more flexible. The former is limited to a small number of goals, while the latter can generate behavior that is consistent with a large number of broader goals. Cohen (2005) describes this as the familiar trade-off between speed and specialization that is faced by engineers in the design of computers.

18.2.2 *An illustrative framework*

We now consider a general framework for the interaction between the emotional and the cognitive system, following Loewenstein and O'Donoghue (2004, 2007). We begin by noting the drawbacks of this approach. Many of the assumptions in the illustrative applications may appear arbitrary, yet they are plausible. Nor can one always cleanly separate observed behavior based on the source, cognitive or emotional, that gives rise to it. This framework cannot be used to fully justify the importance of the dual process model for explaining behavioral phenomena, nor make watertight predictions. Yet, it does take an important first step in this direction and has the potential to serve as the basis of further models that can lead to much needed theoretical developments in the area.

The emotional and cognitive processes/systems are also respectively termed as the *affective* and the *deliberative* processes/systems. The emotional system, because it is automatic and quick, initially comes into play when an individual faces some stimulus (Baumeister and Vohs, 2003). The stimulus could be external (coiled object on a dark street) or internal (drop in blood glucose levels).

The cognitive system, because it is reflective and considers broader goals, can influence the emotional system by the exercise of costly *willpower*. Willpower has three relevant features for our purpose. First, it is a limited resource and can get depleted. Second, cognitive load reduces the ability to exert willpower. Third, willpower can be depleted by stress; for instance, Shiffman and Waters (2004) find that relapse among cigarette smokers was often preceded by a stressful event.²

The second feature of willpower is suggested by the experiments of Shiv and Fedorikhin (1999) who employed two groups of subjects. The first group was subjected to a high cognitive load and was asked to memorize a 7-digit number. The second group was subjected to a low cognitive load and asked to memorize a 2-digit number. The subjects were then required to walk to another part of the building to report their memorized number. On the way, they were presented with a high calorie cake and a fruit salad. A larger fraction of subjects who were subjected to a high cognitive load (63%) chose the cake, relative to the those who were subject to a low cognitive load

¹ For a more detailed discussion of these issues, see Part 5 of the book where we distinguish between changes in the environment that occur with different speeds. Long-run changes induce changes in the structure of the DNA. But for short-run and medium-run changes, other responses, such as gene-culture coevolution, are more efficient.

² For a model of the effect of stress on reducing well-being that works through a cognitive load channel, see Walde (2015).

(42%). There was only a slight difference between the two conditions when the same objects were presented symbolically in the form of photographs.

The cognitive system operates at the cool end of the spectrum and deliberates about broader goals. Thus, as a first pass, assume that it has the same objectives as the rational agent in neo-classical economic models. So, for instance, when the cognitive system faces a risky lottery, it uses expected utility theory to evaluate the lottery, and when it faces an intertemporal stream of consumption, it uses exponential discounting to evaluate the consumption stream.

The emotional system operates at the hot end of the spectrum and focuses on immediate goals. As a first pass, assume that it behaves in a manner that is consistent with the recent evidence in behavioral economics. So it uses loss aversion when presented with gains and losses and uses hyperbolic discounting when evaluating an intertemporal stream of consumption.

The cognitive system cares about the moral and ethical dimensions of a situation, while the emotional system is driven between the two extremes of purely self-regarding and purely altruistic behavior. Which of these extremes is activated depends perhaps on a set of cues which activate such emotions. In related literature, Bernheim and Rangel (2004) also assume that the brain operates in one of two states, a cold state or a hot state. The state depends stochastically on the environmental conditions (which in turn could depend on actions taken in the past). However, their primary focus is on the cold state, which maximizes expected discounted utility, while the hot state is assumed to follow a simple rule. They take as their welfare measure, the preferences of the cold state.³

The schematic outline of the dual process model is shown in Figure 18.1. A stimulus, s , which belongs to a set S , activates the emotional and the cognitive systems. In particular, it activates an emotional/motivational state $m(s)$, and a cognitive state $c(s)$; the set of emotional and cognitive states could themselves be subsets of S . The problem is to choose the optimal behavior x from a set




Figure 18.1 Schematic representation of the dual process model.

³ In Benhabib and Bisin (2005), individuals can either invoke an automatic emotional system that is susceptible to temptation or a costly cognitive system that does not succumb to temptation. They explore an internal control mechanism that takes the form of a simple cutoff rule in a consumption-savings problem. Models of multiple selves (see Part 3 of the book) do not allow for this internal control mechanism.

$X \subseteq \mathbb{R}$, given that the two systems have different objective functions and none fully dominates the other.

In particular, the objective function of the emotional system is given by $U^e(x, m(s))$, which is maximized at $x^e \in X$. We assume that $\frac{\partial}{\partial m} \left(\frac{\partial U^e}{\partial x} \right) > 0$, i.e., an increase in the motivational state increases the marginal utility of the emotional system from increasing behavior, x . The objective function of the emotional system does not depend on the cognitive state, $c(s)$. The cognitive system has the objective function $U^c(x, m(s), c(s))$, which is maximized at $x^c \in X$. We assume that $\frac{\partial}{\partial c} \left(\frac{\partial U^c}{\partial x} \right) > 0$. Since the cognitive system cares about broader goals, it also takes account of the emotional states induced by the stimulus, $m(s)$.

For the sake of expositional ease, we will assume that U^e and U^c are strictly concave in x and continuous in all arguments, and the set X is compact, so that the maximizers x^e, x^c are unique, and continuous functions of the parameters.

Finally, observed behavior is the outcome of the interaction between the emotional and the cognitive systems. For instance, the proximity of tempting foods could create an emotional state of hunger. Faced with such a stimulus at lunchtime, the objectives of the cognitive and emotional systems are likely to be in agreement with each other. However, faced with such a tempting stimulus between lunch and dinner, the cognitive system of an individual on a diet is likely to recommend the opposite course of action to the emotional system. It is this latter situation where the desired actions of the two systems are in conflict, that is the most interesting for our purposes.

What determines the relative strength of activation of the emotional and the cognitive systems, when faced with a cue? Several factors play an important role. These include the following. (i) Physical proximity of the stimulus (Mischel et al., 2003). Food is more tempting if it is placed nearby, or if its consumption by others is observed. (ii) *Changes* versus *levels* of stimuli (Frederick and Loewenstein, 1999). The emotional system is more likely to respond to changes than to levels. (iii) Vividness of stimuli (Loewenstein et al., 2001). We are often more deterred from speeding if we directly observe an accident in graphic detail relative to reading about the annual road fatality statistics.

To economize on notation, we suppress the state, s , in the objective functions of the two systems by assuming a fixed state. We assume that the individual has sufficient time to choose actual behavior so that the eventual choice is made by the cognitive system. Using the standard terminology in economics, one may think of the cognitive system as a *benevolent principal* and the emotional system as the *agent*.

We shall primarily be interested in a static model. However, we first make a few observations on willpower in a dynamic model; index discrete time by t .⁴ Willpower costs exerted by the cognitive system are directly proportional to the extent to which the emotional system is moved from its most preferred behavior at time t , x_t^e , to some $x \neq x_t^e$ that might be in the individual's longer-term interest. Thus, the willpower, w_t , exerted by the cognitive system at time t is given by

$$w_t(x_t) = U_t^e(x_t^e, m) - U_t^e(x_t, m). \quad (18.1)$$

Let W_t be the stock of willpower at time t . The stock of willpower carried forward from time t to $t + 1$ is given by $W_{t+1} = f(w_t, W_t)$, such that $f_1 < 0$ and $f_2 > 0$. Thus, exertion of willpower effort in any period depletes the stock of willpower carried over to the next period ($f_1 < 0$).

⁴ An exercise at the end of the chapter, based on Ozdenoren et al. (2012), provides a more rigorous, but also more specialized, framework of analysis.

For instance, depletion of willpower can lead to a greater incidence of impulse buying (Vohs and Faber, 2007). Furthermore, *ceteris paribus*, a greater current stock of willpower transforms into a greater stock of willpower next period ($f_2 > 0$). These ideas are fundamental to an understanding of the role of willpower in a dynamic analysis (Baumeister and Vohs, 2003).

We shall only focus on the static case below. Omitting time subscripts, (18.1) can be written as

$$w(x) = U^e(x^e, m) - U^e(x, m). \quad (18.2)$$

The objective of the principal, or the cognitive system, is to choose $x \in X$ to maximize

$$V^c(x) = U^c(x, m, c) - \phi(W, \sigma)w(x), \quad (18.3)$$

where $\phi(W, \sigma) \geq 0$ represents per unit willpower cost, and σ is cognitive load/stress. Consistent with the evidence reviewed above, $\phi_1 < 0$ and $\phi_2 > 0$. Thus, per unit willpower costs are decreasing in the stock of willpower but increasing in the cognitive load/stress facing the individual. Substituting (18.2) in (18.3) and noting that x^e is fixed for a given stimulus (hence, omitting the term $U^e(x^e, m)$), we can write the maximization problem of the cognitive system as

$$x^* \in \arg \max_{x \in X} V^c(x) = U^c(x, m, c) + \phi(W, \sigma)U^e(x, m). \quad (18.4)$$

As $\phi(W, \sigma) \rightarrow 0$, $x^* \rightarrow x^c$ and as $\phi(W, \sigma) \rightarrow \infty$, $x^* \rightarrow x^e$. The cognitive system makes the eventual decision, so it is like a chief executive of the brain. However, the objective function in (18.4) also bears another interpretation if we recognize that for a fixed state, $U^c(x^c, m, c)$ and $U^e(x^e, m)$ are fixed numbers. Choosing x to maximize (18.4) is equivalent to minimizing the following.

$$x \in \arg \min [U^c(x^c, m, c) - U^c(x, m, c)] + \phi(W, \sigma) [U^e(x^e, m) - U^e(x, m)]. \quad (18.5)$$

The interpretation of (18.5) is that actual behavior is chosen in order to minimize the weighted sum of two costs. Namely, the cost of moving the emotional and the cognitive systems from their most preferred behaviors; the relative weight placed on the cost faced by the emotional system is $\phi(W, \sigma)$.

Our framework of analysis is based on the following three premises.

- A1:** The cognitive system behaves like the textbook economic agent that operates in a world of emotionless deliberation.
- A2:** Evidence from the behavior of animals who lack the prefrontal cortex and who share their primitive mammalian brain with us is used to infer the preferences of the emotional system.
- A3:** When the required evidence in A2 is not available, then the actual behavior of humans is used to recover the preferences of the emotional self.

We now apply this framework to three different domains: Intertemporal choice, behavior under risk, and social preferences. This gives rise to a range of applications. A partial list includes: the effectiveness of advertising even when it conveys no useful information, breakdown of bargaining when there are mutual gains to be shared, unpleasant divorces, and behavior in financial markets.

INTERTEMPORAL CHOICE

Assume that the cognitive system is driven by long-term payoffs and it uses exponential discounting. Evidence from animal studies suggests extreme myopia in pigeons. There is also evidence of myopia in monkeys (Tobin et al., 1996; Stevens et al., 2005). So the objective of the emotional system is also assumed to be myopic.

Formally, suppose that an individual faces a stream of dated consumption levels, $x_1, x_2, \dots, x_T$. In line with premise A1, the utility of the cognitive system is $U^c = \sum_{j=1}^T \delta^{j-1} x_j$, where δ is the discount factor. Using premise A2 and the evidence from animal behavior, the emotional system is myopic, so its utility is given by $U^e = \lambda x_1$, where $\lambda > 0$. Substituting these expressions in (18.4) we get

$$V^c = \sum_{j=1}^T \delta^{j-1} x_j + \phi(W, \sigma) \lambda x_1. \quad (18.6)$$

Maximizing (18.6) is equivalent to maximizing

$$\tilde{V}^c = x_1 + \beta \sum_{j=2}^T \delta^{j-1} x_j, \text{ where } \beta = \frac{1}{1 + \lambda \phi(W, \sigma)} < 1. \quad (18.7)$$

The preferences expressed in (18.7) are the (β, δ) preferences of *quasi-hyperbolic discounting* (see Part 3 of the book). Depletion of willpower (reduction in W) or an increase in the cognitive load (increase in σ) increases $\phi(W, \sigma)$. Equivalently, the testable prediction of the model is that these changes (reduction in W and increase in σ) decrease β , resulting in more myopic behavior. The Shiv and Fedorikhin (1999) results mentioned above are consistent with this prediction. Vohs and Heatherton (2000) find that willpower depletion among dieters leads them to buy more ice cream. Benjamin et al. (2013) give problems to Chilean high school students with short-term and long-term trade-offs. Cognitive load leads to a reduction in short-term patience. Another prediction of the model that can be derived from its interpretation in (18.5) is that an increase in $\phi(W, \sigma)$ puts greater weight on the relative cost of moving the emotional system away from its most preferred outcome. So factors that enhance the emotional desire such as proximity to stimuli are predicted to lead to greater myopia. The results of Mischel et al. (2003) mentioned above are supportive of this prediction.

BEHAVIOR UNDER RISK

Suppose that an individual faces the lottery

$$L = (x_1, p_1; x_2, p_2; \dots; x_n, p_n),$$

where $x_i \in R$ is the i th outcome, $p_i \geq 0$ the associated probability, and $p_1 + \dots + p_n = 1$. Given premise A1, we assume that the cognitive system uses expected utility theory. We assume the special case of expected value maximization, so

$$U^c = \sum_{i=1}^n x_i p_i. \quad (18.8)$$

What preferences should we assign to the emotional system? These issues were considered in more detail in Section 17.2.2 (Loewenstein, 2001). In studies that measure physiological response to unexpected events, the intensity of the shock rather than the likelihood of receiving it, is

the key. Other studies show that mental images of outcomes are the key to eliciting emotional responses, and that such responses are invariant to probabilities of outcomes (Loewenstein and O'Donoghue, 2007). Evidence that individuals often ignore information on probabilities when purchasing house insurance is also given in the work of Kunreuther et al. (1978).

Ditto et al. (2006) gave 80 undergraduates the option to undertake or avoid the following risky task: Work on a boring task in the lab for an extra 30 minutes with a probability that could be high (more risky) or low (less risky). The prize was chocolate chip cookies. Participants who could see and smell the cookies while they made their decision were less sensitive to risk information. In this case, under high risk and low risk, respectively, 80% and 85% of the players opted to undertake the task. In contrast, when the cookies were merely described (but not seen or smelt), participants were much more sensitive to the level of risk; 45% played the game under high risk and 95% under low risk. Insofar as the sight and smell of the cookies engages the emotional system, this evidence lends further support to the view that information on probabilities is relatively less salient for the emotional system.

In light of this evidence, suppose that the utility function of the emotional system is $U^e = \sum_{j=1}^n x_i/n = \bar{x}$, i.e., it cares only for the unweighted average of the outcomes, ignoring probabilities. Substituting U^c , U^e in (18.4) we get

$$V^c = \sum_{j=1}^n x_i p_i + \phi(W, \sigma) \bar{x}. \quad (18.9)$$

Maximizing (18.9) is equivalent to maximizing

$$\hat{V}^c = \sum_{j=1}^n \pi_i x_i, \text{ where } \pi_i = \frac{p_i + \phi/n}{1 + n\phi}. \quad (18.10)$$

It is evident from (18.10) that

$$p_i \geq \pi_i \text{ as } p_i \geq \frac{1}{n^2} = \hat{p}. \quad (18.11)$$

Hence, when the probability of an event is low enough ($p_i < \hat{p}$), then the decision weight is higher than the probability; the converse is true for high probability events ($p_i > \hat{p}$). This is consistent with an *inverse S-shaped* π . From Part I we know that the probability weighting function is inverse S-shaped. Thus, $\hat{V}^c$ is similar to the objective function of a prospect theory maximizer, as in the original 1979 version of prospect theory (see Part I), but not with rank dependent utility or with the cumulative probability version of prospect theory in Tversky and Kahneman (1992). Noting that $\partial \pi_i / \partial \phi > 0$ and $\phi_1 < 0$, $\phi_2 > 0$ we can predict when π becomes more or less S-shaped. If an individual's willpower is depleted or there is a greater cognitive load, then the decision weights should become more S-shaped.

A pervasive experimental finding when there are both positive and negative outcomes in a lottery, L , is *loss aversion*. Loss aversion implies that losses are more salient than equivalent gains, i.e., the absolute value of disutility from losses is greater than the utility from gains that are of the same magnitude. There is evidence that capuchin monkeys also exhibit loss aversion (Chen et al., 2006; see Part I for details). Using premises A2, A3, we postulate that the emotional system exhibits loss aversion according to the following utility function.

$$U^e = \begin{cases} x_i & \text{if } x_i \geq 0 \\ \lambda x_i & \text{if } x_i < 0 \end{cases}, \quad (18.12)$$

where $\lambda > 1$ is the parameter of loss aversion. Substituting (18.8), (18.12) in (18.4) we get that maximizing V^c is equivalent to maximizing

$$\bar{V}^c = \sum_{j=1}^n v(x_j), \quad (18.13)$$

where

$$v(x_i) = \begin{cases} x_i & \text{if } x_i \geq 0 \\ \hat{\lambda}x_i & \text{if } x_i < 0 \end{cases}, \text{ and } \hat{\lambda} = \frac{p_i + \lambda\phi}{p_i + \phi} > 1.$$

Hence, a loss averse emotional system induces loss aversion in the cognitive system. Furthermore, $\hat{\lambda} < \lambda$ so the cognitive system is less loss averse relative to the emotional system. Noting that $\partial \hat{\lambda}_i / \partial \phi > 0$ and $\phi_1 < 0, \phi_2 > 0$, willpower depletion and an increase in the cognitive load lead to an increase in the loss aversion of the cognitive system, $\hat{\lambda}$. Also, factors such as proximity of the stimuli that enhance the role of the emotional system (increase in ϕ) increase $\hat{\lambda}$. Benjamin et al. (2013) provide some evidence that is consistent with the first prediction. See Loewenstein and O'Donoghue (2007) for evidence supporting the second prediction.

SOCIAL PREFERENCES

Consider two individuals, labeled 1 and 2, who have respective material payoffs of x_1 and x_2 . There is evidence that the cognitive system takes account of the payoffs of others, consistent with some notion of an equitable or ethical payoff (Loewenstein and O'Donoghue, 2007). Thus, the utility of the cognitive system for individual $i = 1, 2$, is assumed to be given by $U_i^c = x_i + w_i^c x_j, j = 1, 2$ and $j \neq i$; $w_i^c > 0$ is the weight put by the cognitive system on the payoff of the other player, x_j .

Animals as well as humans often demonstrate empathy and sympathy towards others. Loewenstein and O'Donoghue (2007) review some relevant evidence; they write (p. 41) "rats who view a distressed fellow rat suspended in air by a harness will press a bar to lower the rat back to safe ground." In another experiment, as one rat is subjected to electric shocks, another rat, who is watching, retreats to a corner, motionless. In yet another experiment, monkeys could pull one of two chains to receive food with one of the chains delivering twice the amount but also an electric shock to another monkey in sight. Two thirds of the monkeys preferred the "less food" chain; one monkey refused to pull either chain for 5 days and another refused for 12 days.⁵ In general, it appears that the emotional system can demonstrate extreme empathy or antipathy towards others. We assume that the preferences of the emotional system of individual $i = 1, 2$, are given by $U_i^e = x_i + w_i^e x_j, j = 1, 2$ and $j \neq i$, and $w_i^e \geq 0$; a positive weight captures altruism and a negative weight captures envy.

Substituting U_i^c, U_i^e in (18.4) we get

$$V^c = x_i + w_i^c x_j + \phi(W, \sigma) [x_i + w_i^e x_j]. \quad (18.14)$$

The cognitive system chooses x_i, x_j to maximize V^c in (18.14); equivalently, it maximizes

$$V^{c*} = x_i + \omega_i x_j, \text{ where } \omega_i = \frac{(w_i^c + \phi w_i^e)}{1 + \phi}.$$

⁵ For a review of these studies, see Preston and de Waal (2002).

It is clear that

$$\omega_i \geq w_i^c \Leftrightarrow w_i^e \geq w_i^c.$$

Hence, the emotional system might induce the cognitive system to assign a lower weight to the payoff of the other player when $w_i^e < w_i^c$. Indeed, if $w_i^e < 0$ and $\frac{w_i^e}{\phi} < |w_i^c|$, then the cognitive system, which is inherently altruistic ($w_i^c > 0$), becomes envious, $\omega_i < 0$, and derives utility by lowering the payoff of the other individual. If, on the other hand, $w_i^e > w_i^c$, then the emotional system amplifies the altruism of the cognitive system. Furthermore, ω_i is increasing (decreasing) in per unit costs of willpower, ϕ , if $w_i^e - w_i^c$ is positive (negative). Hence, depending on the relative empathy triggered in the emotional system, it is possible that a decrease in willpower or an increase in cognitive load/stress (which increase ϕ) can either increase or decrease a concern for the payoff of others.

A NOTE ON WELFARE COMPARISON

Neoclassical economics typically uses revealed preference arguments to make welfare comparisons. An individual's actual choices are used to make inferences about preferences. These preferences can be used to make statements about the well-being of the individual when the environment changes. Central to this exercise is the view that individuals choose options that maximize their well-being.

Behavioral economics, on the other hand, shows that individuals might not always behave in their best interest. Should a social planner stand back and respect individual preferences? In the context of the dual process models, two main candidate functions for behavioral welfare economics are U^c and V^c (see (18.3) for the relation between the two); we can rule out U^e as it is too unstable with respect to the stimulus. However, there are interesting arguments associated with each of them.

The main argument in favor of U^c is that it captures the behavior of the cognitive system from a detached perspective. But there are two important drawbacks of using U^c . First, the set of emotional states is an argument of U^c and so imparts to it a degree of instability. Second, most people would agree that, fixing outcomes, a policy that reduced the need to exercise willpower and so reduced per unit willpower costs, ϕ , would be an improvement. This is not guaranteed by U^c . While the second of these criticisms can be avoided by considering V^c as the appropriate criterion, however, the presence of emotional states as an argument in V^c does not eliminate the first criticism. Thus, advertising by firms might appeal to the emotional system of an individual, so it also has an affect on V^c , yet the advertisements might not always promote the best-long term interests of the individual. For instance, for a long time, advertisements associated cigarette smoking with a macho image, yet the advertised product was known to the cigarette companies to be harmful to the consumers.

18.3 A dual selves model with costly commitment

One possible framework to model emotions is to consider each individual as comprising of a patient long-run (LR) self and a myopic short-run (SR) self. These are models of *dual selves*. An example is the Shefrin and Thaler (1988) model that we considered in Part 3 of the book. Here, we consider another model in this tradition, due to Fudenberg and Levine (2006).

A model that is related to the dual selves model is the model of *multiple-selves* (see Part 3 of the book). Fudenberg–Levine claim that dual selves models are better grounded in psychological evidence, and are analytically simpler. Furthermore, these models typically have a unique equilibrium, while models of multiple-selves often have multiple equilibria. We use a simple consumption–savings model to illustrate the dual selves model.

18.3.1 *The model without commitment*

Consider an infinitely lived consumer who makes a consumption–savings decision in each time period $t \in T = \{1, 2, \dots, \infty\}$. In any period, the consumer has wealth y_t , which can either be consumed, c_t , or saved, s_t . There is no other source of income. There is an exogenous and time-invariant gross return of $R = 1 + r$ on each unit of savings carried over to the next period, hence, $y_{t+1} = R s_t$. Initial wealth y_1 is given exogenously. Suppose that the marginal propensity to save is α_t in period $t \in T$ so that $s_t = \alpha_t y_t$ and $c_t = (1 - \alpha_t) y_t$. The resource constraint at time $t + 1$ is given by

$$y_{t+1} = R \alpha_t y_t, t \in T. \quad (18.15)$$

The objective of the SR self is to myopically maximize current utility. Thus, the SR self chooses α_t to maximize

$$u(y_t, \alpha_t) = \log(1 - \alpha_t) y_t = \log c_t, t \in T. \quad (18.16)$$

In the absence of any control exerted by the LR self, the optimal choice of the SR self, from (18.16), is to choose $\alpha_t = 0$, i.e., consume the entire current income.

The LR self can induce the SR self to choose some non-zero marginal propensity to save, $\alpha_t \neq 0$, by expending some self-control or willpower costs, $C(y_t, \alpha_t)$. Assuming that these costs are linear and proportional to the amount saved, we get

$$C(y_t, \alpha_t) = \gamma [\log(1 - 0) y_t - \log(1 - \alpha_t) y_t] = -\gamma \log(1 - \alpha_t), t \in T. \quad (18.17)$$

In (18.17), $\gamma > 0$ is the marginal cost of inducing the SR self to save extra. Furthermore, C is increasing and convex in α_t .

The LR self uses exponential discounting to maximize the present discounted value of lifetime consumption, subject to the resource constraint in (18.15). Hence, the utility of the LR self is given by

$$U = \sum_{t=1}^{\infty} \delta^{t-1} [u(y_t, \alpha_t) - C(y_t, \alpha_t)], \quad (18.18)$$

where $0 < \delta < 1$ is the discount factor, $u(y_t, \alpha_t)$ is defined in (18.16), and $C(y_t, \alpha_t)$ is defined in (18.17). Comparing (18.16), (18.18), the instantaneous utility functions of the LR self and the SR self are identical, except for the costs of self-control. Substitute (18.16), (18.17) in (18.18) and add and subtract $\gamma \log(1 - \alpha_t) y_t$ from the instantaneous utility function to get

$$U = \sum_{t=1}^{\infty} \delta^{t-1} [(1 + \gamma) \log(1 - \alpha_t) y_t - \gamma \log y_t]. \quad (18.19)$$

The problem of the LR self is to choose the sequence of marginal propensities to save, $\{\alpha_t\}_1^\infty$, in order to maximize (18.19) subject to the constraint (18.15). Fudenberg and Levine (2006) show that the unique solution is a constant, time-invariant, savings rate, α , that lies between zero and one.⁶ When $\alpha_t = \alpha$ for all t , then (18.15) can be solved recursively to give

$$y_t = (R\alpha)^{t-1} y_1, \quad (18.20)$$

where y_1 is the initial level of wealth. Substitute (18.20) in (18.19) we get

$$U = \frac{(1 + \gamma)}{1 - \delta} [\log(1 - \alpha) + \log y_1] + \frac{\delta}{(1 - \delta)^2} \log R\alpha. \quad (18.21)$$

Since we have an interior and unique solution, using $\frac{\partial U}{\partial \alpha} = 0$, we get the optimal value of α ,

$$\alpha = \frac{\delta}{1 + \gamma(1 - \delta)}. \quad (18.22)$$

As γ , the per unit cost of inducing the SR self to save an extra unit increases, it becomes optimal to reduce α , the marginal propensity to save. Furthermore, $\frac{\partial}{\partial \delta} \left(\frac{\partial \alpha}{\partial \gamma} \right) > 0$, thus, as the degree of patience (as reflected in the magnitude of δ) increases, the marginal impact of self-control costs on the savings rate increases.

In the special case of no self-control costs, we have $\gamma = 0$; in this case, using (18.22), the optimal savings rate, denoted by α^* , is

$$\alpha^* = \delta. \quad (18.23)$$

Hence, optimal consumption in the absence of self-control costs is

$$c_t^* = (1 - \delta) y_t, \quad t \in T. \quad (18.24)$$

Substitute (18.23) in (18.21) to get

$$U = \frac{1}{1 - \delta} [\log(1 - \delta) + \log y_1] + \frac{\delta}{(1 - \delta)^2} \log R\delta. \quad (18.25)$$

In the absence of self-control costs, (18.25) expresses lifetime utility of the LR self as a function of initial wealth, the discount rate, and the rate of return on savings.

18.3.2 *Commitment via a cash-in-advance constraint*

So far we have assumed that the LR self does not have access to a commitment technology. For that reason, the LR self can only control the SR self through the costly exercise of willpower. We now allow for the possibility that such a technology is available; essentially we wish to give the

⁶ The proof is available in the supplementary section of an earlier version of the paper by Fudenberg and Levine (2006) and applies to any CRRA utility function of which log utility is a special case.

LR self the power to remove the source of temptation for the SR self. Each period is now divided into two sequential subperiods: the *bank subperiod* and the *nightclub subperiod*.

At any time t , in the first subperiod, the bank subperiod, there is no consumption temptation. The LR self chooses, in this period, to save some amount $s_t = \alpha_t y_t$. This amount is kept in the bank and available only in the next period at time $t + 1$. The rest, $x_t = (1 - \alpha_t)y_t$, is passed on to the second subperiod in time t , as cash. In the nightclub subperiod, the SR self has access to a cash amount, x_t , and it consumes an amount, c_t , such that $0 \leq c_t \leq x_t$; consumption always takes place only in the nightclub subperiod. The discount factor between two successive nightclub subperiods is δ .

Any cash leftover from the nightclub subperiod of period t , i.e., $x_t - c_t \geq 0$, is deposited into the bank and carried over to period $t + 1$, along with the savings, s_t , made in the bank subperiod of period t . There is a time-invariant gross return R on savings, hence, the individual's wealth at the beginning of period $t + 1$ is

$$y_{t+1} = R(s_t + x_t - c_t). \quad (18.26)$$

We look at two possible subcases.

I. Savings is the only source of income

The LR self can costlessly impose any savings rate in any period by controlling the amount of cash available to the SR self to spend; this is the commitment technology. From (18.23), we know that in the absence of self-control costs, $\alpha^* = \delta$. This savings rate can be readily implemented by choosing to pass on the cash

$$x_t = (1 - \delta)y_t, \quad (18.27)$$

to the SR self in the nightclub period. Since the SR self is completely myopic, it will immediately consume all the cash received.

II. Stochastic income augments savings income

At $t = 1$, the cash available for the SR self in the nightclub period is x_1 . Now suppose that in the first, and only the first, nightclub subperiod, the SR self faces additive income uncertainty of the following form: Income is $x_1 + \tilde{z}_1$, where $\tilde{z}_1$ is a random variable with density f . The interpretation is that the income of the SR self is $x_1 + z_1$ with probability $f(z_1)$ for all z_1 in the support of f . Then the budget constraint for the SR self in the first nightclub period is given by

$$c_1 \leq x_1 + z_1. \quad (18.28)$$

We need to study the individual's optimization problem for each possible realization, z_1 . Second period wealth is given by

$$y_2 = R(s_1 + x_1 + z_1 - c_1) \equiv R(y_1 + z_1 - c_1), \quad (18.29)$$

where $s_1 + x_1 \equiv y_1$ and $y_1 + z_1$ gives the actual income in period 1. Net of first period consumption, c_1 , the savings are invested at a gross rate R to give y_2 .

The random income is available only in the first nightclub subperiod, leading to a potential violation of the cash-in-advance constraint for the SR self. Hence, the first best outcome is no longer guaranteed. The SR self in the first nightclub subperiod could have extra cash, z_1 , to play

around with, hence, the LR self, in the first period, might wish to exert costly willpower effort. The cost of this effort is as given in (18.17),

$$C = \gamma [\log(x_1 + z_1) - \log c_1].$$

The first period utility of the LR self is $\log c_1 - \gamma [\log(x_1 + z_1) - \log c_1]$, which can be written as

$$(1 + \gamma) \log c_1 - \gamma \log(x_1 + z_1). \quad (18.30)$$

From the second period onwards, however, there is no need to exert costly willpower because the SR self faces no subsequent income uncertainty and, so, the cash-in-advance constraint is never violated. Hence, the lifetime utility of the LR self starting from period 2 is found by substituting $y_2 = y_1$ in (18.25) (taking account that we are summing from period 2 onwards) to give

$$\frac{\delta}{1 - \delta} [\log(1 - \delta) + \log y_2] + \frac{\delta^2}{(1 - \delta)^2} \log R\delta. \quad (18.31)$$

Putting together (18.30), (18.31), the LR self maximizes

$$(1 + \gamma) \log c_1 - \gamma \log(x_1 + z_1) + \frac{\delta}{1 - \delta} [\log(1 - \delta) + \log y_2] + \frac{\delta^2}{(1 - \delta)^2} \log R\delta, \quad (18.32)$$

subject to (18.29). We ignore the constraint (18.28) for the time being. The first order condition with respect to c_1 gives the optimal value of first period consumption

$$c_1^* = (1 - \theta) (y_1 + z_1), \text{ where } \theta = \frac{\delta}{\delta + (1 + \gamma)(1 - \delta)} < 1. \quad (18.33)$$

From (18.23), (18.24) in the absence of self-control problems, optimal consumption is $(1 - \delta) y_1$. One gets a similar result from (18.33) when willpower costs per unit $\gamma = 0$; here one gets $c_1^* = (1 - \delta) (y_1 + z_1)$.

Whether the individual gets to consume the optimal amount c_1^* , given in (18.33), depends on whether (18.28) binds. If $c_1^* \leq x_1 + z_1$, then c_1^* is implemented. If $c_1^* > x_1 + z_1$, then the first period consumption is constrained by the upper bound $x_1 + z_1$, so all pocket cash is consumed by the SR self in the first nightclub subperiod.

Suppose $c_1^* \leq x_1 + z_1$. Since x_1 is the solution to $\gamma = 0$ we know from (18.27), that $x_1 = (1 - \delta) y_1$. Substituting this in (18.28), if

$$c_1^* \leq (1 - \delta) y_1 + z_1 \quad (18.34)$$

then the optimal consumption plan is implemented. Define by z_1^* the value of z_1 such that (18.34) just binds, i.e.,

$$\left(1 - \frac{\delta}{\delta + (1 + \gamma)(1 - \delta)}\right) (y_1 + z_1^*) = (1 - \delta) y_1 + z_1^*,$$

which gives $z_1^* = \gamma (1 - \delta) y_1$.




Figure 18.2 Willpower costs and optimal consumption.

Figure 18.2 plots the two terms on the left and the right hand sides of (18.34). c_1^* is plotted for two values of $\gamma = 0$ and $\gamma > 0$. When $\gamma = z_1 = 0$, then $c_1^*|_{\gamma=0, z_1=0} = (1 - \delta)y_1$; the slope of $c_1^*|_{\gamma=0}$ is δ . An increase in γ reduces the value of θ in (18.33), hence, increasing the intercept and slope of the curve $c_1^*|_{\gamma>0}$ relative to that of $c_1^*|_{\gamma=0}$. The intersection of the curves $c_1^*|_{\gamma>0}$ and $(1 - \delta)y_1 + z_1$ in (z_1, c_1) space corresponds to $z_1 = z_1^*$.

For all $z_1 < z_1^*$, $c_1^* > (1 - \delta)y_1 + z_1$, hence, consumption is bounded above by $(1 - \delta)y_1 + z_1$ and the SR self does not deposit any cash in the bank to be carried forward to the next period. For all $z_1 > z_1^*$, $c_1^* < (1 - \delta)y_1 + z_1$, hence, the SR self deposits surplus cash into the bank for the next period. In the absence of self-control problems ($\gamma = 0$) we have $z_1^* = 0$, hence, the SR self does not spend all the increment in income, for any level of $z_1 > 0$.

Consider the case $\gamma > 0$. It follows that in the range of small windfalls, $z_1 < z_1^*$, an increase in unanticipated income is fully consumed (because $c_1^* > (1 - \delta)y_1 + z_1$). On the other hand, in the range of large windfalls, $z_1 > z_1^*$, an increase in unanticipated income is not fully consumed (because $c_1^* < (1 - \delta)y_1 + z_1$).

A final implication is that in an appropriately modified model where the windfall gains are deposited in a bank, the marginal propensity to consume out of these gains would be lower. The reason is that the LR self cares about future time periods and, hence, would desire to smooth out consumption.

Fudenberg–Levine show that by assuming a convex cost of delay in the model above, they can explain other behavioral findings. These include the Allais paradox and the positive relation between cognitive load and risk aversion. Benjamin et al. (2013) give empirical evidence to show that as cognitive load increases, people become more risk averse. Cognitive load was induced by asking the subjects to remember a seven digit number when they play lotteries in Allais-type paradoxes. This can also provide an explanation for the Allais paradox. However, this does not explain why the Allais paradox should arise in the absence of cognitive load (see Part 1 of the book).

Fudenberg–Levine also fit their model to estimated macroeconomic parameters such as the interest rate and the coefficient of relative risk aversion. However, some of the results are quite sensitive to the parameters used. It is also not clear how this model would explain a range of

other behavioral regularities that rely on reference dependence, loss aversion, and non-linear probability weighting (see Part I of the book for details).

18.4 Information asymmetries between emotions and cognition

In this section, we consider a fairly speculative framework that models the interaction between a myopic emotional system/module and a farsighted cognitive system/module in the presence of information asymmetries between the two (Brocas and Carillo, 2008).⁷ To borrow an analogy from principal-agent models, the cognitive module is the principal and the emotional module is the agent.

Suppose that there are only two time periods, $t = 1, 2$. In each period, the individual's consumption and labor are denoted, respectively, by

$$c_t \geq 0, l_t \in [0, L], \quad (18.35)$$

where L is an upper limit on available labor. The emotional system is self-regarding and myopic, so it cares only about instantaneous utility, which takes the quasi linear form

$$U_t(c_t, l_t; \theta_t) = \theta_t u(c_t) - l_t; \theta_t > 0. \quad (18.36)$$

θ_t is a parameter of preferences that enhances the marginal utility of consumption. We assume that θ_t is *iid* over time and takes two possible values, θ_L, θ_H , with associated probabilities $p, 1 - p$, and $0 < \theta_L < \theta_H$.⁸ Furthermore, $u' > 0, u'' < 0$.

In contrast to the emotional system, the cognitive system is utilitarian and farsighted, with preferences given by $U_1 + U_2$. The cognitive system is treated as the *principal* and the emotional system in period $t = 1, 2$ is the *agent* in that period.

The individual can transform one unit of labor into one unit of income, so the wage rate is 1. Saving, s , made in period 1, earns a risk-free per unit rate of return, $r > 0$. The budget constraints in the two periods are respectively

$$c_1 + s \leq l_1 \text{ and } c_2 \leq l_2 + (1 + r)s. \quad (18.37)$$

Second period consumption can be increased not just by working hard in period 2 but also by working hard and saving more in the first period. This feature will be important in the model. Since the marginal utility of consumption is always positive, the budget constraints hold with equality. Eliminating s from the two budget constraints, we get the intertemporal budget constraint

$$c_1 + \frac{c_2}{(1 + r)} = l_1 + \frac{l_2}{(1 + r)}. \quad (18.38)$$

⁷ For a brief survey of these models that also considers a wider range of issues such as a better understanding of memory and multitasking, see Brocas and Carillo (2014).

⁸ Brocas and Carillo (2008) consider a continuous distribution of θ . However, our two types case is able to illustrate the main points.

18.4.1 The full information equilibrium

Under full information, the principal (the cognitive system) knows the values of θ_1, θ_2 . This enables it to design contracts $\{c_1(\theta_1), l_1(\theta_1)\}, \{c_2(\theta_2), l_2(\theta_2)\}$, respectively, for the two agents, namely, the emotional systems at dates $t = 1$ and $t = 2$.⁹ The principal chooses c_1, l_1, c_2, l_2 in order to maximize the following objective function

$$(\theta_1 u(c_1) - l_1) + (\theta_2 u(c_2) - l_2), \quad (18.39)$$

subject to the feasibility constraints in (18.35) and the intertemporal budget constraint in (18.38). Ignoring (18.35) for the time being, the Lagrangian is

$$Z = (\theta_1 u(c_1) - l_1) + (\theta_2 u(c_2) - l_2) + \lambda \left[l_1 + \frac{l_2}{(1+r)} - c_1 - \frac{c_2}{(1+r)} \right],$$

where λ is the Lagrangian multiplier on the intertemporal budget constraint. The first order conditions for an interior solution to c_1, c_2 are

$$\frac{\partial Z}{\partial c_1} = 0 \Leftrightarrow \theta_1 u'(c_1) = \lambda. \quad (18.40)$$

$$\frac{\partial Z}{\partial c_2} = 0 \Leftrightarrow \theta_2 u'(c_2) = \frac{\lambda}{1+r}. \quad (18.41)$$

The partial derivatives of the Lagrangian with respect to l_1, l_2 are,

$$\frac{\partial Z}{\partial l_1} = -1 + \lambda, \quad (18.42)$$

$$\frac{\partial Z}{\partial l_2} = -1 + \frac{\lambda}{1+r}. \quad (18.43)$$

There is no interior solution to the labor supply in both periods. Suppose to the contrary that $l_1, l_2 \in (0, L)$. Then (18.42), (18.43) hold with equality. But this implies that $r = 0$, a contradiction. There is a boundary solution $l_1 = L$ at which $\frac{\partial Z}{\partial l_1} > 0$. The intuition is that one unit of labor in period 1 generates one unit of income (and one unit of disutility). This can be converted to $1 + r$ units of second period consumption. Since there is no discounting between the two periods, and $r > 0$, it is optimal that the full information solution to first period labor supply attains its upper bound, $l_1^F = L$ and the solution to l_2 is an interior one, so $\lambda = 1 + r$.

Substituting λ in (18.40), (18.41) gives the solution for c_1 and c_2 . Given $l_1^F = L$, l_2 can then be recovered from (18.38). Thus, the solution under full information is given by

$$\theta_1 u'(c_1^F) = 1 + r. \quad (18.44)$$

$$\theta_2 u'(c_2^F) = 1. \quad (18.45)$$

$$l_1^F = L, l_2^F = c_1^F(1+r) + c_2^F - L(1+r). \quad (18.46)$$

⁹ This is clearly a massively reduced-form and fairly strong assumption, the microfoundations of which are not entirely established.

The conditions (18.44), (18.45) equate the marginal utility of consumption in periods 1 and 2, respectively, with the opportunity cost. Each unit of consumption in period 1 leads to $1+r$ units of forgone second period consumption, while the opportunity cost of an extra unit of consumption in period 2 is one unit of forgone leisure. From (18.46), we can check if the solution to l_2 is an interior one.

18.4.2 The asymmetric information equilibrium

Now suppose that only the emotional system knows the true value of θ_t , $t = 1, 2$. Using backward induction, we first solve the period 2 problem. In period 2, the preferences of the emotional system and the cognitive system coincide because there is no future to worry about. The second period optimization problem is to choose c_2, l_2 to maximize

$$\theta_2 u(c_2) - l_2,$$

subject to the budget constraint

$$c_2 \leq l_2 + (1+r)(l_1 - c_1),$$

where we have used $s = (l_1 - c_1)$. Since $u' > 0$, the budget constraint binds. We can substitute l_2 from the binding budget constraint into the objective function to get an unconstrained optimization problem in c_2 . The optimal solution, $c_2^*(\theta_2)$, can be found from the first order condition

$$\theta_2 u'(c_2^*) = 1, \quad (18.47)$$

i.e., marginal utility from consumption equals its opportunity cost. The optimal second period labor supply, $l_2 = l_2^*$, can now be found from the budget constraint

$$l_2^*(\theta_2) = c_2^*(\theta_2) + (l_1 - c_1)(1+r). \quad (18.48)$$

In the first period, the preferences of the two systems are in conflict. Hence, the full information contracts are not incentive compatible under asymmetric information. Under full information, $\theta_1 u'(c_1^F) = 1+r$ and $l_1(\theta_L) = l_1(\theta_H) = L$. Since $0 < \theta_L < \theta_H$ it follows that $c_1^F(\theta_L) < c_1^F(\theta_H)$. Type θ_L will then benefit by misstating type to be θ_H in order to receive extra consumption without working harder. Hence, the principal (cognitive system) needs to design incentive compatible contracts.

Denote the contracts under asymmetric information by $\{c_1(\theta_L), l_1(\theta_L)\}$ and $\{c_1(\theta_H), l_1(\theta_H)\}$. While private contracts are enforceable by law or by the oversight of third parties, what prevents the cognitive system from reneging on the promised consumption, once the contractual amount of labor has been put in? There are no clear cut answers to these difficult questions. Brocas and Carillo (2008) conjecture that other areas of the brain might be involved in such enforcement and claim supporting evidence from neuroscientific studies.

We can use the revelation principle to find the first period optimal contracts $\{c_1(\theta_L), l_1(\theta_L)\}$ and $\{c_1(\theta_H), l_1(\theta_H)\}$ as the solution to the following problem.

$$\text{Max}_{\{c_1(\theta_1), l_1(\theta_1)\}} E_{\theta_1} [\theta_1 u(c_1(\theta_1)) - l_1(\theta_1)] + E_{\theta_2} [\theta_2 u(c_2^*(\theta_2)) - l_2^*(\theta_2)]. \quad (18.49)$$

Subject to (18.35), (18.47), (18.48), and the two incentive compatibility constraints for types θ_L, θ_H , respectively,

$$IC_L : \theta_L u(c_1(\theta_L)) - l_1(\theta_L) \geq \theta_L u(c_1(\theta_H)) - l_1(\theta_H). \quad (18.50)$$

$$IC_H : \theta_H u(c_1(\theta_H)) - l_1(\theta_H) \geq \theta_H u(c_1(\theta_L)) - l_1(\theta_L). \quad (18.51)$$

The objective function in (18.49) reflects uncertainty arising from the two periods. The operators E_{θ_1} and E_{θ_2} denote, respectively, expectations with respect to θ_1 and θ_2 ; each variable takes the value θ_L with probability p and θ_H with probability $1 - p$. The IC_L constraint requires that type θ_L cannot do better by mimicking type θ_H , while the IC_H constraint rules out similar mimicking by type θ_H . A notable difference from standard mechanism design problems is the absence of participation constraints, as it is not clear what they would imply in the current context.

The most important insight of the model can be obtained directly from the two incentive compatibility constraints. As in standard mechanism design models, IC_H binds.¹⁰ From the binding IC_H constraint we get

$$l_1(\theta_H) = \theta_H [u(c_1(\theta_H)) - u(c_1(\theta_L))] + l_1(\theta_L). \quad (18.52)$$

Substituting (18.52) in IC_L we get that $u(c_1(\theta_H)) \geq u(c_1(\theta_L))$, or $c_1(\theta_H) \geq c_1(\theta_L)$, so the level of consumption is monotonic in the type. From IC_L , we also get that

$$\theta_L [u(c_1(\theta_H)) - u(c_1(\theta_L))] + [l_1(\theta_L) - l_1(\theta_H)] \leq 0. \quad (18.53)$$

Since $u(c_1(\theta_H)) \geq u(c_1(\theta_L))$, a necessary condition for (18.53) to hold is that

$$l_1(\theta_L) \leq l_1(\theta_H). \quad (18.54)$$

This brings us to the crux of the solution. Type θ_L would like to claim to be type θ_H if it cared about consumption alone (because $c_1(\theta_H) \geq c_1(\theta_L)$). Anticipating this incentive, the principal cleverly designs the contracts in such a way that (18.54) holds, so if type θ_L wants to consume more, then it must also work more. Since more work entails greater disutility, if $l_1(\theta_H)$ is sufficiently high, then type θ_L no longer has an incentive to misrepresent its type. The authors interpret this finding as explaining behaviors such as: “I will spend the weekend in Palm Springs only if I finish the paper by Friday” or “If I decide to watch the soccer game today, then I must first complete the referee report today.”

Consider the general case when θ_1 and θ_2 are distributed continuously over the interval $[\underline{\theta}, \bar{\theta}]$. The density of θ_1 is $f(\theta_1)$ and the distribution function is $F(\theta_1)$. In this case we get the following result.

Proposition 18.18 (*Brocas and Carillo, 2008*): *Under asymmetric information and a continuous distribution of θ_1, θ_2 the optimal contracts are as follows.*

- (i) *At time $t = 1$: There exists a cutoff point $\theta_c < \bar{\theta}$ such that the optimal contracts for all types $\underline{\theta} \leq \theta_1 < \theta_c$, denoted by $\{c_1^*(\theta_1), l_1^*(\theta_1)\}_{\underline{\theta}}^{\theta_c}$, are given by*

¹⁰ Brocas and Carillo (2008) demonstrate this property when θ has a continuous distribution.

$$\left[\theta_1 + r \left(\theta_1 + \frac{F(\theta_1)}{f(\theta_1)} \right) \right] u'(c_1^*(\theta_1)) = 1 + r, \quad (18.55)$$

while the optimal amount of labor elicited in the first period is

$$l_1^*(\theta_1) = L - \left[\bar{\theta} u(c_1^*(\bar{\theta})) - \theta_1 u(c_1^*(\theta_1)) - \int_{\theta_1}^{\bar{\theta}} u(c_1^*(x)) dx \right] < L. \quad (18.56)$$

For all $\theta_c \leq \theta_1 \leq \bar{\theta}$ there is pooling of contracts. Each of these types is given the contract $\{c_1^*(\theta_c), L\}$, where $c_1^*(\theta_c)$ is the optimal consumption given in (18.55), and evaluated at $\theta_1 = \theta_c$.

- (ii) At time $t = 2$: The principal allows the time 2 agent to choose any consumption–labour pair that satisfies the intertemporal constraint (18.38). The agent then, voluntarily, chooses

$$u'(c_2^*(\theta_2)) = \frac{1}{\theta_2} \text{ and } l_2^*(\theta_2) = (1 + r)(c_1^*(\theta_1) - l_1^*(\theta_1)) + c_2^*(\theta_2); \theta_2 \in [\underline{\theta}, \bar{\theta}].$$

Since $\theta_1 + r \left(\theta_1 + \frac{F(\theta_1)}{f(\theta_1)} \right) > \theta_1$, it follows by comparing (18.44) with (18.55) that $c_1^*(\theta_1) > c_1^F(\theta_1)$, hence, consumption in the first period under asymmetric information is larger than that under full information. This reflects the information rents derived by the better informed agent in period 1. From (18.46), (18.56), under asymmetric information, the first period emotional system is asked for a lower amount of labor relative to that under full information. Hence, relative to full information, first period consumption is larger and first period labor supply is smaller. Thus, it is as if an endogenous rate of time preference develops under asymmetric information despite the fact that the cognitive self undertakes no time discounting.

18.5 Strategic ignorance, confidence, and self-esteem

In this section we consider models of multiple-selves in which individuals may have asymmetric information about some aspects of their own characteristics; for instance, about their willpower, ability, or motivations. Individuals might engage in introspection to discover themselves, or use their performance in a task to discover their underlying motivations or abilities. In this scenario, individuals might decide to strategically transmit or hide information to their own future selves, or even manage their degree of awareness of personal qualities. Among the set of phenomena explained by this model are the roles of strategic ignorance, selective memory transmission, and self-handicapping.

18.5.1 Strategic ignorance

Consider a three period version of the multiple-selves model in Carrillo and Mariotti (2000), where time $t = 1, 2, 3$; the self in period t is denoted by self t . Each self, in each period, is risk neutral and makes a simple binary consumption choice $x_t \in \{0, 1\}$; 0 denotes abstinence and 1 denotes consumption of a free good (introducing a cost of acquiring the good does not alter the qualitative predictions). To simplify further, we assume that self 3 chooses $x_3 = 0$. Current consumption increases instantaneous utility of self t but it reduces the utility of self $t + 1$ by $e > 1$

units (negative externality) with some probability $\theta \geq 0$; e.g., the good could be some unhealthy food/lifestyle.¹¹ θ is distributed uniformly over $[0, 1]$ with cumulative density F ; all selves are aware of the relevant distribution. In addition to the consumption decision, self 1 has to decide whether to costlessly learn the true value of θ .

The instantaneous utility of self t is given by

$$u_t = x_t - y_t, \quad (18.57)$$

where y_t is the externality caused by the consumption decision of self $t - 1$, so

$$y_t = \begin{cases} 0 & \text{if } x_{t-1} = 0 \\ \theta e & \text{if } x_{t-1} = 1 \end{cases}, \quad (18.58)$$

with the assumption that $x_3 = 0$. Preferences take the special *quasi-hyperbolic* $(\beta, 1)$ form, $0 < \beta < 1$ (see Part 3 of the book), so the utility of self t is

$$U_t = \begin{cases} u_1 + \beta(u_2 + u_3) & \text{if } t = 1 \\ u_2 + \beta u_3 & \text{if } t = 2 \\ u_3 & \text{if } t = 3 \end{cases}. \quad (18.59)$$

From (18.57), (18.58), (18.59) and the assumption that $x_3 = 0$, we can recursively solve backwards to work out the intertemporal utility of the successive selves.

$$U_3 = -\theta e x_2; \quad x_2 \in \{0, 1\}. \quad (18.60)$$

$$U_2 = (x_2 - \theta e x_1) + \beta(-\theta e x_2) = x_2(1 - \beta\theta e) - \theta e x_1; \quad x_1, x_2 \in \{0, 1\}. \quad (18.61)$$

$$U_1 = x_1 + \beta(x_2 - \theta e x_1) + \beta(-\theta e x_2) = x_1(1 - \beta\theta e) + x_2\beta(1 - \theta e); \quad x_1, x_2 \in \{0, 1\}. \quad (18.62)$$

Assume that $\frac{1}{\beta e} < 1$. We also make the tie-breaking assumption that when self t is indifferent between consuming and not consuming, he chooses to consume.

Using (18.62), the following three cases arise for self 1.

1. $0 \leq \theta \leq 1/e$: If $0 \leq \theta \leq 1/e$, then $1 - \beta\theta e > 0$ and $1 - \theta e \geq 0$. Thus, self 1 would like to consume in period 1 ($x_1 = 1$) and would also like self 2 to consume ($x_2 = 1$).
2. $\frac{1}{e} < \theta \leq \frac{1}{\beta e}$: If $\frac{1}{e} < \theta \leq \frac{1}{\beta e}$, then $1 - \beta\theta e \geq 0$ and $1 - \theta e < 0$. Thus, self 1 prefers to consume ($x_1 = 1$), and prefers self 2 to abstain from consumption ($x_2 = 0$).
3. $\frac{1}{\beta e} < \theta \leq 1$: In this case, self 1 prefers to abstain in both periods ($x_1 = x_2 = 0$).

With the present bias implied by $(\beta, 1)$ preferences, the preferences of self 2 are different from the preferences of self 1. Furthermore, self 1 cannot bind the choices of self 2.

4. From (18.61), it is apparent that self 2 will consume ($x_2 = 1$) if $1 - \beta\theta e \geq 0$, or $0 < \theta \leq \frac{1}{\beta e}$. Otherwise, if $\frac{1}{\beta e} < \theta \leq 1$, then self 2 abstains, $x_2 = 0$.

For self 1, if θ is high enough, so that x_2 causes a large enough negative externality in period 3, then self 1 prefers $x_2 = 0$ (as in case 2 above). However, self 2 is relatively biased towards x_2 on

¹¹ Carrillo and Mariotti (2008) also consider the more general case when current consumption creates a negative externality for all future selves.

account of his current bias and so always prefers $x_2 = 1$ unless $\frac{1}{\beta e} < \theta \leq 1$. We summarize these results in the following proposition.

Proposition 18.19 *Consider the split of the domain of θ into three regions.*

- (i) *If $\theta \in [0, \frac{1}{\beta e}]$, then there is no conflict between the choices of selves 1 and 2.*
- (ii) *If $\frac{1}{\beta e} < \theta \leq \frac{1}{\beta e}$, then there is strategic conflict between the two selves; self 1 prefers $x_2 = 0$, while self 2 prefers $x_2 = 1$.*
- (iii) *If $\frac{1}{\beta e} < \theta \leq 1$, then selves 1 and 2 prefer $x_2 = 0$.*

Let us now assume that self 1 can learn the true value of θ . As soon as the value of θ is learnt, self 2 is also privy to it; if the value is not learnt then self 2 can, at best, rely on the knowledge of the underlying distribution, F . We assume that memory cannot be unlearned, although we relax this assumption in Section 18.5.2 below.

Self 1 then knows that if $\theta \leq \frac{1}{\beta e}$, then there is no conflict with self 2 and the outcome is $x_1 = x_2 = 1$. On the other hand, if self 1 learns that $\frac{1}{\beta e} < \theta$, then there is no conflict either, and the outcome is $x_1 = x_2 = 0$. Suppose that self 1 decides not to learn the true value of θ . Since all selves are risk neutral and utility is linear in θ , the following result directly follows from our discussion.

Proposition 18.20 *(No learning and risk neutrality): If $E[\theta] \in [0, \frac{1}{\beta e}]$, then $x_1 = x_2 = 1$. However, if $\frac{1}{\beta e} < E[\theta]$, then $x_1 = x_2 = 0$.*

Let us now compute the utility of self 1 from the learning and the no learning options.

1. *Learning:* Denote the expected utility of self 1 from learning the true value of θ by U_1^L . Note that U_1^L is the utility under the learning option, but before the value of θ has been revealed. Then,

$$U_1^L = F\left(\frac{1}{\beta e}\right) \left[1 + \beta - 2\beta e E\left[\theta \mid \theta \leq \frac{1}{\beta e}\right] \right], \quad (18.63)$$

where $F\left(\frac{1}{\beta e}\right)$ is the cumulative probability that θ is at most equal to $\frac{1}{\beta e}$. In this case, selves 1 and 2 prefer to consume in their respective periods. The term in the square braces follows from risk neutrality, when (18.62) is evaluated at $x_1 = x_2 = 1$. With the complementary probability $1 - F\left(\frac{1}{\beta e}\right)$, none of the selves wants to consume and there is no externality caused to any self.

2. *No learning:* Denote the expected utility of self 1 if she chooses not to learn the true value of θ by U_1^N ; in this case self 2 only knows the distribution F of θ . Then, using Proposition 18.20 we get

$$U_1^N = \begin{cases} 1 + \beta - 2\beta E[\theta] e & \text{if } E[\theta] \leq \frac{1}{\beta e} \\ 0 & \text{if } E[\theta] > \frac{1}{\beta e} \end{cases}. \quad (18.64)$$

We get the following main result.

Proposition 18.21

- (a) *If $E[\theta] \leq \frac{1}{\beta e}$, then $U_1^N < U_1^L$, so the individual prefers to be strategically informed about θ .*

(b) If $E[\theta] > \frac{1}{\beta e}$, then $U_1^N > U_1^L$ if $E\left[\theta \mid \theta \leq \frac{1}{\beta e}\right] > \frac{1+\beta}{2\beta e}$.

Proof: Since θ is distributed uniformly over $[0, 1]$, it follows that

$$E[\theta] = \frac{1}{2}, E\left[\theta \mid \theta \leq \frac{1}{\beta e}\right] = \frac{1}{2\beta e}, F\left(\frac{1}{\beta e}\right) = \frac{1}{\beta e}. \quad (18.65)$$

(a) Suppose that $E[\theta] \leq \frac{1}{\beta e}$. Substituting (18.65) in (18.63), (18.64) it follows that

$$U_1^N < U_1^L \Leftrightarrow 0 < (\beta e - 1) \left(1 - \frac{1}{e}\right),$$

which is true because of the assumptions (i) $\frac{1}{\beta e} < 1$ so $\beta e > 1$, and (ii) $e > 1$.

(b) Suppose that $E[\theta] > \frac{1}{\beta e}$. In this case $U_1^N = 0$, so the condition $U_1^N > U_1^L$ reduces to $U_1^L < 0$. The stated result now follows directly from (18.63). ■

The intuition for Proposition 18.21 arises from the interval $\theta \in \left(\frac{1}{e}, \frac{1}{\beta e}\right]$ in which self 1 prefers $x_2 = 0$ but self 2 prefers $x_2 = 1$. Self 1 prefers to be strategically uninformed if this leads to self 2 abstaining from consumption. This occurs when $E[\theta] > \frac{1}{\beta e}$. However, this is only a necessary condition because ignorance chosen by self 1 also entails costs. For instance, it might be that $\theta > \frac{1}{\beta e}$, in which case selves 1 and 2 both prefer $x_2 = 0$. The condition $E\left[\theta \mid \theta \leq \frac{1}{\beta e}\right] > \frac{1+\beta}{2\beta e}$ in Proposition 18.21 ensures that conditional on $E[\theta] > \frac{1}{\beta e}$, θ is closer to $\frac{1}{\beta e}$ rather than to 0. The finding by Lundborg and Lindgren (2004) that individuals with higher perceived risks are less likely to be smokers offers some partial support for this theory.

18.5.2 Confidence, motivation, and self-esteem

Bénabou and Tirole (2002, 2003) consider situations where the subsequent selves of an individual might have some asymmetric information (*imperfect self-knowledge*) about the current self; we distinguish between *good* and *bad* types. A good type has *high levels of willpower*. Depending on the context, a good type might keep to deadlines, or see tasks through to completion once they have been initiated, or undertake a high level of productive effort. In contrast, a *bad type* has *inadequate willpower*. The current self might wish to undertake costly actions to convince subsequent selves of being a good type. However, the actions by a current self to convince subsequent selves of his type must be incentive compatible.

There are at least three main building blocks of the work by Bénabou and Tirole in this regard. The first two, *imperfect self-knowledge* and *imperfect willpower*, have been mentioned above in the other models in this chapter. The third building block is *imperfect recall and motivated cognition*. Imperfect recall implies, roughly, that an individual might manage, wilfully or otherwise, how past actions are remembered. The thrust of the argument is not that memory can simply be wilfully suppressed, but rather that the awareness of events in the past can be manipulated.

Consider for instance the following examples. (i) Our memory of past actions can often be self-serving. We tend to remember our successes more than our failures. We might, wrongly, ascribe our lack of success to external factors, or blame others for it. (ii) The intensity of past visceral states in our memories might diminish gradually with time. (iii) By taking down written records of an event, sharing them with friends and rehearsing the records, offers at least partial protection

from memory loss. But in many cases, we choose not to. (iv) We can control the vividness of the memory of an event by choosing the degree of attention we give to it. We sometimes pay insufficient attention to events that are unpleasant. At the extreme, we can sometimes get drunk to try to forget unpleasant events. Memory also decays naturally with time, especially if not consciously rehearsed; the latter is of course an endogenous decision.

We operationalize this as follows. At any instant of time, the current self chooses some action and the probability with which the subsequent selves can recall this action. It is possible that the current self chooses to engage in wilful self-deception. Subsequent selves are aware of the incentive for self-deception, hence, good news might get discounted in equilibrium, especially if it is likely to arise from too much self-deception. However, some self-deception might actually help. For instance, if the decision of the future selves to undertake an activity depends on the belief in their abilities, or in their “self-confidence,” then the current self might undertake actions that strengthen this belief.

There are at least three main reasons why individuals might wish to have a more optimistic and confident view of themselves, relative to a more realistic one. First, confidence in one’s abilities might directly enhance utility. Second, it might have a signaling value; self-belief might be a signal to others of one’s abilities. Third, it might have a motivational value for one’s future selves.

We now examine the motivational value of self-confidence in the context of the three period model of Bénabou and Tirole (2002). Consider a risk-neutral individual with an economic life-span of three periods $t = 1, 2, 3$. There are multiple-selves, so there is a distinct self of the individual in each period. Each self uses hyperbolic discounting and has $(\beta, 1)$ preferences, $0 < \beta < 1$. The decisions faced by the different selves are as follows.

1. Self 1 chooses an action that affects her flow payoff u_1 and the date 2 information structure. Denote by $\theta \in [0, 1]$, a random variable, such as self-confidence or ability, over which the beliefs of self 1 are given by the distribution function $F_1(\theta)$. Self 1 can alter the date 2 information structure by engaging in costly self-deception that lowers current payoff u_1 , but generates a more optimistic belief in abilities for subsequent selves, in particular, for self 2.
2. Self 2 can undertake some action or project that has current costs $c \in [0, 1]$ in terms of effort or resources at date 2. The project is successful with probability θ . If successful, the project yields a payoff of 1 unit that is obtained at date 3. The beliefs of self 2 are given by the distribution function $F_2(\theta)$.

Define the mean of θ with respect to the distribution function $F_i(\theta)$ by $\bar{\theta}_i = \int_0^1 \theta dF_i(\theta)$. Individuals are risk neutral and payoff functions are linear in θ , so the mean, $\bar{\theta}_i$, encompasses everything that the individual needs to know about the randomness in the model. In other words, $\bar{\theta}_i$ is a sufficient statistic for $F_i(\theta)$.

The expected payoff of self 2, $E_2[U]$, is given by

$$E_2[U] = \begin{cases} -c + \beta E_2\theta = -c + \beta \bar{\theta}_2 & \text{if activity is undertaken} \\ 0 & \text{if activity is not undertaken} \end{cases}, \beta \in (0, 1), \quad (18.66)$$

where E_2 is the expectation operator with respect to the distribution $F_2(\theta)$. (18.66) shows that the cost of the activity is borne in the current period but the reward is received in the next period. Self 2 undertakes the activity if $-c + \beta \bar{\theta}_2 > 0$, or¹²

¹² We do not consider the case when a self is indifferent between two choices but this can be readily incorporated by assuming some tie breaking rule.

$$\bar{\theta}_2 > \frac{c}{\beta} > c. \quad (18.67)$$

If self 2 knew the true value of θ , then she would undertake the action only if

$$\theta > \frac{c}{\beta} > c. \quad (18.68)$$

The expected payoff of self 1, $E_1 [U]$, depends on the choice made by self 2 and is given by

$$E_1 [U] = \begin{cases} u_1 + \beta(-c + \bar{\theta}_2) & \text{if self 2 undertakes the activity} \\ u_1 & \text{if self 2 does not undertake the activity} \end{cases}. \quad (18.69)$$

Self 1 prefers self 2 to undertake the activity if $u_1 + \beta(-c + \bar{\theta}_2) > u_1$, or

$$\bar{\theta}_2 > c. \quad (18.70)$$

If self 1 knew θ , then she would want the action to be undertaken if $\theta > c$. Comparing (18.67), (18.70) the crucial interval of disagreement between selves 1 and 2 is $c < \bar{\theta}_2 \leq \frac{c}{\beta}$. We summarize the main result from the discussion so far.

Lemma 18.2 *If θ is unknown to both selves, then the following hold. (1) If $c < \bar{\theta}_2 \leq \frac{c}{\beta}$, self 2 will not undertake the activity, but self 1 would like the activity to be undertaken. (2) If $\bar{\theta}_2 < c$, then none of the selves want the activity to be undertaken. (3) If $\bar{\theta}_2 > \frac{c}{\beta}$, then both selves would like the activity to be undertaken.*

MAINTAINING THE OPTIMAL DEGREE OF CONFIDENCE

Suppose that self 1 were to learn about the true value of θ . Should self 1 choose to transmit this information to self 2? The answer depends on the value of θ as well as on the credibility of information transmission. Suppose that information is completely credibly transmitted, hence, immediately believed by self 2. Thus, if self 1 chooses to reveal information, he cannot lie about the state, θ . We distinguish between two main cases.

Case I ($\bar{\theta}_2 > \frac{c}{\beta}$): From (18.67), self 2 will undertake the activity in the absence of any information from self 1. Using (18.69), the expected utility of self 1 from not passing on any information to self 2 is given by

$$E_1 [U] = u_1 + \beta(-c + \bar{\theta}_1), \quad (18.71)$$

where we have used the fact that in the absence of any transmitted information, $\bar{\theta}_1 = \bar{\theta}_2$.

On the other hand, if information is passed on to self 2, then, using (18.68), self 2 will only take the action when $\theta > \frac{c}{\beta}$. Hence, from an ex-ante perspective, before self 1 makes the decision to know θ , the expected utility of self 1 from the decision to pass on the information is

$$E_1 [U] = u_1 + \beta \int_{c/\beta}^1 (-c + \theta) dF_1(\theta). \quad (18.72)$$

Subtracting (18.71) from (18.72), the *information disclosure value* to self 1, denoted by D , can be written as

$$D = \beta \left[\int_0^c (c - \theta) dF_1(\theta) - \int_c^{c/\beta} (\theta - c) dF_1(\theta) \right] = \beta [G - L], \quad (18.73)$$

where

$$G = \int_0^c (c - \theta) dF_1(\theta) \text{ and } L = \int_c^{c/\beta} (\theta - c) dF_1(\theta)$$

denote, respectively, the gain and loss to self 1 from information disclosure; we explain in more detail below.

1. Suppose that self 1 learns $\theta < c$. Because $\bar{\theta}_2 > \frac{c}{\beta}$, self 2 will undertake the action in the absence of the information (overconfidence), but not if the information were to be revealed. Both selves do not wish the project to be undertaken if $\theta < c$, so self 1 reveals the information. By doing so, self 1 covers for the *risk of overconfidence* on the part of self 2. This corresponds to the term G in (18.73).
2. Suppose that self 1 learns $c < \theta < \frac{c}{\beta}$. In this case, self 1 would like the action to be undertaken, but because $\bar{\theta}_2 > \frac{c}{\beta}$, self 2 will undertake the action anyway, even in the absence of the information. Providing self 2 more information, i.e., that $c < \theta < \frac{c}{\beta}$ would prevent it from undertaking the action. Hence, there is a loss to self 1 from having an informed self 2, thus, self 1 prefers to *maintain the confidence* of self 2. This corresponds to the term L in (18.73).
3. Finally, if $\theta > \frac{c}{\beta}$, then self 1 is indifferent to revealing or not revealing the information because self 1 wishes the project to be undertaken and self 2 would undertake it anyway because $\bar{\theta}_2 > \frac{c}{\beta}$.

If the “risk of overconfidence motive” is stronger, then self 1 transmits information. If, however, the “confidence maintenance motive” is relatively stronger, then self 1 does not transmit information. Note that when $\beta = 1$, so the individual uses exponential discounting, then $D = G > 0$, so self 1 always transmits information; in this case the objectives of selves 1 and 2 are fully aligned. As β decreases and the present bias of self 2 increases, L increases and D falls. Thus, if the present bias of self 2 is higher, then the confidence maintenance motive dominates and so information is less likely to be transmitted.

Case II ($\bar{\theta}_2 < \frac{c}{\beta}$): Using (18.67), self 2 will not undertake the activity in the absence of additional information. In this case, self 2 will undertake the activity only if self 1 reveals $\theta > \frac{c}{\beta}$. Hence, the expected gain to self 1 from providing information to self 2 is given by

$$D = \beta \int_{c/\beta}^1 (\theta - c) dF_1(\theta) > 0.$$

Since $D > 0$, so information is always transmitted. In this case, the individual always looks for feedback on his ability.

DEGREES OF SELF-CONFIDENCE

Suppose that we have two individuals with different degrees of initial self-confidence. Which one is more eager to transmit information to its future selves? Consider two individuals whose initial

beliefs about the distribution of θ are given, respectively, by the distribution functions $F(\theta)$ and $G(\theta)$; the associated density functions are $f(\theta)$ and $g(\theta)$. The domain of θ is $[0, 1]$.

Definition 18.6 (*Confidence*): The individual with the distribution function $F(\theta)$ is said to be relatively more confident as compared to the individual with the distribution function $G(\theta)$ if the monotone likelihood ratio property is satisfied

$$\frac{f(\theta_L)}{g(\theta_L)} < \frac{f(\theta_H)}{g(\theta_H)} \quad \forall \theta_L < \theta_H \text{ in } [0, 1]. \quad (18.74)$$

Definition 18.6 requires that the more confident individual places a relatively higher probability on the higher value of θ . A straightforward implication of (18.74) is

$$F(\theta) < G(\theta) \quad \forall \theta \in [0, 1], \quad (18.75)$$

i.e., the distribution function of a more confident individual, first order stochastically dominates the distribution function of a less confident individual. Let $\bar{\theta}_1^F, \bar{\theta}_1^G$ denote, respectively, the expected values taken with respect to F and G . We assume that

$$\bar{\theta}_1^F > \bar{\theta}_1^G > c/\beta. \quad (18.76)$$

From (18.76), in the absence of additional information (so $\bar{\theta}_1^F = \bar{\theta}_2^F$ and $\bar{\theta}_1^G = \bar{\theta}_2^G$), self 2 will, undertake the activity, irrespective of the degree of initial confidence.

From (18.73), $D \geq 0$ if and only if

$$\int_0^c (c - \theta) dF_1(\theta) \geq \int_c^{c/\beta} (\theta - c) dF_1(\theta) \Leftrightarrow \int_0^{c/\beta} (\theta - c) dF_1(\theta) \leq 0.$$

Integrating by parts we get

$$|(\theta - c)F_1(\theta)|_0^{c/\beta} - \int_0^{c/\beta} dF_1(\theta) \leq 0,$$

or

$$\int_0^{c/\beta} \frac{F_1(\theta)}{F_1(c/\beta)} d\theta \geq \left(\frac{1 - \beta}{\beta} \right) c. \quad (18.77)$$

A similar expression applies for the individual who holds beliefs given by the distribution function G_1 . Using (18.74), we have

$$\frac{F_1(\theta)}{F_1(c/\beta)} < \frac{G_1(\theta)}{G_1(c/\beta)} \quad \forall \theta \leq c/\beta. \quad (18.78)$$

From (18.77), (18.78) it follows that (18.77) is more likely to hold for the less confident individual who has the initial distribution function G_1 . Hence, and equivalently, D is more likely to be positive for the less confident individual, i.e., she is more likely to transmit information. A more confident individual has more to lose by transmitting information.

SELF-HANDICAPPING AND SELF-ESTEEM

A puzzling aspect of behavior is that individuals sometimes take actions that are not conducive to good performance. Examples include, anxiety before an exam, withholding of effort, inadequate preparation, and setting of tasks for oneself where one is almost sure to fail. Such self-handicapping has been suggested by psychologists as a deliberate strategy to maintain self-esteem.

Suppose that $\bar{\theta}_2 > \frac{c}{\beta}$ so that in the absence of any information, self 2 will undertake the action. Suppose also that the efficient action is for self 1 to transmit information to self 2. Not transmitting any information (self-handicapping) is the alternative strategy for self 1. Suppose that the cost of such self-handicapping is zero. From (18.77), self-handicapping is the better alternative if

$$\int_0^{c/\beta} \frac{F_1(\theta)}{F_1(c/\beta)} d\theta < \left(\frac{1-\beta}{\beta} \right) c. \quad (18.79)$$

Using (18.78), it is more likely that (18.79) will hold for the more confident individual (who uses F_1 rather than G_1). Thus, more confident individuals are more likely to self-handicap; because they are more confident they have more to lose from learning the true value of their ability. This conclusion is not robust, however, to inclusion of costs of self-handicapping. More confident individuals are also less likely to receive bad signals, increasing the returns from self-handicapping; for a fuller analysis along these lines, see Bénabou and Tirole (2002).

AWARENESS MANAGEMENT

Suppose, for reasons mentioned above, one desires to have a positive self-assessment. How can this be achieved? One possibility is that we may be hardwired, on account of evolutionary reasons, to have positive self-assessment. However, many individuals suffer from a poor self-image, while others have to work for it, hence, this explanation is not entirely satisfactory. We focus instead on a simple strategic game between successive selves that stresses the roles of *motivation* and *cognition*.

Self 1 can, to an extent, affect the probability that self 2 remembers a piece of information. Hence, the *motivation* of self 1 might be to create greater memory awareness about signals that are more consistent with the long-term goals. In the absence of this greater awareness, the future selves who have “current biased preferences” may jeopardize these long-term goals. However, each self is rational and aware of the possibilities of self-deception by other selves. Thus, self 2 is *cognizant* of the incentives of self 1 to promote recall of events conducive to long-term goals.

Let $\lambda \in [0, 1]$ be the probability that an event, which occurs at time $t = 1$ can be remembered by self 2 at time $t = 2$. We look only at the case where the manipulation of memory by self 1 is costless.¹³ Evidence suggests that only a part of the total stock of information in one’s memory can be retrieved for conscious processing at the time of decision making. There are several possible mechanisms that self 1 can make use of. Rehearsed information is more readily available, as is information that one has direct behavioral experience of. Self 1 can also avoid situations that are directly reminiscent of unpleasant and forgettable events; for instance, by tearing up pictures of a cheating spouse, or in extreme cases taking recourse to drugs and alcohol. The information content of unpleasant news can also be questioned; for instance, on getting a low grade, a student could question if the marker has adequately understood her work. Hence, we are interested in the management of awareness of memories.

¹³ See Bénabou and Tirole (2002) for the case where such manipulation is costly.

Self 2 is aware of the incentives of self 1 to manipulate memory. Such skepticism on the part of self 2 is reflected in the use of Bayes' rule to update information.¹⁴ The technology of self-deception is as follows.

At $t = 1$, self 1 receives a binary signal σ about her ability, θ . With probability q the signal is uninformative, i.e., $\sigma = \phi$ and with probability $1 - q$, the signal is "bad news," i.e., $\sigma = L$. Hence, "good news" is defined as no news. Define

$$\theta_L \equiv E[\theta \mid \sigma = L] < E[\theta \mid \sigma = \phi] \equiv \theta_H. \quad (18.80)$$

Denote the signal recollected by self 1 as $\hat{\sigma} = \{\phi, L\}$. Memories cannot be manufactured, so $\sigma = \phi \Rightarrow \hat{\sigma} = \phi$ (if the signal is uninformative, then the recollected signal is also uninformative). However, if $\sigma = L$, then this signal is remembered only with an endogenous probability λ ,

$$\lambda = \Pr[\hat{\sigma} = L \mid \sigma = L]. \quad (18.81)$$

Hence, $1 - \lambda$ is the probability with which bad news is hidden from self 2.

Self 2 solves a Bayesian inference problem on observing the signal sent by self 1. She assesses the *reliability* of a *no news* message, r , by Bayes' rule:

$$r \equiv \Pr[\hat{\sigma} = \phi \mid \sigma = \phi] = \frac{q}{q + (1 - q)(1 - \lambda)}. \quad (18.82)$$

The numerator of (18.82) is the prior probability of the signal ϕ and the denominator is the total probability with which the signal ϕ arises. If $\lambda = 1$, then a no news signal is completely reliable, i.e., $r = 1$. The individual's *degree of self-confidence*, $\theta(r)$, is defined as

$$\theta(r) = r\theta_H + (1 - r)\theta_L. \quad (18.83)$$

So far, we have assumed that the cost of effort c of self 1, is deterministic. We now relax this assumption; let the cost of the project, c , be a random variable drawn from $[\underline{c}, \bar{c}]$ with distribution $H(c)$ and density $h(c) > 0$. We assume that

$$\frac{\underline{c}}{\beta} < \theta_L, \quad \theta_H < \frac{\bar{c}}{\beta}. \quad (18.84)$$

Analogous to (18.67), self 2 undertakes the activity if

$$E[\theta \mid \hat{\sigma}] > \frac{c}{\beta} \Leftrightarrow c < \beta E[\theta \mid \hat{\sigma}]. \quad (18.85)$$

Hence, the expected payoff of self 1 is

$$E_1[U] = u_1 + \beta \int_0^{\beta E[\theta \mid \hat{\sigma}]} (-c + E[\theta \mid \sigma]) dH(c). \quad (18.86)$$

¹⁴ Evidence is not consistent with individuals using Bayes' rule. However, here this is not important for the qualitative results, as we show below.

Suppose that self 2 was naive, and trusted the signal of self 1 without engaging in rational Bayesian updating. What expectation $E[\theta | \hat{\sigma}]$ would self 1 like self 2 to have? Analogous to (18.70), self 1 would like self 2 to undertake the activity when

$$E[\theta | \sigma] > c. \quad (18.87)$$

By manipulating the naive expectations of self 2 to be $E[\theta | \hat{\sigma}] = E[\theta | \sigma] / \beta$, the upper limit of integration in (18.86) becomes $E[\theta | \sigma]$. Thus, the expected payoff of self 1 is maximized (it is as if self 1 can induce self 2 to only undertake the action when it is in the former's self-interest). However, if self 2 is not naive this will not work.

Denote the expected utility of self 1 when the adverse information (signal L) is *concealed* by $U_C(\theta_L)$ and when it is *revealed* by $U_R(\theta_L)$. We are interested in the perfect Bayesian equilibria (PBE) of a signaling game between selves 1 and 2 in which they respectively choose λ and r (the optimal values are denoted by λ^*, r^*).

We begin by giving some intuition about the decisions of the two selves. If the bad news is hidden from self 2, it raises self-confidence from θ_L to $\theta(r^*)$ (recall that from (18.80), these are expected values). From (18.85), this additional self-confidence induces self 2 to work in the additional states of the world (a state corresponds to a cost, c) where

$$\beta\theta_L < c < \beta\theta(r^*). \quad (18.88)$$

However, the enhancement of self-confidence in this manner has a cost and a benefit. Analogous to (18.73), the net gain/loss from concealing information about the bad state is given by

$$U_C(\theta_L) - U_R(\theta_L) = \beta \left[\int_{\beta\theta_L}^{\theta_L} (\theta_L - c) dH(c) - \int_{\theta_L}^{\beta\theta(r^*)} (c - \theta_L) dH(c) \right]. \quad (18.89)$$

The first term in (18.89) is the *gain from confidence building* because over the range $\beta\theta_L < c < \theta_L$ self 1 would like self 2 to undertake effort but the latter does not wish to. The second term is the *loss from overconfidence*; over these states of the world, self 2 would like to undertake effort but self 1 does not wish self 2 to do so. The overconfidence cost is higher, the more reliability (higher r^*) self 2 places on the signal from self 1 because from (18.83), $\theta(r^*)$ is increasing in r^* .

Suppose that self 1 gets bad news, $\sigma = L$. The problem of self 1 is to choose the recall probability λ in order to maximize

$$\lambda U_R(\theta_L) + (1 - \lambda) U_C(\theta_L). \quad (18.90)$$

given the reliability, r^* , induced by his action.

Definition 18.7 A Perfect Bayesian Equilibrium is a pair of strategies (λ^*, r^*) where $\lambda^* \in [0, 1]$, $r^* \in [q, 1]$ such that:

- (a) Self 1 chooses the recall probability, λ^* , optimally to maximize (18.90), given the assessment reliability of self 2, r^* .
- (b) Self 2's choice of r^* is consistent with Bayes' rule, (18.82), given the optimal choice of self 1.

Using (18.90), the first order condition to self 1's problem is

$$U_R(\theta_L) - U_C(\theta_L) = 0. \quad (18.91)$$

Using (18.83), (18.89) in (18.91), we get

$$\beta \int_{\beta\theta_L}^{\beta(r\theta_H+(1-r)\theta_L)} (\theta_L - c) dH(c) = 0. \quad (18.92)$$

The main implications of (18.92) are given in Proposition 18.22 below.

Proposition 18.22 (Bénabou and Tirole, 2002): *There exist values of $\beta : \beta_L, \beta_H \in (0, 1)$ and $\beta_L < \beta_H$ such that the following hold.*

- (a) *For $\beta > \beta_H$, the unique equilibrium is $\lambda^* = 1$, i.e., no news is hidden from self 2.*
- (b) *For $\beta < \beta_L$ the unique equilibrium is $\lambda^* = 0$, i.e., all (bad) news is hidden from self 2.*
- (c) *For $\beta \in [\beta_L, \beta_H]$ there are three possible equilibria, $\lambda^* \in \{0, \Gamma(\beta), 1\}$ where $\Gamma(\beta)$ is a decreasing function from 1 to 0 as β rises from β_L to β_H .*

The intuition behind Proposition 18.22 is straightforward. When β is very high, overconfidence is the main concern, thus, bad news is transmitted to self 2 to dampen her overconfidence. On the other hand, when β is low, the confidence of self 2 needs to be built up, so $\lambda^* = 0$. For intermediate values of β , a range of values $\lambda^* \in (0, 1)$ is optimal, so one gets multiple equilibria; see Figure 18.3.

The multiple equilibria can be explained as follows. Suppose that self 1 is very concealing of information (low λ^*). This makes self 2 discount such information even more and so the risk that self 2 is overconfident increases. This makes self 1 even more keen to conceal information. On the other hand if self 1 truthfully reveals most information (high λ^*) self 2 considers it very reliable (using Bayes' rule). This gives an incentive for self 1 to be even more truthful.




Figure 18.3 The optimal recall probability as β varies.

Source: Roland Bénabou and Jean Tirole, "Self-confidence and personal motivation." *The Quarterly Journal of Economics* (2002) 117(3), by permission of Oxford University Press.

In the fully censored equilibrium, $\lambda^* = 0$, no information is transmitted. In this “babbling equilibrium,” the incentive for information suppression is different from Carillo and Mariotti (2000) in that self 1 does not want to suppress good news. But in trying to suppress only the bad news, she invariably ends up suppressing the good news (which, in this case, is no news). This might do more harm than good. To see this, the payoff to self 1 upon receiving good news is

$$U_R(\theta_H) = \beta \int_0^{\beta\theta(r^*)} (\theta_H - c) dH(c), \quad (18.93)$$

where r^* is the “reliability” inferred by self 2. Note that

$$\beta \int_0^{\beta\theta(r^*)} (\theta_H - c) dH(c) < \beta \int_0^{\beta\theta_H} (\theta_H - c) dH(c) \quad (18.94)$$

when all information is truthfully revealed to self 2. Thus, a child who is praised by her parents in most circumstances will never be sure if she really did well, even when she actually does well. Using (18.90), (18.93) the expected payoff of self 1, before receiving the ability signal is

$$W(\lambda^*, r^*) = qU_R(\theta_H) + (1 - q)[\lambda U_R(\theta_L) + (1 - \lambda)U_C(\theta_L)].$$

Thus, if the strategy is to truthfully reveal all information to self 2, the payoff of self 1 is $W(1, 1)$. It is possible that $W(\lambda^*, r^*) < W(1, 1)$.

Example 18.12 Let $q = 1/2$, $\theta_L = 0$, $\theta_H = 1$, and $\bar{c} < 1$. Suppose $\lambda^* = 0$, so by Bayes’ rule $r^* = q$. Note that

$$W(\lambda^*, r^*) - W(1, 1) = (1 - q) \int_{\beta\theta_L}^{\beta\theta(r^*)} (\theta_L - c) dH(c) - q \int_{\beta\theta(r^*)}^{\beta\theta_H} (\theta_H - c) dH(c). \quad (18.95)$$

Substituting $r^* = q = 1/2$, $\theta_L = 0$, $\theta_H = 1$ in (18.95) we get:

$$W(\lambda^*, r^*) - W(1, 1) = -\frac{1}{2} \left[\int_0^{\beta/2} c dH(c) + \int_{\beta/2}^{\beta} (1 - c) dH(c) \right]. \quad (18.96)$$

Since $\bar{c} < 1$ we get $W(\lambda^*, r^*) < W(1, 1)$, hence, the individual will not benefit from hiding information from subsequent selves.

PERSONAL EVALUATION BIASES AND BAYES’ RULE

In experimental evidence, most individuals overestimate their abilities in absolute terms and their abilities relative to others. Bénabou and Tirole (2002) cite survey evidence that shows 94% of college professors thought they were better than their average colleague. We now turn to the question of the rationality of such beliefs.

Consider a large population that, using the law of large numbers, allows us to identify q , $1 - q$ as the proportions of the population that, respectively, receives the good ($\sigma = \phi$) and the bad ($\sigma = L$) signals. Suppose that less than half the population receives the good news, i.e., $q < 1/2$.

Selves 1 belonging to the proportion q can only signal good news to Selves 2. Suppose that the proportion $1 - q$, which receives the bad news, uses an identical optimal recall probability

$\lambda^* \in (0, 1)$; a symmetry argument or a unique equilibrium outcome can be used to justify this. Denote by r^* , the corresponding reliability placed by self 2 on the signal transmitted by self 1.

A fraction $1 - \lambda^*$ of Selves 2 of individuals who had bad news in period 1 forget about the bad news. Thus, the proportion of Selves 2 who overestimate their ability to be high, but actually have low ability, is $(1 - q)(1 - \lambda^*)$. The overestimation of their ability is $\theta(r^*) - \theta_L$; using (18.83) this is given by $r^*(\theta_H - \theta_L)$. Hence, the total proportion of period 2 selves who think that they have high ability is given by

$$q + (1 - q)(1 - \lambda^*).$$

If $\lambda^* < 1$, then this proportion is greater than the initial proportion ($q < 1/2$) that thinks that its ability is high. In fact if λ^* is not too high in the sense that $\lambda^* < \frac{1}{2(1-q)}$, then a majority of the population now, rationally, thinks that it has high ability and an ability greater than the average.

We have seen elsewhere in this book that Bayes' rule is not an accurate description of the human inference process; for instance it neglects the base rate. In an earlier paper Bénabou and Tirole (1999) use a modified, less sophisticated form of the Bayes' rule as compared to (18.82); in particular, they use

$$r \equiv \Pr[\hat{\sigma} = \phi \mid \sigma = \phi] = \frac{q}{q + \pi(1 - q)(1 - \lambda)}, \quad (18.97)$$

where $\pi \in [0, 1]$ is a measure of the individual's cognitive sophistication. So, if $\pi = 0$, then $\Pr[\hat{\sigma} = \phi \mid \sigma = \phi] = 1$, i.e., self 2 naively trusts the signal transmitted by self 1. At the other extreme, if $\pi = 1$, then the decision maker is a fully rational Bayesian decision maker. Intermediate values of π allows for more limited cognitive sophistication. The unattractiveness of $\pi = 0$ has already been alluded to before; the assumption of a complete lack of introspection is clearly a strong one. Bénabou and Tirole show that the qualitative results carry through for intermediate values of π . Hence, there does not seem to be a gain from using the alternative version of Bayes' rule in (18.97) in this case.

WILLPOWER AND PERSONAL RULES

Bénabou and Tirole (2004) look at the related issue of signaling one's willpower to their future selves. Each self now lives over two subperiods; making the decision to undertake an activity in the first and the decision to quit or continue the activity in the second. In subperiod 1, Self 1 could undertake an activity that has some limited current benefits, for instance, abstinence from smoking or starting a paper. In the second subperiod, continuation of the activity becomes much more onerous. For instance, the craving created by abstinence from smoking since the morning must have worsened, or in the case of writing the paper, fatigue could have set in. In the absence of any consideration for the future, the individual would perhaps give up on the activity and not realize the full fruits of the task started in subperiod 1. However, doing so might signal to self 2 that one's willpower is low. This has potentially negative consequences for the decision that self 2 makes in undertaking the task in her first subperiod. If self 1 is aware of these long-term consequences of her decision she might continue with the activity in subperiod 2. Such models can account for personal rules such as I will quit smoking from now on.

Exercises for Part 6

1. (Laibson, 2001) Consider the model in Section 17.3. Prove Proposition 17.1, i.e., show in the context of the no-cues model that the optimal decision rule of the individual is

$$a = \begin{cases} 1 & \text{if } x > \tilde{x} \\ 0 & \text{if } x < \tilde{x} \end{cases}.$$

Also show that the threshold value $\tilde{x}$ is increasing in the outside option, $\bar{u}$, the discount factor, δ , and the compensatory process weighting factor, θ .

2. (Caplin and Leahy, 2001) Consider the *psychological expected utility model* in Section 17.5.
 - (a) (Possibility and probability) There is evidence that economic agents might be influenced by *possibility* rather than *probability*; for instance, Monat et al. (1972), Bankhart and Elliot (1974), Loewenstein et al. (2001). Suppose that there is an announcement that the Federal Reserve is concerned about inflation. Why might this alter the investment behavior of economic agents independent of the probability of an interest rate increase?
 - (b) (Time course of anxiety) The time path of anxiety tends to be U shaped; see for instance, Breznitz (1984). Fear is greatest when individuals are first informed of an aversive event. It then diminishes for some time before increasing sharply in anticipation of the event. In light of this evidence what do you expect to be the effect on price of those assets whose risks have only been recently revealed relative to those whose risks have been known for some time?
 - (c) Suppose that the Federal Reserve comes across some bad piece of news such as the possible shutdown of a major bank. Assume that economic agents suffer from anxiety. Show why it might now be optimal to hide the bad news for longer relative to the standard case without anxiety.
3. (Caplin and Leahy, 2004)¹⁵ Hopefulness, suspense, and anxiety are feelings related to uncertainty about the future. An anticipatory feeling is a feeling that anticipates the outcome of an event. These feelings are embedded here in an otherwise standard two-period consumption and saving problem. Consider an individual who is young and works in period t and is retired in period $t + 1$; labor income is the only source of income and there are no inheritances or bequests. The budget constraint when young (period t) is given by

$$c_t + s_t = w_t,$$

where w_t is the wage, c_t is consumption and s_t is savings in period t . The budget constraint when old (period t) is given by

$$c_{t+1} = (1 + r_{t+1}) s_t,$$

where the return r_{t+1} on savings is the only source of uncertainty in the model. Let $\mu \equiv E_t [1 + r_{t+1}]$, $\sigma^2 \equiv \text{var}_t [1 + r_{t+1}]$. Preferences of the individual are given by

$$U_t = E_t [\gamma [-\phi \ln a_t + \ln c_t] + (1 - \gamma) \ln c_{t+1}],$$

¹⁵ I am grateful to Klaus Wälde for suggesting this problem.

where a_t captures the utility-reducing anticipatory feeling in period t , $\gamma, \phi \in [0, 1]$ are preference parameters. Following Caplin and Leahy (2004), we model anticipatory feelings as a reduction in mean and an increase in the variance of uncertain consumption c_{t+1} . In order to allow for positive and negative values of anxiety, we write anxiety as

$$\ln a_t = \zeta \text{var}_t c_{t+1} - (1 - \zeta) E_t c_{t+1}, \quad 0 \leq \zeta \leq 1. \quad (6.E1)$$

E_t and var_t are, respectively, the expectation and variance operators, conditional on the information set at time t .

- (a) Compute the first order condition and interpret it in the two polar cases, $\mu = 0$ and $\sigma^2 = 0$.
- (b) Consider a more general version of the objective function

$$U_t = E_t [\gamma u(a_t, c_t) + (1 - \gamma) u(c_{t+1})].$$

Compute the first order condition for the optimal saving decision and provide an economic interpretation.

4. (Loewenstein et al., 2003) Extend the two-period model in Section 17.6.2 to three periods, $T = 3$. Recall that (c_1^b, c_2^b, c_3^b) is the optimal profile of consumption in time period $t = 1$ when the individual suffers from projection-bias ($\alpha > 0$). As time passes by, the individual's tastes change in a way that she did not predict. The individual has a chance to reoptimize in period 2. Denote by (c_2^{bb}, c_3^{bb}) the optimal consumption profile when the individual reoptimizes in time period $t = 2$ and suffers from projection bias. Show that if $c_1^b > s_1$, then we get the following result. The sign of v''' determines completely the sign of $c_2^{bb} - c_2^b$. In particular, if $v''' > 0$ (which holds for CARA and CRRA preferences) projection bias will repeatedly lead people to adjust their consumption plans upwards ($c_2^{bb} > c_2^b$). If $v''' < 0$, then $c_2^{bb} < c_2^b$ and if $v''' = 0$, then $c_2^{bb} = c_2^b$. Does this result provide an explanation for why individuals engage in inadequate savings?
5. (Fudenberg and Levine, 2006) Consider the four-period problem in O'Donoghue and Rabin (1999) with (β, δ) preferences, adapted to the case of a dual-self model. An agent has 4 periods to do a report. In the original problem, the pecuniary cost in successive periods of completing the report is 3, 5, 8, 13, and the benefit from doing the report is v . Assume now that the total cost of doing the report in any period is $1 + \gamma$ times the pecuniary cost, where $\gamma > 0$ is the self-control cost per unit. In any period, the decision to act is taken by the short-run self. Show the following.
 - (a) If $\gamma < 10/3$, the agent acts to write the report in the first period. If prevented from writing the report in period 1, then any of the subsequent short-run selves are willing to write the report.
 - (b) If $\gamma > 10/3$, then the agent never finds it optimal to act.
 - (c) The knife edge case, $\gamma = 10/3$, is the only case where the agent is willing to act in intermediate periods. [Hint: Solve the game backwards starting in period 4, being careful to include the self-control costs.]
6. (Lipman and Pesendorfer, 2013) Show that classical preferences (Definition 17.4) are a special case of self-control preferences (see Proposition 17.13) if either (i) v is an affine transformation of u , i.e., $v = \alpha u + \beta$ for some $\alpha \geq 0$ and $\beta \in \mathbb{R}$, or (ii) u is a constant function.

Show also that when $\alpha < 0$, $\beta \in \mathbb{R}$, and $v = \alpha u + \beta$, the individual has a sole desire for commitment.

7. (Lipman and Pesendorfer, 2013) Consider the setup of Example 17.11 for (β, δ) preferences. Benhabib et al. (2010) propose an alternative to (β, δ) preferences by introducing a fixed cost of delay. Thus, in period 1, the consumption pair $(0, l)$ is preferred over $(s, 0)$ if $\delta w(l) - f > w(s)$, where f is a fixed cost of delay associated with the later outcome, l , and $\delta > 0$ is the discount factor. What functional forms for commitment utility and temptation utility can be used to represent these preferences?
8. (Lipman and Pesendorfer, 2013) Consider the experimental results of Dana et al. (2006) for dictator games that we considered in Part 2 of the book. The dictator splits \$10 between himself and a passive recipient. Many dictators prefer to take \$9 and not play the game; in this case, the recipient is given no information about the experimenter's interaction with the dictator. How can you transform this into a problem with Gul–Pesendorfer temptation preferences?
9. (Lipman and Pesendorfer, 2013; Dekel et al., 2009) Consider the model of temptation preferences in Section 17.7.
 - (a) Let b denote broccoli, c denote chocolate, and p denote potato chips; c and p are the two snacks. Suppose that we have $C = \{b, c, p\}$. A dieting individual expresses the following first period preferences: $\{b\} \succ_1 \{b, c\} \succ_1 \{b, p\} \succ_1 \{b, c, p\}$ (the middle two preferences can be reversed without altering the conclusion). The preferences imply that the individual prefers to have no snacks to having at least one snack and finds having both snacks on the menu to be the worst option. These preferences appear to be plausible. Show that these expressed preferences are not consistent with the Gul–Pesendorfer and the Strotz preferences.
 - (b) Let h be a healthy dish; d a small dessert; and D a large dessert. Let $C = \{h, d, D\}$. A dieting individual expresses the following first period preferences: $\{h, d\} \succ_1 \{d\}$ and $\{h, d, D\} \succ_1 \{h, D\}$. The first preference says that the individual would like to have the option of keeping to his diet. The second preference says that the individual may wish to keep d as a compromise option on the menu. These preferences also appear plausible. However, show that they are not consistent with the Gul–Pesendorfer (u, v) preferences.
 - (c) How would you incorporate the plausible preferences expressed in (a) and (b) into a theory of temptation preferences? [Hint: One possibility is that preferences violate WARP. A second possibility, explored formally in Dekel and Lipman (2012), is that unlike in the Gul–Pesendorfer model, there might be uncertainty. For instance, the preference $\{b, p\} \succ_1 \{b, c, p\}$ in (a) might arise because the individual is uncertain whether he will succumb to temptation. Having two tempting snacks, c and p , increases the probability that he will succumb to temptation, so he prefers it the least.]
10. (Ozdenoren et al., 2012) Consider a cake eating problem in continuous time $[0, T]$ in which a consumer maximizes utility by choosing an optimal path of consumption. The instantaneous utility function at time t is given by a differentiable function $U(c(t))$ with $U(0) = 0$, $U_c > 0$, and $U_{cc} < 0$. We do not consider addictive behaviors so U does not depend on past consumption. The size of the cake remaining at time t is given by $R(t)$. The initial size of the cake, $R(0)$, is given. The main difference from the standard cake eating problem in economics comes from the assumption that the consumer is endowed with a given *stock of willpower* $W(0)$. It is possible to restrain consumption but at the cost of a depletion in willpower. The rate of willpower depletion is given by $f(W(t), c(t))$ with $f_W \leq 0$, and $f_c < 0, f_{cc} > 0$ for $c \in [0, \bar{c})$. The convexity in c indicates that the more the individual

restrains his consumption, the faster does willpower deplete. f is also assumed to be twice differentiable everywhere except at $\bar{c}$, continuous at $\bar{c}$, and $f_{cW} \geq 0$. Finally, $f > 0$ for $c \in [0, \bar{c})$, and $f = 0$ for $c \in [\bar{c}, \infty)$ for some $\bar{c} > 0$. Any willpower, $W(t')$, left over after the individual has finished the cake is devoted to an alternative activity that generates utility $m(W(t'))$; m is a constant function, and

$$t' = \sup \{t \in [0, T] : R(t) > 0\}.$$

Based on the description above, the optimization problem of the individual (Problem P-1) can, thus, be written as follows. The individual chooses $c(t)$ to maximize

$$V(0) = \int_0^T e^{-\rho t} U(c(t)) dt + e^{-\rho t'} m(W(t')) \quad (\text{P-1})$$

subject to

$$\begin{aligned} \dot{R}(t) &= -c(t) \\ \dot{W}(t) &= \begin{cases} -f(W(t), c(t)) & \text{if } R(t) > 0 \\ 0 & \text{otherwise} \end{cases} \\ R(T) &\geq 0; W(T) \geq 0; R(0) = \bar{R} \geq 0; W(0) = \bar{W} \geq 0. \end{aligned}$$

ρ is the subjective rate of discount, and $\bar{W}$ is the initial level of willpower.

(a) Derive the solution to Problem P-1 by forming the appropriate Hamiltonian. Show that the first order condition for consumption at time t can be written as

$$e^{-\rho t} U'(c(t)) + \lambda(t) (-f_c) = \alpha(t),$$

where $\lambda(t)$ is the Lagrangian multiplier on the $\dot{W}(t)$ constraint and $\alpha(t)$ is the Lagrangian multiplier on the $\dot{R}(t)$ constraint. Interpret the first order condition, in particular, the second term on the LHS that sets this condition apart from the standard conditions that ignore issues of willpower.

(b) Suppose that there is no discounting, i.e., $\rho = 0$. If $c(t) > 0$ for $t \leq s \in [0, T]$, show that

$$\dot{c}(t) = \lambda(t) \frac{f_W f_c - f_c W f}{U'' - \lambda(t) f_{cc}}. \quad (6.E2)$$

When $\lambda = 0$, we get the classical conclusion that so long as $c(t) > 0$, $\dot{c}(t) = 0$. However, when $\lambda > 0$, whether we have an increasing or decreasing profile of consumption depends on the sign of the numerator in (6.E2).

(c) This part demonstrates the importance of the initial level of willpower. Denote by W_H the minimum level of $W(0)$ such that $c(t) = \bar{R}/T$ is feasible for all $t \in [0, T]$. Denote the optimal consumption path by $c^*(t)$. Show the following. (i) If $\bar{W} \geq W_H$, then the cake is exhausted and $c^*(t) = \bar{R}/T$ for all $t \in [0, T]$. (b) If $\bar{W} < W_H$, then the cake and willpower are exhausted ($R(s) = 0$, $W(s) = 0$ for some $s \in [0, T]$) and when $c(t) > 0$, whether we have an increasing or decreasing profile of consumption depends on the sign of the numerator in (6.E2).

11. (Krussel et al., 2010) Consider a representative agent two-period model with time $t = 1, 2$. Consumption in period t is denoted by c_t . The commitment ordering over outcomes is given by $\tilde{u}(c_1, c_2) = u(c_1) + \delta u(c_2)$ where $\delta > 0$ is the discount factor and the temptation ordering is given by $\tilde{v}(c_1, c_2) = \gamma (v(c_1) + \beta \delta v(c_2))$ where $\beta \in (0, 1]$. Hence, while the commitment ordering takes the standard exponential discounting form, the temptation ordering takes the hyperbolic discounting form. Preferences are represented as in the Gul–Pesendorfer model, i.e.,

$$\max_{c_1, c_2} \{\tilde{u}(c_1, c_2) + \tilde{v}(c_1, c_2)\} - \max_{c_1, c_2} \tilde{v}(c_1, c_2). \quad (6.E3)$$

Each individual is endowed with one unit of labor in each period and an endowment of capital equal to $k_1 > 0$ at the beginning of period 1. The factors of production, labor and capital, are used by perfectly competitive firms to produce output. The production function of firms in period $t = 1, 2$ is given by $f(\bar{k}_t)$ where $\bar{k}_t$ is the aggregate per capita capital stock. Hence, in any period, the factor incomes to labor and capital are given respectively by $w(\bar{k}_t)$ and $r(\bar{k}_t)$. Let $R(\bar{k}_t) = 1 + r(\bar{k}_t)$ denote the gross return on capital. The budget constraint of the representative agent is given by

$$B(k_1, \bar{k}_1, \bar{k}_2) \left\{ (c_1, c_2) : c_1 = R(\bar{k}_1)k_1 + w(\bar{k}_1) - k_2, c_2 = R(\bar{k}_2)k_2 + w(\bar{k}_2) \right\}. \quad (6.E4)$$

- (a) Show that the consumer saves the least if second period preferences were simply $\tilde{v}(c_1, c_2)$, saves the most if second period preferences were $\tilde{u}(c_1, c_2)$, and saves an intermediate amount when preferences are given by $\tilde{u}(c_1, c_2) + \tilde{v}(c_1, c_2)$.
 (b) Now consider autarky in which the individual uses his own backyard technology, which we assume is given by $f(\bar{k}_t)$ and there is no recourse to private markets. Show that the same allocation is achieved as under (a). However, utility is higher under autarky. [Hint: The agent cannot trade along the market line under autarky and is limited to the opportunity set which lies under the market line. Hence, there is lower temptation under autarky].
 (c) Suppose now that the government undertakes a balanced budget fiscal policy. It levies a tax τ_y on total income and a tax τ_i on first period investment. This gives the modified budget constraint

$$\begin{aligned} B(k_1, \bar{k}_1, \bar{k}_2) &= \left\{ (c_1, c_2) : c_1 = \left[R(\bar{k}_1)k_1 + w(\bar{k}_1) \right] (1 - \tau_y) - (1 + \tau_i)k_2, c_2 \right. \\ &= \left. R(\bar{k}_2)k_2 + w(\bar{k}_2) \right\} \end{aligned}$$

and the balanced budget constraint of the government is $\tau_y \left[R(\bar{k}_1)k_1 + w(\bar{k}_1) \right] = -\tau_i \bar{k}_2$. Suppose that the government is a benevolent central planner who wishes to maximize the agent's indirect utility by optimally choosing τ_y, τ_i . Derive the optimal tax rates and show that it is optimal for the government to impose an investment subsidy (i.e. $\tau_i < 0$). [The intuition is that an investment subsidy increases the cost of first period temptation that would arise from the possibility of reducing savings.]

12. (Amador et al., 2006) Consider the following three-period economy, where the time index $t = 0, 1, 2$. Society moves in period $t = 0$ to choose an appropriate constitution, while the government moves in period $t = 1$ to choose some level of public spending. Society's

welfare is given by $\theta U(g) + W(c)$, where g is the level of government spending and c is private consumption. The government wants to maximize instead $\theta U(g) + \beta W(c)$, where $0 < \beta \leq 1$ and $1 - \beta$ is a measure of the strength of temptation towards government spending. The resource constraint is given by $B = \{(c, g) \in \mathbb{R}_+^2 : c + g \leq y\}$. A fiscal constitution is a subset $C \subset B$ that constrains the government to choose $(c, g) \in C$. The government has private information on θ . Restrict attention to pure commitment mechanisms with no transfers across types. An optimal constitution is some subset of C which maximizes the welfare of society, given the subsequent bias of the government towards g . Show that it is always optimal to limit government spending. How does the optimal solution trade off commitment versus flexibility?

REFERENCES FOR PART 6

- Alesina, A., Di Tella, R., and MacCulloch, R. (2004). Inequality and happiness: are Europeans and Americans different? *Journal of Public Economics* 88: 2009–42.
- Amador, M., Werning, I., and G. M. Angeletos (2006). "Commitment vs. flexibility. *Econometrica* 74(2): 365–96.
- Ameriks, J., Caplin, A., Leahy, J., and Tyler, T. (2007). Measuring self-control problems. *American Economic Review* 97: 966–72.
- Angelucci, M. and De Giorgi, G. (2009). Indirect effects of an aid program: how do cash transfers affect ineligible's consumption? *American Economic Review* 99(1): 486–508.
- Bankhart, C. and Elliot, R. (1974). Heart rate and skin conductance in anticipation of shocks with varying probability of occurrence, *Psychophysiology* 11: 160–174.
- Battigalli, P., Dufwenberg, M., and Smith, A. (2015). Frustration and anger in games. Working Paper No. 539. IGER – Università Bocconi.
- Baumeister, R. F., De Wall, C. N., and Zhang, L. (2007). Do emotions improve or hinder the decision making process? In K. D. Vohs, R. F. Baumeister, and G. Loewenstein (eds.), *Do Emotions Help Or Hurt Decision Making? A Hedgesfoxian Perspective*. New York: Russell Sage Foundation Press.
- Baumeister, R. F., Heatherton, T. F., and Tice, D. M. (1994). *Losing Control: How and Why People Fail at Self-Regulation*. San Diego, CA: Academic Press.
- Baumeister, R. F. and Vohs, K. D. (2003). Willpower, choice, and self-control, in G. Loewenstein, D. Read, and R. F. Baumeister (eds.), *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice*. New York: Russell Sage Foundation, 201–16.
- Becchetti L., Pelloni, A., and Rossetti, F. (2008). Relational goods, sociability, and happiness. *Kyklos* 61(3): 343–63.
- Bechara, A., Damasio, A. R., Damasio, H., and Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition* 50(1–3), 7–15.
- Bechara, A., Damasio, H., Damasio, A. R., and Lee, G. P. (1999). Different contributions of the human amygdala and ventromedial prefrontal cortex to decision-making. *Journal of Neuroscience* 19: 5473–81.
- Bechara, A., Damasio, H., Tranel, D., and Damasio, A. R. (1997). Deciding advantageously before knowing the advantageous strategy. *Science* 275(5304), 1293–95.
- Bechara, A., Tranel, D., Damasio, H., and Damasio, A. R. (1996). Failure to respond autonomically to anticipated future outcomes following damage to prefrontal cortex. *Cerebral Cortex* 6: 215–25.
- Becker, G. S. (1996). *Accounting for Tastes*. Cambridge, MA: Harvard University Press.
- Becker, G. S. and Murphy, K. M. (1988). A theory of rational addiction. *Journal of Political Economy* 96(4): 675–701.
- Becker, G. S., and Murphy, K. M. (1993). A simple theory of advertising as a good or a bad, *Quarterly Journal of Economics* 108: 941–64.
- Beja E. Jr. (2014). Income growth and happiness: reassessment of the Easterlin Paradox, Munich Personal RePEc Archive Paper No. 53360.
- Bell, D., (1985). Disappointment in decision making under uncertainty. *Operations Research* 33: 1–27.
- Bénabou, R. and Pycia, M. (2002). Dynamic inconsistency and self-control. *Economics Letters* 77(3): 419–24.
- Bénabou, R. and Tirole, J. (1999). Self-confidence: intrapersonal strategies. IDEI mimeo.
- Bénabou, R. and Tirole, J. (2002). Self-confidence and personal motivation. *Quarterly Journal of Economics* 117: 871–915.
- Bénabou, R. and Tirole, J. (2003). Intrinsic and extrinsic motivation. *Review of Economic Studies* 67: 529–44.
- Bénabou, R. and Tirole, J. (2004). Willpower and personal rules. *Journal of Political Economy* 112: 848–87.
- Benhabib, J. and Bisin, A. (2005). Modeling internal commitment mechanisms and self-control: a neuroeconomics approach to consumption-saving decisions. *Games and Economic Behavior* 52(2): 460–92.
- Benhabib, J., Bisin, A., and Schotter, A. (2010). Present-bias, quasi-hyperbolic discounting, and xed costs. *Games and Economic Behavior* 69(2): 205–23.

- Benjamin, D. J., Brown, S. A., and Shapiro, J. M. (2013). Who is “Behavioral”? Cognitive ability and anomalous preferences. *Journal of the European Economic Association* 11(6): 1231–255.
- Benjamin, D. J., Heffetz, O., Kimball, M. S., and Rees-Jones, A. (2012). What do you think would make you happier? What do you think you would choose? *American Economic Review* 102(5): 2083–110.
- Benjamin, D. J., Heffetz, O., Kimball, M. S., and Szembrot, N. (2014). Beyond happiness and satisfaction: toward well-being indices based on stated preference. *American Economic Review* 104(9): 2698–735.
- Bentham, J. (1789). *An Introduction to the Principle of Morals and Legislations* (reprinted Oxford: Blackwell 1948).
- Bernheim, B. D. and Rangel, A. (2004). Addiction and cue-triggered decision processes. *American Economic Review* 94(5): 1558–90.
- Berns, G., Cohen, D., and Mintun, M. (1997). Brain regions responsive to novelty in the absence of awareness. *Science*, 276: 1272–5.
- Blanchflower, D. G. and Oswald, A. (2004). Well-being over time in Britain and the USA. *Journal of Public Economics* 88(7–8): 1359–86.
- Bodner, R. and Prelec, D. (2003). The diagnostic value of actions in a self-signaling model, in I. Brocas and J. Carrillo (eds.), *The Psychology of Economic Decisions, Volume 1*. Oxford: Oxford University Press.
- Bolton, G. E. and Ockenfels, A. (2000). ERC – a theory of equity, reciprocity, and competition. *American Economic Review* 90(1): 166–93.
- Breznitz, S., (1984) *Cry Wolf*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Brickman, P. and Campbell, D. T. (1971). Hedonic relativism and planning the good society. In M. H. Appleby (ed.), *Adaptation Level Theory*. New York: Academic Press, pp. 287–302.
- Brickman, P., Coates, D., and Janoff-Bulman, R. (1978). Lottery winners and accident victims: is happiness relative? *Journal of Personality and Social Psychology* 36: 917–27.
- Brocas, I. and Carrillo, J. (2008). The brain as a hierarchical organization. *American Economic Review* 98(4): 1312–46.
- Brocas, I. and Carrillo, J. (2014). Dual-process theories of decision making: a selective survey. *Journal of Economic Psychology* 41: 45–54.
- Caplin, A. and Eliaz, K. (2003). AIDS policy and psychology: a mechanism-design approach. *The RAND Journal of Economics* 34(4): 631–46.
- Caplin, A. and Leahy, J. (2001). Psychological expected utility and anticipatory feelings. *Quarterly Journal of Economics* 116: 55–79.
- Caplin, A. and Leahy, J. (2004). The supply of information by a concerned expert. *Economic Journal* 114(497): 487–505.
- Carillo, J. D. and Mariotti, T. (2000). Strategic ignorance as a self-disciplining device. *Review of Economic Studies* 67: 529–44.
- Carter, S. and McBride, M. (2013). Experienced utility versus decision utility: putting the “S” in satisfaction. *Journal of Socio-Economics* 42: 13–23.
- Chaiken, S. and Trope, Y. (eds.) (1999). *Dual-Process Theories in Social Psychology*. New York: Guilford Press.
- Cheibub, J. A., Gandhi, J., and Vreeland, J. R. (2010). Democracy and dictatorship revisited. *Public Choice* 143: 67–101.
- Chen, M. K., Lakshminarayanan, V., and Santos, L. R. (2006). How basic are behavioral biases? Evidence from capuchin monkey trading behavior. *Journal of Political Economy* 114(3): 517–37.
- Clark, A. E. (2001). What really matters in a job? Hedonic measurement using quit data. *Labour Economics* 8: 223–42.
- Clark, A. E. (2003). Unemployment as a social norm: psychological evidence from panel data, *Journal of Labor Economics* 21: 323–51.
- Clark, A. E. (2006). A note on unhappiness and unemployment duration. *Applied Economics Quarterly* 52(4): 291–308.
- Clark, A. E., Georgellis, Y., and Sanfey, P. (2001). Scarring: the psychological impact of past unemployment. *Economica* 68(270): 221–41.
- Clark, A. E. and Oswald, A. J. (1994). Unhappiness and unemployment, *Economic Journal* 104: 648–59.
- Clark A. E. and Oswald A. (1996). Satisfaction and comparison income. *Journal of Public Economics* 61: 359–81.
- Clark, A. E. and Senik, C. (2010). Who compares to whom? The anatomy of income comparisons in Europe. *The Economic Journal* 120: 573–94.
- Cohen, J. D. (2005). The vulcanization of the human brain: a neural perspective on interactions between cognition and emotion. *Journal of Economic Perspectives* 19(4): 3–24.

- Coleman, J. S. (1993). Social capital in the creation of human capital. *American Journal of Sociology* 94: 95–120.
- Damasio, A. R. (1994). *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: G.P. Putman.
- Damasio, A. R. (1996). The somatic marker hypothesis and the possible functions of the pre-frontal cortex. *Philosophical Transactions of the Royal Society London Series B: Biological Sciences* 351: 1413–20.
- Dana, J. D., Dalian, M., Cain, D. M., and Dawes, R. M. (2006). What you don't know won't hurt me: costly (but quiet) exit in a dictator game. *Organizational Behavior and Human Decision Processes* 100(2): 193–201.
- Dawkins, R. (2006). *The Selfish Gene (30th Anniversary Edition)*. New York City: Oxford University Press.
- De Neve, J. E. (2011). Functional polymorphism (5-HTTLPR) in the serotonin transporter gene is associated with subjective well-being: evidence from a US nationally representative sample. *Journal of Human Genetics* 56: 456–9.
- De Neve, J. E., Christakis, N. A., Fowler, J. H., and Frey, B. S. (2012). Genes, economics, and happiness. *Journal of Neuroscience, Psychology, and Economics* 5: 193–211.
- Deaton, A. (2008). Income, health, and well-being around the world: evidence from the gallup world poll. *Journal of Economic Perspectives* 22: 53–72.
- Dekel, E. and Lipman, B. L. (2012). Costly self-control and random self-indulgence. *Econometrica* 80(3): 1271–302.
- Dekel, E., Lipman, B., and Rustichini, A. (2001). Representing preferences with a unique subjective state space. *Econometrica* 69(4): 891–934.
- Dekel, E., Lipman, B., and Rustichini, A. (2009). Temptation-driven preferences. *Review of Economic Studies* 76(3): 937–71.
- DellaVigna, S. and Malmendier, U. (2002). Overestimating self-control: evidence from the health club industry, Working Paper, University of California at Berkeley and Stanford University.
- Di Tella, R., Haisken-De New, J., and MacCulloch, R. J. (2010). Happiness adaptation to income and to status in an individual panel. *Journal of Economic Behavior and Organization* 76(3): 834–52.
- Di Tella, R. and MacCulloch, R. J. (2006). Some uses of happiness data in economics. *Journal of Economic Perspectives* 20(1): 25–46.
- Di Tella, R., MacCulloch, R. J., and Oswald, A. J. (2001). Preferences over inflation and unemployment: evidence from surveys of happiness. *The American Economic Review* 91: 335–41.
- Di Tella, R., MacCulloch, R. J., and Oswald, A. J. (2003). The macroeconomics of happiness. *Review of Economics and Statistics* 85: 809–27.
- Diener, E. and Chan, M. Y. (2011). Happy people live longer: subjective well-being contributes to health and longevity. *Applied Psychology: Health and Well-Being* 3(1): 1–43.
- Diener, E., Lucas, R. E., and Scollon, C. N. (2006). Beyond the hedonic treadmill: revising the adaptation theory of well-being. *American Psychologist* 61(4): 305–14.
- Diener, E. and Seligman, M. E. P. (2004). Beyond money: toward an economy of well-being. *Psychological Science in the Public Interest* 5(1): 1–31.
- Diener, E., Tay, L., and Oishi, S. (2013). Rising income and the subjective well-being of nations. *Journal of Personality and Social Psychology* 104(2): 267–76.
- Ditto, P. H., Pizzaro, D. H., Epstein, E. B., Jacobson, J. A., and MacDonald, T. K. (2006). Visceral influences on risk-taking behavior. *Journal of Behavioral Decision Making* 19: 1–15.
- Duesenberry, J. (1949). *Income, Saving, and the Theory of Consumer Behavior*. Cambridge, MA: Harvard University Press.
- Dynan, K. (2000). Habit formation in consumer preferences: evidence from panel data. *American Economic Review* 110: 391–406.
- Easterlin, R. A. (1974). Does economic growth improve the human lot? Some empirical evidence. In P. A. David and M. W. Reder (eds.), *Nations and Households in Economic Growth: Essays in Honor of Moses Abramovitz*. New York: Academic Press, pp. 89–125.
- Easterlin, R. A. (1995). Will raising the incomes of all increase the happiness of all? *Journal of Economic Behavior and Organization* 27(1): 35–47.
- Easterlin, R. A. (2001). Income and happiness: towards a unified theory. *Economic Journal* 111(473): 465–84.

- Easterlin, R. A. (2003). Explaining happiness. *Proceedings of the National Academy of Science* 100(19): 11176–83.
- Easterlin, R. (2013). Happiness and economic growth: the evidence. Discussion Paper No. 7187, Bonn: Institute for the Study of Labor.
- Easterlin, R., Mcvey, L., Switek, M., Sawangta, O., and Zweig, J. (2010). The happiness paradox revisited. *Proceedings of the National Academy of Sciences* 107(52): 22463–8.
- Elster, J. (1998). Emotions in economic theory. *Journal of Economic Literature* 36: 47–74.
- Epstein, L. (2006). An axiomatic model of non-Bayesian updating. *Review of Economic Studies* 73: 413–36.
- Epstein, L. and Kopylov, I. (2007). Cold feet. *Theoretical Economics* 2: 231259.
- Epstein, L., Noor, J., and Sandroni, A. (2008). Non-Bayesian updating: a theoretical framework. *BE Journal of Theoretical Economics* 3(2): 193–229.
- Fleurbaey, M. (2009). Beyond GDP: the quest for a measure of social welfare. *Journal of Economic Literature* 47(4): 1029–75.
- Fleurbaey, M. and Blanchet, D. (2013). *Beyond GDP. Measuring Welfare and Assessing Sustainability*. Oxford: Oxford University Press.
- Frank, R. H. (1988). *Passions Within Reasons: The Strategic Role of the Emotions*. New York: Norton.
- Frederick, S. and Loewenstein, G. (1999). Hedonic adaptation: In D. Kahnemann, E. Diener, and N. Schwarz (eds.), *Well Being: The Foundations of Hedonic Psychology*. New York: Russell Sage Foundation Press, pp. 302–29.
- Fredrickson, B. L. and Kahneman, D. (1993). Duration neglect in retrospective evaluations of affective episodes. *Journal of Personality and Social Psychology* 65: 45–55.
- Freud, S. (1924). *The Ego and the Id*, translated by Joan Riviere. New York: W. W. Norton, 1962.
- Frey, B. S. and Benz, M. (2008). Being independent is a great thing: subjective evaluations of self-employment and hierarchy. *Economica* 75(298): 362–83.
- Frey, B. S. and Stutzer, A. (2000). Happiness, economy and institutions. *Economic Journal* 110(446): 918–38.
- Frey, B. S. and Stutzer, A. (2002). What can economists learn from happiness research? *Journal of Economic Literature* 40(2): 402–35.
- Frijda, N. H. and Mesquita, B. (1994). The social roles and functions of emotions. In S. Kitayama and H. R. Markus (eds.), *Emotion and Culture: Empirical Studies of Mutual Influence*. Washington, DC: American Psychological Association.
- Frijters, P., Haiksen-Denew, J. P., and Shields, M. A. (2004). Money does matter! Evidence from increasing real income and life satisfaction in East Germany following reunification. *American Economic Review* 94: 730–40.
- Fudenberg, D. (2006). Advancing beyond advances in behavioral economics. *Journal of Economic Literature* 44: 694–711.
- Fudenberg, D. and Levine, D. K. (2006). A dual self model of impulse control. *American Economic Review* 96(5): 1449–76.
- Fudenberg, D. and Levine, D. K. (2011). Risk, delay, and convex self-control costs. *American Economic Journal: Microeconomics* 3: 34–68.
- Gardner, J. and Oswald, A. J. (2007). Money and mental wellbeing: a longitudinal study of medium-sized lottery wins. *Journal of Health Economics* 26: 49–60.
- Gilbert, D. T., Gill, M. J., and Wilson, T. D. (2002). The future is now: temporal correction in affective forecasting. *Organizational Behavior and Human Decision Processes* 88: 430–44.
- Glimcher, P. W., Camerer, C., Poldrack, R., and Fehr, E. (eds.) (2009). *Neuroeconomics: Decision Making and the Brain*. London: Academic Press.
- Goldie, P. (2009). *The Oxford Handbook of Philosophy of Emotion*, Oxford: Oxford University Press.
- Graham, C., Eggers, A. and Sukhtankar, S. (2004). Does happiness pay? An initial exploration based on panel data from Russia. *Journal of Economic Behavior and Organization* 55: 319–42.
- Gruber, J. and Koszegi, B. (2001). Is addiction “rational”? Theory and evidence. *Quarterly Journal of Economics* 116(4): 1261–303.
- Gul, F. and Pesendorfer, W. (2001). Temptation and self control. *Econometrica* 69(6): 1403–36.
- Gul, F. and Pesendorfer, W. (2004). Self-control and the theory of consumption. *Econometrica* 72(1): 119–58.
- Gul, F. and Pesendorfer, W. (2005). The revealed preference theory of changing tastes. *Review of Economic Studies* 72: 429–48.
- Gul, F. and Pesendorfer, W. (2007). Harmful addiction, *Review of Economic Studies* 74: 147–72.

- Güven, C., Senik, C. and Stichnoth, H. (2012). You can't be happier than your wife. Happiness gaps and divorce. *Journal of Economic Behavior and Organization* 82(1): 110–30.
- Helliwell, J. F. (2007). Well-being and social capital: does suicide pose a puzzle?. *Social Indicators Research* 81: 455–96.
- Helliwell J. and Putnam R. D. (2004). The social context of well-being. *Philosophical Transactions of the Royal Society (London), Series B [Internet]* 359: 1435–46.
- Inglehart, R., Foa, R., Peterson, C., and Welzel, C. (2008). Development, freedom, and rising happiness. *Perspectives on Psychological Science* 3(4): 264–85.
- Jevons, W. (1905). *Essays on Economics*. London: Macmillan.
- Kahneman, D., Fredrickson, B. L., Schreiber, C. A., and Redelmeier, D. A. (1993). When more pain is preferred to less: adding a better end. *Psychological Science* 4: 401–5.
- Kahneman, D. and Thaler, R. H. (2006). Anomalies: utility maximization and experienced utility. *Journal of Economic Perspectives* 20(1): 221–34.
- Kahneman, D., Wakker, P. P., and Sarin, R. (1997). Back to Bentham? Explorations of experienced utility. *Quarterly Journal of Economics* 112: 375–405.
- Kassenboehmer, S. C. and Haisken-DeNew, J. P. (2009). You're fired! The causal negative effect of entry unemployment on life satisfaction. *Economic Journal* 119(536): 448–62.
- Kenny, C. (1999) Does growth cause happiness, or does happiness cause growth? *Kyklos* 52(1): 3–26.
- Knabe, A. and Rätzl, S. (2011). Scarring or scaring? The psychological impact of past unemployment and future unemployment risk. *Economica* 78(310): 283–93.
- Knight, J. and Gunatilaka, R. (2012). Income, aspiration and the hedonic treadmill in a poor society. *Journal of Economic Behavior and Organization* 82(1): 67–81.
- Knight, J., Song, L., and Gunatilaka, R. (2009). SWB and its determinants in rural china. *China Economic Review* 20: 635–49.
- Kreps, D. M. (1979). A representation theorem for preference for flexibility. *Econometrica* 47: 565–77.
- Kreps, D. and Porteus, E. (1978). Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* 46: 185–200.
- Krusell, P., Kuruscu, B., and Smith, A. A. (2010). Temptation and taxation. *Econometrica* 78(6): 2063–84.
- Kuhn, P., Kooreman, P., Soetevent, A., and Kapteyn, A. (2011). The effects of lottery prizes on winners and their neighbors: evidence from the Dutch postcode lottery. *American Economic Review* 101(5): 2226–47.
- Kunreuther, H., Ginsberg, R., Miller, L., Sagi, P., Slovic, P., Borkan, B., and Katz, N. (1978). *Disaster Insurance Protection: Public Policy Lessons*. New York: Wiley.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 112: 443–77.
- Laibson, D. (2001). A cue-theory of consumption. *Quarterly Journal of Economics* 116: 81–119.
- Layard, R. (2003). Happiness: Has Social Science a Clue? Lionel Robbins Memorial Lectures 2002/3. Lecture given at the London School of Economics, London, March 3–5.
- Layard, R. (2005). *Happiness: Lessons from a New Science*. New York: Penguin Press.
- Lipman, B. L. and Pesendorfer, W. (2013). Temptation. In D. Acemoglu, M. Arellano, and E. Dekel, (eds.), *Advances in Economics and Econometrics: Tenth World Congress, Volume 1*. Cambridge: Cambridge University Press.
- Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. *Economic Journal* 97: 666–84.
- Loewenstein, G. (1996). Out of control: visceral influences on behavior. *Organizational Behavior and Human Decision Processes* 65(3): 272–92.
- Loewenstein, G. (2000). Emotions in economic theory and economic behavior. *American Economic Review* 90(2): 426–32.
- Loewenstein, G. (2010). Insufficient emotion: soul-searching by a former inductor of strong emotions. *Emotion Review* 2(3): 234–9.
- Loewenstein, G. and O'Donoghue, T. (2004). Animal spirits: affective and deliberative processes in economic behavior. Unpublished Paper.
- Loewenstein, G. and O'Donoghue, T. (2007). The heat of the moment: modeling interactions between affect and deliberation. Unpublished paper.

- Loewenstein, G., O'Donoghue, T., and Rabin, M. (2003). Projection bias in predicting future utility. *Quarterly Journal of Economics* 118: 1209–48.
- Loewenstein, G., Weber, E., Hsee, C., and Welch, N. (2001). Risk as feelings. *Psychological Bulletin* 127: 267–86.
- Loomes, G. and Sugden, R. (1986). Disappointment and dynamic consistency in choice under uncertainty. *Review of Economic Studies* 53: 271–82.
- Luechinger, S., Meier, S., and Stutzer, A. (2010). Why does unemployment hurt the employed? Evidence from the life satisfaction gap between the public and the private sector. *Journal of Human Resources* 45(4): 998–1045.
- Lundborg P. and Lindgren, B. (2004). Do they know what they are doing? Risk perceptions and smoking behaviour among Swedish teenagers. *Journal of Risk and Uncertainty* 28: 261–86.
- Luttmer, E. (2005). Neighbors as negatives: relative earnings and wellbeing. *Quarterly Journal of Economics* 120: 963–1002.
- Massey, D. S. (2002). A brief history of human society: the origin and role of emotion in social life. *American Sociological Review* 67(1): 1–29.
- Meier, S. and Stutzer, A. (2008). Is volunteering rewarding in itself? *Economica* 75(297): 39–59.
- Meshulam M., Winter, E., Ben-Shakhar, G., and Aharon, I., (2012). Rational emotions. *Social Neuroscience* 7(1): 11–17.
- Metcalf, J. and Mischel, W. (1999). A hot/cool-system analysis of delay of gratification: dynamics of willpower. *Psychological Review* 106(1): 3–19.
- Mischel, W., Ayduk, O., and Mendoza-Denton, R. (2003). Sustaining delay of gratification over time: a hot-cool systems perspective. In G. Loewenstein, D. Read, and R. F. Baumeister (eds.), *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice*. New York: Russell Sage Foundation, pp. 175–200.
- Monat, A., Averill, J., and Lazarus, R. (1972). Anticipatory stress and coping reactions under various conditions of uncertainty. *Journal of Personality and Social Psychology* 24: 237–53.
- Myers, D. G. (2000). The funds, friends, and faiths of happy people. *American Psychologist* 55(1): 56–67.
- Noor, J. (2011). Temptation and revealed preference. *Econometrica* 79(2): 601–44.
- O'Donoghue, T. and Rabin, M. (1999). Doing it now or later. *American Economic Review* 89(1): 103–24.
- OECD (2001). *The Well-Being of Nations: The Role of Human and Social Capital*. Paris: OECD Centre for Educational Research and Innovation.
- Oswald, A. and Powdthavee, N. (2008). Does happiness adapt? A longitudinal study of disability with implications for economists and judges. *Journal of Public Economics* 92(5–6): 1061–77.
- Oswald, A., Proto, E., and Sgroi, D. (2015). Happiness and productivity. *Journal of Labor Economics* 33(4): 789–822.
- Oswald, A. and Wu, S. (2011). Well-being across America: evidence from a random sample of one million U.S. citizens. *Review of Economics and Statistics* 93(4): 1118–34.
- Ozdenoren, E., Salant, S., and Silverman, D. (2012). Willpower and the optimal control of visceral urges. *Journal of the European Economic Association* 10(2): 342–68.
- Pavlov, I. P. (1904). Sur la secretion psychique des glandes salivaires. *Archives Internationales de Physiologie* 1: 119–35.
- Phelps E. and Pollak, R. A. (1968). On second-best national saving and game-equilibrium growth. *Review of Economic Studies* 35: 185–99.
- Pollak, R. A. (1970). Habit formation and dynamic demand functions. *Journal of Political Economy* 78: 745–63.
- Preston, S. D. and de Waal, F. B. M. (2002). Empathy: its ultimate and proximate bases. *Behavioral and Brain Sciences* 25(1): 1–71.
- Read, D. and van Leeuwen, B. (1998). Predicting hunger: the effects of appetite and delay on choice. *Organization Behavior and Human Decision Processes* 76(2): 189–205.
- Redelmeier, D. and Kahneman, D. (1996). Patients' memories of painful medical treatments: real-time and retrospective evaluations of two minimally invasive procedures. *Pain* 116: 3–8.
- Rode, M., Knoll, B., and Pitlik, H. (2013). Economic freedom, democracy, and life satisfaction. In J. Gwartney, R. Lawson, and J. Hall, *Economic Freedom of the World: 2013 Annual Report*. Fraser Institute, pp. 215–33.
- Rolls, E. T. (2014). *Emotion and Decision Making Explained*. Oxford: Oxford University Press.

- Ruhm, C. J. (2000). Are recessions good for your health? *Quarterly Journal of Economics* 115(2): 617–50.
- Sacks, D. W., Stevenson, B., and Wolfers, J. (2013). Growth in subjective wellbeing and income over time. Unpublished.
- Sarracino, F. (2010). Social capital and subjective well-being trends: comparing 11 Western European countries. *Journal of Socio-Economics* 39(4): 482–517.
- Schneider, W. and Shiffrin, R. M. (1977a). Controlled and automatic human information processing: detection, search, and attention. *Psychological Review* 84: 1–66.
- Schneider, W. and Shiffrin, R. M. (1977b). Controlled and automatic human information processing: 2. Perceptual learning, automatic attending and a general theory. *Psychological Review* 84: 127–90.
- Sen, A. (1985). *Commodities and Capabilities*. Oxford: Oxford University Press.
- Sen, A. (1992). *Inequality Reexamined*. New York: Russell Sage Foundation; Cambridge, MA: Harvard University Press.
- Senik, C. (2014). Wealth and happiness. *Oxford Review of Economic Policy* 30(1): 92–108.
- Shefrin, H. and Thaler, R. (1988). The behavioral life-cycle hypothesis. *Economic Inquiry* 26(4): 609–43.
- Shiffman, S. and Waters, A. J. (2004). Negative affect and smoking lapses: a prospective analysis. *Journal of Consulting and Clinical Psychology* 72(2): 192–201.
- Shiv, B. and Fedorikhin, A. (1999). Heart and mind in conflict: the interplay of affect and cognition in consumer decision making. *Journal of Consumer Research* 26: 278–92.
- Sinaceur, M. and Tiedens, L. Z. (2006). Get mad and get more than even: when and why anger expression is effective in negotiations. *Journal of Experimental Social Psychology* 42: 314–22.
- Slovic, P., Finucane, M., Peters, E., and MacGregor, D. G. (2002). The affect heuristic. In T. Gilovich, D. Griffin, and D. Kahneman (eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment*. New York: Cambridge University Press, pp. 397–420.
- Smith, D., Loewenstein, G., Jepson, C., Jankovich, A., Feldman, H., and Ubel, P. (2008). Mispredicting and misremembering: patients with renal failure overestimate improvements in quality of life after a kidney transplant. *Health Psychology* 27(5): 653–58.
- Smith, K. and Dickhaut, J. (2005). Economics and emotion: institutions matter. *Games and Economic Behavior* 52: 316–35.
- Stevens, J. R., Hallinan, E. V., and Hauser, M. D. (2005). The ecology and evolution of patience in two new world monkeys. *Biology Letters* 1: 223–6.
- Stevenson, B. and Wolfers, J. (2008). Economic growth and subjective well-being: reassessing the Easterlin paradox. *Brookings Papers on Economic Activity* 1–87.
- Stevenson, B. and Wolfers, J. (2013). Subjective well-being and income: is there any evidence of satiation? *American Economic Review: Papers & Proceedings* 103(3): 598–604.
- Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies* 23(3): 165–80.
- Stutzer A. (2004). The role of income aspirations in individual happiness. *Journal of Economic Behaviour and Organization* 54(1): 89–109.
- Stutzer, A. and Lalive, R. (2004). The role of social work norms in job searching and subjective well-being. *Journal of the European Economic Association* 2(4): 696–719.
- Thaler, R. and Shefrin, H. (1981). An economic theory of self-control. *Journal of Political Economy* 89(2): 392–406.
- Tobin, H., Logue, A. W. Chelonis, J. J., and Ackerman, K. T. (1996). Self-control in the monkey macaca fascicularis. *Animal Learning and Behavior* 24(2): 168–74.
- Tooby, J. and Cosmides, L. (1992). The psychological foundations of culture. In J. Barkow, L. Cosmides, and J. Tooby (eds.), *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. New York: Oxford University Press.
- Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4): 297–323.
- Van Praag, B. M. S. (1968). *Individual Welfare Functions and Consumer Behavior: A Theory of Rational Irrationality*. Amsterdam: North-Holland.
- Van Praag, B. M. S. (1971). *Individual Welfare Functions and Consumer Behavior-A Theory of Rational Irrationality*. Dissertation, Contributions to Economic Analysis, No. 57. Amsterdam: North-Holland Publishing Company.

- Van Praag, B. M. S. and Ferrer-i-Carbonell, A. (2011). Happiness economics: a new road to measuring and comparing Happiness. *Foundations and Trends in Microeconomics* 6(1): 1–97.
- Varey, C. and Kahneman, D. (1992). Experiences extended across time: evaluation of moments and episodes. *Journal of Behavioral Decision Making* 5: 169–86.
- Veblen, T. ([1899] 1973). *The Theory of the Leisure Class: An Economic Study of Institutions. Introduction John Kenneth Galbraith*. Boston, MA: Houghton Mifflin.
- Veenhoven R. (1991). Is happiness relative? *Social Indicators Research* 24(1): 1–34.
- Veenhoven, R. (2000). Freedom and happiness: a comparative study in 44 nations in the early 1990's. In E. Diener and E. M. Suh (eds.) *Culture and Subjective Wellbeing*. Cambridge, MA: MIT Press, pp. 257–88.
- Veenhoven, R. and Vergunst, F. (2013). The Easterlin illusion: economic growth does go with greater happiness, Erasmus University Rotterdam, MPRA Paper No. 43983.
- Viscusi, W. K. (1990). Do smokers underestimate risks? *Journal of Political Economy* 98: 1253–69.
- Vohs, K. D. and Faber, R. J. (2007). Spent resources: self-regulatory resource availability affects impulse buying. *Journal of Consumer Research* 33: 537–47.
- Vohs, K. D. and Heatherton, T. F. (2000). Self-regulatory failure: a resource-depletion approach. *Psychological Science* 11(3): 249–54.
- Wälde, K. (2015). Stress and coping: An economic approach. Mimeo, Gutenberg University Mainz.
- Wilson, E. O. (2012). *The Social Conquest of Earth*. New York: Liveright Publishing Corporation.
- Winkelmann, L. and Winkelmann, R. (1998). Why are the unemployed so unhappy? Evidence from panel data. *Economica* 65: 1–15.
- Winter, E. (2014). *Feeling Smart: Why Our Emotions Are More Rational Than We Think*. New York: Public Affairs.
- Winter, E., Méndez-Naya, L., and García-Jurado, I. (2014). Equilibrium and strategic emotions. Mimeo.
- Wolfers, J. (2003). Is business cycle volatility costly? Evidence from surveys of subjective well-being, *International Finance* 6(1): 1–26.

PART 7

Bounded Rationality

INTRODUCTION TO PART 7

Neoclassical economics assumes fully rational, Bayesian, subjective expected utility maximizing decision makers who have limitless cognitive abilities and follow the rules of classical statistics. In Part 7 of the book, we consider the evidence, theoretical approaches, and applications of a competing view of human nature that is based on *bounded rationality*.

Chapter 19 mainly focuses on the *heuristics and biases program* of Daniel Kahneman and Amos Tversky. Heuristics are any decision rules used by individuals that are *fast* (in terms of the time required to take a decision) and *frugal* (in terms of the information requirements). Extensive empirical evidence shows that people do use simple heuristics and their behavior often systematically deviates from the classical rules of statistics. Such individuals may be said to be *boundedly rational*.

We describe a range of heuristics and the biases relative to the normative benchmarks. Many biases arise from the tendency to believe that small samples possess the statistical properties of large samples. This gives rise to the *law of small numbers*, which is the basis of the *representativeness heuristic*, the *gambler's fallacy*, and the *hot hands fallacy*.

Other heuristics do not necessarily invoke the law of small numbers. The *conjunction fallacy* arises from inadequate attention given to the set inclusion relation; the *availability heuristic* arises from drawing inferences based on readily available information; the *affect heuristic* arises from attention to the emotional dimension of a decision; the *anchoring heuristic* arises from tying one's inferences to anchors that are often irrelevant; *base rate underweighting* arises from giving inadequate attention to the base rate in Bayes' rule; *conservatism* arises from underweighting the likelihood of a sample; *hindsight bias* arises from discrepancies between predictive and postdictive guesses; *confirmation bias* arises from selective attention to events that is biased towards confirming one's initially held position; *false consensus* arises when people overestimate the extent to which others share their beliefs. Biases also arise from ignoring statistical phenomena such as *regression to the mean* and failing to distinguish between *necessary and sufficient conditions*.

The heuristics and biases program is one of the major achievements in social science. Some of the challenges to this program can be easily dismissed; these include subject errors and misperceptions or a misunderstanding on the part of Kahneman and Tversky of probability theory. Other challenges have sharpened the approach. For instance, presenting subjects with information in a frequency-format rather than a probability-format reduces some of the biases, yet the vast majority of subjects continue to exhibit these biases. Furthermore, a great deal of information given out by governments and firms is based on the probability-format. Experts are not exempt from the biases, although they are better at rationalizing their choices relative to non-experts.

We discuss two other approaches to bounded rationality. The first is Herbert Simon's approach that advocates *satisficing behavior* as a response to limited cognitive abilities and to the problem of infinite regress in search costs. In this approach, if the payoffs of individuals fall below their endogenous *aspiration levels*, they experiment with alternative strategies, otherwise they are passive. This is the sense in which they *satisfice*. The second approach, whose intellectual leader is Gerd Gigerenzer, compares the performance of fast and frugal heuristics against selected optimization techniques. Despite the perceived intellectual chasm between the heuristics and biases program and these approaches, we argue that the two approaches are complementary.

The heuristics and biases approach is not just a collection of empirical facts. A growing array of theoretical models is now able to organize the evidence on heuristics and biases. Prominent in this literature are models that formalize the law of small numbers and explain the gambler's and hot hands fallacies; reconcile base rate neglect with conservatism; model hindsight bias to explain financial volatility; explain overconfidence and confirmation bias; use aspiration adaptation to explain the results from well-known games; and explain limited attention and the existence of shrouded attributes in equilibrium.

We close the chapter by discussing the emerging topic of *limited attention*. Increasing evidence is being documented for limited attention, e.g., among the poor, taxpayers, eBay users, and used car buyers. Exciting theoretical models are also emerging to explain the empirical facts.

Chapter 20 considers *mental accounting*. Richard Thaler, who is a pioneer in this field, defines mental accounting as *the entire process of coding, categorizing, and evaluating events*. In neoclassical economics, all sources of income and expenditure are consolidated into a single financial account and money is completely fungible across different sources of income or expenditure. In actual practice, however, individuals form a *system of mental accounts* and the degree of fungibility of money across these accounts varies. For instance, individuals are reluctant to raid a children's educational account or consume out of home equity. In addition, people may have mental accounts for holidays, weekend drinks, book expenditures, and charitable contributions.

In life-cycle models, one implication of mental accounts and the lack of perfect fungibility is that individuals may have an account-specific marginal propensity to consume. Thus, the effects of defined pension plans, bonuses, dividends, and windfall gains depend on which mental account is used to debit or credit these income flows. A mental account that is in debt is said to be in *red*, otherwise it is said to be in *black*. Evidence indicates that individuals are particularly averse to mental accounts that are in the red; this is a form of *debt-aversion*. It can explain a range of phenomena, for instance, the *disposition effect*, namely, the tendency to hold on to loss making assets for too long and sell winners early.

Consumption that is paid for, say, a vacation, can be enjoyed as if it were free. On the other hand, the thought of future payments reduces the current flow of utility, such as the aversion to a ticking taxi meter. We incorporate these considerations in a theoretical model of mental accounting. The resulting insights can explain why we prefer to prepay for a vacation but are often willing to buy consumer durables on credit. Other phenomena that can be explained in this framework include the growing use of debit cards, popularity of fixed fee payments, and gift giving among couples with joint accounts.

Among other mental accounting phenomena that we consider are the following. Individuals often engage in *narrow bracketing* of problems, e.g., taxi drivers often target daily earnings, which prevents them from taking advantage of intertemporal labor-leisure trade-offs; *coherent arbitrariness* shows that initial preferences are malleable but gradually become coherent; the distinction between *transaction utility* and *acquisition utility*; real effects of *sunk costs*; dominance of *in-kind transfers* over *cash transfers*; the *diversification heuristic* and the formation of suboptimal portfolios; the *silver lining effect*; the *house money effect*; the *mental accounting of goals*; and the *behavioral economics of retirement savings* that relies on default options (opt-in or opt-out), initial contribution rates, and automatic escalation of contributions.

Chapter 21 considers a set of issues that would typically be considered in a course in *behavioral finance* and *behavioral corporate finance*. Financial markets have traditionally been held up by economists as a model of efficient markets. Thus, the demonstration of bounded rationality in financial markets is particularly important. Undergraduates in microeconomics courses are typically taken through the following set of arguments that justify an exclusive focus on fully rational

economic agents. (i) Irrationality, i.e., any departure from classical full rationality, should cancel out in the aggregate because positive errors are as likely to occur as negative errors. (ii) Irrational traders will eventually be out-competed by rational traders, so the former will exit the market. While plausible, the real test of these arguments is empirical, and they fail this test.

The *efficient markets hypothesis* (EMH) is the cornerstone of classical finance. It predicts that at any instant in time, prices of financial assets equal their fundamental values. However, the evidence decisively rejects EMH. For instance, stale information affects current prices; public information is not readily incorporated into stock prices; the flow of information in financial markets is not instantaneous; there can be huge movements in stock prices despite no discernible news; and prices of substitute securities in Siamese twin companies typically differ from the predicted ratio.

The evidence contradicting EMH has also led to a better understanding of the behavior of financial market participants. In turn, this has led to the development of an impressive body of theoretical literature that demonstrates why financial markets need not be efficient. Indeed, *noise traders*, who trade on *investor sentiment* can outperform rational arbitrageurs. *Professional arbitrage* is also subject to biases, and often operates under the constraint of *performance based arbitrage* that prevents it from equilibrating markets. *Gradual flow of information* and *investor heterogeneity* have implications for the *momentum* in asset prices as well as *trading volume*. Managerial biases such as *overconfidence* and *attribution bias* also lead to inefficient outcomes. These predictions are borne by the evidence.

The theoretical and empirical results in Chapter 21 make a powerful case for reconsidering the classical benchmarks in finance. We close with the following quote from Thaler (1999).

Behavioral finance is no longer as controversial a subject as it once was. As financial economists become accustomed to thinking about the role of human behavior in driving stock prices, people will look back at the articles published in the past 15 years and wonder what the fuss was about. I predict that in the not-too-distant future, the term “behavioral finance” will be correctly viewed as a redundant phrase. What other kind of finance is there? In their enlightenment, economists will routinely incorporate as much “behavior” into their models as they observe in the real world. After all, to do otherwise would be irrational.

CHAPTER 19

Judgment Heuristics

19.1 Introduction

The modern *heuristics and biases* research program commenced with the seminal articles by Tversky and Kahneman (1971, 1974). A central insight of this program is that individuals are not fully rational Bayesian decision makers, but their behavior is not random either. Individuals often ignore a substantial part of the available information or use it selectively to implement simple *rules of thumb* or *heuristics*. Most heuristics are *fast*, in terms of computation time required, and *frugal* in the use of information. These heuristics do not optimize in the classical sense, hence, their performance is not necessarily optimal. However, when there is uncertainty about the optimal solution, their performance can be at least as good or even better than some more complex methods (Gigerenzer et al., 1999).

To motivate their approach, Tversky and Kahneman (1974) give the following example. Individuals often judge distance by the heuristic of how clear it is; sharper objects are thought to be closer. This heuristic is fast and frugal. While this heuristic often works well, it might also sometimes give rise to misleading perceptions. For instance, when atmospheric conditions are hazy, distances might be overestimated. Conversely, when conditions are very clear, distances might be underestimated; for a revised account of this example, see Kahneman and Frederick (2005).

The heuristics and biases approach often considers problems that require some probabilistic assessments or inferences. Classical economics assumes that individuals respect the classical laws of probability and statistical inference. In contrast, the evidence indicates that in these situations, individuals employ a range of *judgment heuristics*. The heuristics and biases approach demonstrates a systematic departure of the actual behavior of individuals from the predictions of the laws of classical statistics.¹ It is notable that even statistically sophisticated researchers and experts typically rely on these judgment heuristics.² The heuristics and biases approach is quite distinct from the formal decision theory approach in Part 1 of the book. It does not necessarily rely on any explicit optimization method, although most formal applications use a mixture of optimization methods in conjunction with simple rules of thumb.

We begin with the *representativeness heuristic* in Section 19.2 (Tversky and Kahneman, 1974). When individuals use this heuristic, they compare the likeness of a sample, even if the sample is

¹ There are many excellent surveys and collected readings. Particularly relevant to economists are Camerer (1995), Gilovich et al. (2002), Kahneman et al. (1982), Kahneman (2003), and Kahneman (2011). For a commendable effort at using some of this machinery to analyze development issues, see WDR (2015).

² See for instance, Gilovich et al. (2002), Kahneman (2003), Kahneman (2011), and Tetlock (2006).

small, with known features of the population distribution. The evidence shows that relative to a Bayesian, users of the representativeness heuristic assign too high a probability that small samples share the properties of large samples (Camerer 1987, 1990, 1995). Such individuals behave as if they obeyed a *law of small numbers*, rather than the statistically correct, *law of large numbers*.

Reliance on the law of small numbers has several implications. Individuals find it difficult to produce sequences of outcomes of random variables, and their constructed sequences show up too much *negative autocorrelation*, which is also sometimes known as *alternation bias* or *negative recency* (Tversky and Kahneman, 1974; Bar-Hillel and Wagenaar, 1991; Rapoport and Budescu, 1997). A belief in negative autocorrelation, when there is none in the underlying distribution, in betting domains is known as the *gambler's fallacy*. For instance, gamblers are often reluctant to bet on the immediately previous winning numbers in a lottery, yet the conditional and unconditional mathematical probabilities of these numbers is unchanged. Evidence from betting behavior gives strong support to the gambler's fallacy (Metzger, 1984; Clotfelter and Cook, 1993; Terrell, 1994; Croson and Sundali, 2005), as does the evidence from the behavior of investors in stock markets (Odean, 1998).

The converse of the gambler's fallacy, i.e., a belief in *positive autocorrelation* in outcomes, when there is none, is known as the *hot hands fallacy*. When subjects observe too many consecutive outcomes of one kind, they come to believe that the realization of a similar outcome is more probable. For instance, basketball players who score heavily in a game are often perceived to be on a hot streak even when there is no evidence of underlying hot hands (Gilovich et al., 1985; Tversky and Gilovich, 1989a,b; Camerer, 1989). Similar beliefs are often assigned to winning football teams. Recent evidence indicates that for basketball and baseball, there might be an underlying hot hands phenomenon (Green and Zwiebel, 2013; Bocskocsky et al., 2014; Miller and Sanjurjo, 2015), yet we argue that it does not invalidate the hot hands fallacy.

In contrast to outcomes from sports, there is good support for the hot hands fallacy when the underlying stochastic process is inanimate. For instance, lotto stores that sell winning lotto tickets experience unusually high sales for up to 40 weeks following the lotto win (Guryan and Kearney, 2008). Individuals are willing to pay more for experts who are believed to possess hot hands (Powdthavee and Riyanto, 2015). Players who win in casinos are less likely to quit, believing that they are on a winning streak (Croson and Sundali, 2005). When the underlying stochastic process relies on human skills, one is more likely to observe a gambler's fallacy and when the process is inanimate, a hot hands fallacy is more likely to be observed (Ayton and Fischer, 2004).

The gambler's fallacy and the hot hands fallacy are both likely to be caused by the representativeness heuristic. An individual who initially believes the underlying process to be random (so zero autocorrelation in successive draws) predicts too much negative autocorrelation to try to make outcomes appear random (gambler's fallacy). However, when too many outcomes of the same type are observed, the individual starts to believe that the underlying process is positively autocorrelated, even when small samples do not statistically justify this conclusion (hot hands fallacy).

Finally, we give a useful formalization of the law of small numbers due to Rabin (2002). This model also explains the existence of *overconfidence* and the gambler's fallacy.

Section 19.3 considers the *conjunction fallacy*, namely, the tendency to assign higher probabilities to the intersection of two sets relative to any of the two sets alone. The most well known and critically discussed illustration of this phenomenon is the *Linda problem* (Tversky and Kahneman, 1983). After describing a set of features of a hypothetical female called Linda (e.g., concerned with social justice, participated in anti-nuclear demonstrations) subjects in experiments are asked to judge Linda's profession. A large number of subjects chose the composite description

“Linda is a bank teller and active in the feminist movement” relative to either of the two descriptions separately. Since the probability of the intersection cannot ever be larger than any of the individual sets, this is termed as the conjunction fallacy.

Hertwig and Gigerenzer (1999), two long-standing critics of the Kahneman–Tversky approach, claimed that the conjunction fallacy disappears if it is posed in terms of natural frequencies. However, in a between-subjects design, when the Linda problem was presented in a natural frequency format, the conjunction fallacy survived (Kahneman and Tversky, 1996). The conjunction fallacy was shown to be reduced in the presence of incentives and joint decisions made by groups of two and three individuals (Charness et al. 2010). However, other experimental evidence shows that the conjunction fallacy survives intact not only in the presence of monetary incentives but also when the same problem is framed in alternative ways (Bonini et al., 2004; Stolarz-Fantino et al., 2003). Another potential explanation of the conjunction fallacy is given in terms of the potential confound between the word “and” and the logical operator $\wedge$ (Camerer, 1995; Hertwig et al., 2008). However, other research shows that while the conjunction fallacy may be lessened slightly due to these reasons, it cannot be eliminated (Tentori et al., 2004; Tentori and Crupi, 2012).

Section 19.4 considers the *availability heuristic*. Individuals judge the probability of an event by the ease with which similar events are accessible or “available” to the mind (Tversky and Kahneman, 1974). Availability may depend on how salient the event is, or how vivid or lasting an impression it made, or even how easy it is to construct hypothetical scenarios. Media coverage is one important determinant of the availability of events (Lichtenstein et al., 1978; Eisenman, 1993). An influential early empirical study established a firm link between the perceived frequency of deaths by various causes and proxies for the availability of such events in memory (Lichtenstein et al., 1978). These results have been replicated and extended (Hertwig et al., 2005; Pachur et al., 2012). Indeed, there is a close relation between the availability heuristic and the *affect heuristic* that we consider later (Pachur et al., 2012).

Section 19.5 considers the *affect heuristic* and highlights the link between emotions and heuristics (Zajonc, 1980; Slovic et al., 2002; Loewenstein et al., 2001). The basic idea is that we classify and tag events and memories by their *affect*, for instance, positive affect (enjoyable holiday) or negative affect (painful skiing accident). In making judgments and decisions, we are then influenced by the pool of these positive and negative affects or the *affect-pool*. For instance, in buying a car, we might focus a great deal on the affect that the car has on us, for instance, a sleek car, a cool car, or an enviable car. The affect heuristic substitutes a hard question (what do I think about it?) by an easier question (what do I feel about it?) that allows for quicker decision making (Kahneman, 2011).

Evidence shows that our response to various hazards, such as smoking and nuclear power, depends on the affect, or *dread* in this case, that they invoke (Fischhoff et al., 1978; Finucane et al., 2000). Experts also exhibit the affect heuristic; the perceived riskiness of various hazardous substances is closely related to the affect that they have on experts (Slovic et al., 1999). Presenting the data in terms of natural frequencies to experts enhances the affect of an event (Slovic et al., 2000).

Anchoring, considered in Section 19.6, is one of the most robust and important heuristics that people use. Anchoring occurs when people are influenced in their judgments and decisions by an initial suggestion, or *anchor*, that may be irrelevant to the question asked. Anchors may be self-generated or provided by the experimenter. Anchoring has been observed in a wide variety of phenomena. Furthermore, once individuals fix an anchor in their minds, they adjust too slowly, and insufficiently, towards the target; this is particularly the case with self-generated anchors (Epley and Gilovich, 2001). However, when the anchors are provided by the experimenter, then

individuals typically engage System 1 of the brain (see Section 19.11) in gauging the accuracy of the anchor. System 1 is biased towards finding confirmatory evidence, which leads to anchoring.

Early studies on anchoring considered the relation between stimulus and sensation. However, Tversky and Kahneman (1974) revolutionized the field by focusing on numeric anchors, which is now the established method of eliciting anchoring effects. In a famous experiment, they rigged a wheel of fortune with numbers 1–100 to come up with either 10 or 65 (low and high anchors). They first asked subjects if the number of African nations in the United Nations was larger or smaller than the anchor (comparative judgment task). Then they asked subjects to guess the number of African nations in the United Nations (absolute judgment task). Although, in this case, the anchors are irrelevant to the questions, they had a significant effect on the subjects' guesses.

Experienced judges are also subject to the anchoring effect (Englich and Mussweiler, 2001). Anchors that are clearly implausible also have a significant effect on subjects' choices (Strack and Mussweiler, 1997; Mussweiler and Strack, 2000a). Field data show that estate agents are unduly influenced by the initially suggested list price for a house (Northcraft and Neale, 1987). Caps on damages awarded to litigants serve as an anchor and have an effect on the rate of pretrial settlements (Babcock and Pogarsky, 1999) as well as on the level of damages (Pogarsky and Babcock, 2001).

Anchoring is one of the most robust heuristics and has been documented in a wide range of contexts such as poetry reading sessions (Ariely et al., 2006); violation of a basic implication of efficient markets (Mussweiler and Schneller, 2003); and first offers in price negotiation (Galin-sky and Mussweiler, 2001). Anchoring has also been used to explain a range of other behavioral phenomena such as hindsight bias, non-linear probability weighting, and preference reversals.

Section 19.7 shows that the behavior of individuals is inconsistent with Bayes' rule, particularly when they face probabilistic information. In particular, individuals engage in *base rate neglect*, which is the tendency to give insufficient attention to the prevalence of a particular trait or feature in the overall population. Consider the well-known cab problem (Kahneman and Tversky, 1972; Bar-Hillel, 1980; Tversky and Kahneman, 1980). Cabs in a city are either green (85%) or blue (15%). One of the cabs was involved in an accident last night. A witness with a reliability rate of 80% came forward to identify that the cab involved in the accident was blue. What is the probability that a blue cab was involved? The correct answer, using Bayes' rule, is 41.4%, while the modal experimental response was 80%. Subjects appear to ignore the fact that only 15% of the cabs are blue (the base rate); thus, they engage in base rate neglect.

The base rate underweighting phenomenon is robust to several perturbations, although some might lessen it. For instance, base rate underweighting is found even when base rates are made more salient (Ajzen, 1977; Bar-Hillel, 1980), or when subjects are highly educated (Casscells et al., 1978), or when individuals have higher measured intelligence (Stanovich and West, 2000). A much better predictor of base rate underweighting is the important distinction between *causal* and *incidental* base rates (Ajzen, 1977). An example of a causal base rate is the instruction: "there are an equal number of Blue and Green taxis but 85% of the taxis involved in accidents are Green." In this case, base rate neglect is reduced (Kahneman, 2011). Furthermore, when information on base rates is combined with other non-causal information, then base rate neglect may also be reduced (Bar-Hillel, 1980).

As with many other heuristics, the claim is made that presenting the information in a *frequency format* rather than a *probability format* reduces base rate neglect (Brase, 2002b; Cosmides and Tooby, 1996; Gigerenzer and Hoffrage, 1995; Pinker, 1997). For most empirical studies that use the frequency format, the percentage of subjects who do not engage in base rate underweighting is about 40% (Barbey and Sloman, 2007). While this is an improvement over the probability

format, the case for the efficacy of frequency format is overstated. Even when data are given in natural frequencies, but the set inclusion relation is not made obvious, the benefits of natural frequency representation drop significantly (Giroto and Gonzalez, 2001). An even more practical consideration in the debate over natural frequencies versus the probability format, based on anecdotal evidence, is that most real-world economic information of interest is arguably presented in a probability format.

In contrast to base rate underweighting, the phenomenon of *conservatism* refers to *underweighting of the likelihood of a sample* (Camerer, 1995). In Section 19.7.3, we offer a reconciliation of base rate neglect with conservatism by distinguishing between the *strength* and the *weight* of *evidence* (Griffin and Tversky, 1992). For instance, in a reference letter, strength refers to the warmth and support for the applicant, while weight refers to the standing and credibility of the referee. Individuals typically give relatively greater prominence to strength. This is a form of overconfidence or base rate neglect. However, when the strength of the letter is low relative to the weight one might feel underconfident, which is a form of conservatism.

In Section 19.8, we consider *hindsight bias*. This occurs when individuals report, ex-post, that actual events coincided with their initial predictions of the events, i.e., *they knew it all along*. Hindsight bias occurs in a between-subjects (Fischhoff, 1975) and a within-subjects design (Fischhoff and Beyth, 1975). Hindsight bias differs from learning. Under learning, one learns to make better judgments about the future. Under hindsight bias, one misperceives one's own predictions made in the past, in a manner that gives rise to a lower predictive error. The main explanation of hindsight bias is that the human mind has only partial ability to recall what it used to believe in the past, particularly when one changes one's mind (Kahneman, 2011). Hindsight bias might also arise as a coping mechanism to make sense of an unpredictable world, or to signal one's competence to others (Rachlinski, 1998).

Hindsight bias is particularly important in legal judgments because jurors and judges need to determine the culpability of a defendant's action when they already know it has caused harm (Rachlinski, 1998). Hindsight bias has been found in very diverse contexts. A meta-study shows that experts are subjected to the hindsight bias (Guilbault et al., 2004). In Section 19.8.4, we apply hindsight bias to a problem in finance (Biais and Weber, 2009). We show that under hindsight bias, decision makers reduce estimates of volatility and engage in less corrective activity of incorrect models. Empirical evidence bears out the predictions of the model.

Section 19.9 considers *confirmation bias* in which individuals engage in selective interpretation of evidence so that it supports their initial beliefs. This is sometimes known as *biased assimilation* in psychology (Nickerson, 1998). There is extensive evidence of confirmation bias (Lord et al., 1979; Darley and Groos, 1983; Munro et al., 2002). Confirmation bias is not mitigated by greater cognitive abilities (Stanovich and West, 2007, 2008a,b). We also sketch a formal model of confirmation bias that can also explain the prevalence of overconfidence and underconfidence (Rabin and Schrag, 1999).

Section 19.10 considers a range of other judgment heuristics that include the following.³ People often do not take account of the phenomenon of *regression to the mean* (Galton, 1886; Tversky and

³ It is clearly not possible to consider all possible heuristics. For instance, we omit a discussion of the heuristics used to form macroeconomic expectations. Rational expectations still lie at the heart of modern macroeconomic theory and macroeconomic forecasting. Fuster et al. (2010) review evidence that is not consistent with rational expectations and outline an alternative that they call *natural expectations*. Beshears et al. (2013b) show that individuals often find it quite hard to recognize the dynamics of mean-reverting processes, an ability that is essential for existing macroeconomic models of forecasting. Ericson et al. (2015) propose a heuristic approach to choosing among time–outcome pairs that outperforms exponential and hyperbolic discounting. These approaches are likely to gain increasing importance in the future, possibly with profound effects on macroeconomics.

Kahneman, 1974). If outcomes are influenced by luck, then better (worse) than average outcomes in the current period should be expected to revert to more modest (better) levels in the future. For instance, on average, golfers with high (low) initial scores on day 1 of a tournament report lower (higher) scores on day 2 (Kahneman, 2011). People are often not very good at distinguishing between *necessary and sufficient conditions*, a distinction that is critical in many economic models (Wason, 1968). *Attribute substitution* occurs when one attribute of an object is substituted by another (Kahneman and Frederick, 2005). For instance, the answers to the two questions: “How happy are you with your life in general?” and “How many dates did you have last month?” depend on the order in which the questions are asked (Strack et al., 1988).

Section 19.11 focuses on one potential explanation of judgment heuristics. It is useful to think of the human brain as made up of two broad, but purely hypothetical, systems; System 1 and System 2 (Stanovich and West, 2000, 2002). System 1, or the emotional system, is fast, automatic, reactive, and requires little mental effort; it is also never turned off. System 2, or the cognitive system, is slow, deliberative, uses scarce cognitive resources, and takes account of the long-run consequences of our actions. Kahneman (2011) offers a powerful and cogent case for explaining judgment heuristics using this two systems approach.

Since System 1 is automatic, most judgment heuristics arise from its response. Even when an unusual situation is encountered and the services of System 2 are called upon, the agenda is often set by System 1, which controls the events that are recalled. System 1 calls the shots most of the time. This is revealed by the biases exhibited by trained statisticians in dealing with questions on probability (Tversky and Kahneman, 1973), and the extensive evidence from other expert judgments that is summarized in Section 19.18. System 1 is relatively more likely to be engaged when decisions have to be made quickly, statistical training is low, and cognitive ability is lower (Kahneman and Frederick, 2005; Oechssler et al., 2009).

In Section 19.12, we use the model of *thinking and persuasion* as a possible formalization of judgment heuristics (Shleifer et al., 2008). The idea is that, on account of cognitive limitations, individuals group different events into categories (Mullainathan, 2002). They then treat all events in the same category in an identical manner, so that there is across-category heterogeneity and within-category homogeneity; such thinkers may be termed as *coarse thinkers*. We consider the possibility that there is a persuader who controls the message received by a coarse thinker. It is never optimal for a persuader to send a misleading message to a classical Bayesian thinker but it might be optimal to mislead a coarse thinker. This model can address several puzzling phenomena such as advertising and different forms of persuasion by politicians, the media, firms, and experts. Empirical evidence supports the existence of persuasions from many domains (DellaVigna and Gentzkow, 2010).

In Section 19.13, we consider *mental models* that people form about the world around them and their place in it. These models can take several forms such as beliefs about causal relations, social identities, categorizing disparate information into coarse categories, and social worldviews (Denzau and North, 1994; WDR, 2015). Mental models are pervasive, help economize on cognitive costs, are transmitted from generation to generation, once formed they are typically inertial, and can both assist and deter optimal decisions, depending on their accuracy (Datta and Mullainathan, 2014). Humans are hardwired with some mental models, while in other cases, history plays an important role in their formation (Nunn, 2008; Alesina et al., 2013; Acemoglu et al., 2013). We also consider attempts to transform incorrect mental models into more useful ones (Datta and Mullainathan, 2014).

Section 19.14 discusses Herbert Simon’s seminal approach to bounded rationality (Simon, 1978). The idea is that in the presence of cognitive limitations, individuals should not be expected

to engage in the optimization exercises that neoclassical economic theory assumes to be the default (*substantive rationality*). As Simon (1978) famously noted: “But there are no direct observations that individuals or firms do actually equate marginal costs and revenues.” Thus, he recommends concentrating on *procedural rationality*, which focuses on the quality of the process of decisions. As Reinhard Selten has argued on many occasions, the problem is that economic problems are typically hard in the sense that they are *NP-Complete*, a term that is borrowed from the computer science literature.⁴ In an NP-Complete problem, the number of steps needed to solve a problem in any algorithm of the solution, grows exponentially with the size of the problem (e.g., the travelling salesman problem⁵). So, even for moderately difficult problems of this sort, the computing time required to solve the problem is very high. It then becomes a leap of faith that Homo economicus can solve these problems or act as if he/she can.

An example of a process that satisfies procedural rationality is *satisficing behavior*, where individuals have an *aspiration level* and adjust gradually in its direction through a sequence of steps (Simon, 1955, 1978, 2000; Selten, 2001). The word “satisficing” is a neologism that alludes to the fact that such decision procedures are satisfactory and they suffice. Empirical evidence is supportive of the theory (Caplin et al., 2011); subjects appear to search sequentially and stop searching when their aspiration level is achieved.

We demonstrate how the cooperative outcome can be achieved in the prisoner’s dilemma game as one’s aspirations adjust over time (Karandikar et al., 1998; Cho and Matsui, 2005; Börgers and Sarin, 2000). To understand the basic idea, suppose that in the prisoner’s dilemma game, players have an aspiration level that lies in between the payoffs from pure defection and pure cooperation. Then, whenever one or both players defect, one of the players finds that the payoff is below the aspiration level. Such a player then experiments by choosing another strategy. These experiments eventually lead to the cooperative outcome where all experimentation ceases.

Section 19.14.2 considers the approach of Gerd Gigerenzer and colleagues on fast and frugal heuristics. The focus of this work is to try to establish the following (Gigerenzer et al., 1999; Gigerenzer and Selten 2001): Heuristics are not mistakes, but a rational response to bounded rationality, and heuristics often outperform selected optimization methods. In some cases these heuristics are derived from evolutionary adaptation. For instance, individuals use the *gaze heuristic*, to catch a ball by employing a constant angle between the ball and the gaze. However, there should be no presumption that evolution has prepared us with the relevant heuristics needed to optimally invest in alternative pension plans or choose among alternative stocks.

We consider a range of fast and frugal heuristics, but mainly the *recognition heuristic*, and the *take-the-best* heuristic. Using the recognition heuristic, between two stocks, stock market participants invest in the stock that they recognize, relative to the one that they do not. There is some evidence that this method outperforms several investment strategies, such as purely random selection of stocks, expert judgment, and a selection of statistical selection techniques (Borges et al., 1999; Czerlinski et al., 1999). The recognition heuristic has also been used to determine which city has a higher population (Gigerenzer and Goldstein, 1996). However, doubts have

⁴ See, for instance, a video of the entertaining panel discussion on behavioral economics at the 2011 Lindau Nobel Laureate Meeting in Economic Sciences. The participants were George A. Akerlof, Robert J. Aumann, Eric S. Maskin, Daniel L. McFadden, Edmund S. Phelps, and Reinhard Selten.

⁵ Suppose that a salesman is provided with a list of cities and the distances between each pair of cities. The objective, starting with a home city, is to find the shortest possible route such that the salesman visits each city exactly once and returns to the home city.

also been raised about the robustness of these results, and contrary evidence found (Oppenheimer, 2003).

There are several concerns about this literature. First, it demonstrates the superiority of fast and frugal heuristics in domains where the relevant optimization benchmark is typically unclear. For instance, in predicting stock market prices, it is well known that the efficient markets hypothesis fails, so what is the relevant optimal solution to compare with the outcome arising from the heuristics? Furthermore, these heuristics are experimenter-provided. No compelling evidence is provided whether people do actually use these heuristics. Furthermore, in many complicated economic problems, it is not clear what these heuristics should be. In contrast, the Kahneman–Tversky approach addresses squarely the issue of which heuristics are used in different environments and the relevant optimization benchmark is unambiguous.

Section 19.15 considers what has come to be known as the *great rationality debate* in psychology, whose main protagonists are viewed as Kahneman–Tversky on the one hand, and Gerd Gigerenzer/colleagues on the other (Stanovich and West, 2000; Stanovich, 2012). The Kahneman–Tversky approach came under intense attack in the years following the publication of their work and gave rise to the rationality debate (Cohen, 1981; Einhorn and Hogarth, 1981; Lopes, 1991; Gigerenzer, 1991, 1993, 1994; Gigerenzer and Hoffrage, 1995).

This debate is often misplaced and highly muddled for, at least, the following three reasons.⁶

First, there is a basic misunderstanding of the positions of the protagonists. Kahneman–Tversky show that individuals fall short of the unambiguous statistical benchmarks required in the rational actor model in economics. Gigerenzer–colleagues show that in some domains, judgment heuristics perform quite satisfactorily against selected statistical benchmarks. Of course, both Kahneman–Tversky and Gigerenzer–colleagues may be right, and probably are right.

Second, it has been argued that data presented in terms of natural frequencies as compared to probabilities reduces or eliminates biases relative to the best statistical benchmark (Gigerenzer et al., 1988; Gigerenzer, 1991, 1993, 1994; Gigerenzer and Hoffrage, 1995). But this position is vastly overstated (Kahneman and Tversky, 1996). Even in the best case scenarios, 60% or more of the subjects continue to exhibit base rate neglect, across studies, in the presence of frequency information (Barbey and Sloman, 2007). Even for the conjunction fallacy that has been held up as the leading example of the facilitating role of the frequency format, a between-subjects design reveals a high degree of conjunction bias (Kahneman and Tversky, 1996). There are many heuristics such as the anchoring heuristic that are hardly influenced by the choice between frequencies and probabilities.

Third, the high degree of heterogeneity among subjects who exhibit biases contradicts conformity with a single rational response (Stanovich and West, 2000). Finally, there are a large number of other misconceptions about the Kahneman–Tversky position relating to context, framing effects, errors, subject misperceptions, and evolutionary adaptation that have already been well addressed and do not have a firm basis (Kahneman and Tversky, 1996; Stanovich and West, 2000; Kahneman, 2011). Kahneman and Tversky's (1996, p. 584) frustration is well captured in this quote: "The position described by Gigerenzer is indeed easy to refute but it bears little resemblance to ours. It is useful to remember that the refutation of a caricature can be no more than a caricature of a refutation."

⁶ We omit here the debate over which mental model (e.g., the dual systems model) best provides foundations for heuristics and biases (Stanovich and West, 2000). These issues are currently more central to psychologists than to economists.

Section 19.16 considers the rationale and implications of attributes of goods/services that might be *shrouded*. Attributes are said to be *shrouded* if they are partly hidden, but can be discovered by individuals at some cost. For instance, while opening a bank account, the minimum balance fees, fees for bounced checks, or other bank surcharges, might not have been very visible in the information provided. There are no advantages to firms from shrouding information in the classical case with fully rational consumers and competitive firms (Jovanovic, 1982; Milgrom, 1981). However, one does observe many examples of shrouded attributes. For instance, charging individuals high shipping costs (shrouded attribute) but low base prices (unshrouded attribute) on eBay increases sales and revenues (Hossain and Morgan, 2006). We follow up the empirical evidence with a formal model of shrouding of attributes (Gabaix and Laibson, 2006). In the presence of myopic consumers who pay insufficient attention to shrouded attributes, it may be optimal for firms to shroud information.

Section 19.17 considers the implications of *limited attention* that reflects limited cognitive abilities. Classical economics assumes full attention, which is robustly rejected by the evidence. The problem of limited attention is relatively more serious for the poor whose attention is diverted by basic problems of food, shelter, and clean water (Datta and Mullainathan, 2014; WDR, 2015). This diversion of attention is a form of *cognitive tax* (Shah et al., 2012; Mullainathan and Shafir, 2013). It reduces the quality of attention towards other important decisions; this is sometimes known as the problem of *limited bandwidth* (Banerjee and Mullainathan, 2008). The WDR (2015) is an excellent introduction to these issues.

In other contexts, consumers pay less attention to taxes that are not listed on price stickers but added at cash registers; displaying information about taxes on price stickers reduces sales (Chetty et al., 2009). Individuals pay limited attention to the rules of earned income tax credit (Hotz and Scholz, 2003). People pay less attention to electronic toll collections, as compared to manual toll collections, allowing operators to levy higher tolls (Finkelstein, 2009). Car sales in used markets are subjected to a left digit bias in odometer readings. Thus, car sales and car prices are sensitive to integer multiples of the 10,000 mile threshold (Lacetera et al., 2012; Englmaier et al., 2014). We discuss the recent work on the *sparse max operator* that formalizes a way of providing differential attention to various attributes of a good (Gabaix, 2014). We close with brief comments on the rational inattention models (Sims, 1998, 2003).

Section 19.18 collects a range of empirical results to show that experts exhibit biases in the same manner as non-experts (Tetlock, 2002; Ben-David et al., 2013; WDR, 2015). Hence, market experience does not eliminate non-standard behaviors.

19.2 The law of small numbers

In statistics, the *weak law of large numbers* is often expressed as follows.

Theorem 19.1 (*Weak Law of Large numbers; Hogg et al., 2005*) Let $\{X_n\}$ be a sequence of independently and identically distributed random variables, with a common mean μ and finite variance σ^2 . Denote the sample mean based on n observations by $\bar{X}(n) = \sum_{i=1}^n X_i/n$. Then $\bar{X}(n)$ “converges in probability” to the population mean, μ , i.e., $\lim_{n \rightarrow \infty} P[|\bar{X}(n) - \mu| \geq \varepsilon] = 0 \forall \varepsilon > 0$.

The weak law of large numbers can be strengthened to the *strong law of large numbers*, which guarantees almost sure convergence of the sample mean to the population mean (Chung, 1968). Notice that the law of large numbers is a limiting result and holds only for very large sample sizes.

By contrast, individuals are said to subscribe to the *law of small numbers* if they believe that for any sample size, the sample mean is identical to the population mean. Such people believe that the sample proportions mimic the population proportions for any sample size.

Consider the following experiment reported in Tversky and Kahneman (1974). A town is served by two hospitals. In the larger hospital, about 45 babies are born each day, and in the smaller hospital, about 15 babies are born each day. As you know, about 50% of all babies are boys. However, the exact percentage varies from day to day. For a period of 1 year, each hospital recorded the days on which more than 60% of the babies born were boys. Which hospital do you think recorded more such days?

In the sample 53 students said that both hospitals are equally likely to have recorded such days and an equal split of 21 students chose the larger and the smaller hospital, respectively. The correct answer is the smaller hospital. The reason is that the larger hospital has a larger sample size, so it is likely to better reflect the population proportion of 50%.

19.2.1 Representativeness heuristic

Tversky and Kahneman (1974) explain the *representativeness heuristic* as follows. “What is the probability that object A belongs to class B? What is the probability that event A originates from process B? . . . In answering such questions, people typically rely on the representativeness heuristic, in which probabilities are evaluated by the degree to which A is representative of B, that is, by the degree to which A resembles B.”

In one experiment, they give a description of a hypothetical individual named Steve (p. 1124): “Steve is very shy and withdrawn, invariably helpful, but with little interest in people, or in the world of reality. A meek and tidy soul, he has a need for order and structure, and a passion for detail.” Subjects in experiments were then asked to judge the probability of various professions that Steve might be engaged in, such as farmer, physician, librarian, salesman, and airline pilot. Subjects in experiments typically judged Steve’s profession by matching his description with their stereotype of that profession (representativeness). However, people do not take account of the fact that there are many more farmers than librarians; this is base rate neglect that is considered in greater detail, below.

Consider a more formal example of the representativeness heuristic (Camerer, 1987, 1990, 1995). Suppose that there are two urns, $i = A, B$, and one is chosen randomly by nature. Urn A has 1 red and 2 black balls. Urn B has 2 red and 1 black ball. It is common knowledge that nature chooses urn A with probability 0.6, so $P(A) = 0.6$ and $P(B) = 0.4$. A sequence of three balls is drawn *with replacement* from one of the urns. Subjects in experiments do not know which urn the balls are drawn from.

Let x denote the number of red balls that come up in the sample of three balls, $x = 0, 1, 2, 3$. Suppose that the sample, based on three draws, turns out to be $x = 1$. What is the posterior probability that the sample came from urn A? Using Bayes’ rule, this can be computed as follows.

$$P(A | x = 1) = \frac{P(x = 1 | A) P(A)}{P(x = 1 | A) P(A) + P(x = 1 | B) P(B)}. \quad (19.1)$$

In a single draw of a ball from an urn, there are two possibilities (red ball or black ball), hence, this is a Bernoulli experiment. Let p_A, p_B be the respective probabilities of drawing a red ball from urns A and B. Then, the thrice repeated Bernoulli experiment has a binomial distribution given by

$$P(x | i) = \binom{3}{x} p_i^x (1 - p_i)^{3-x}; i = A, B \text{ and } x = 0, 1, 2, 3. \quad (19.2)$$

A simple calculation shows that $P(x = 1 | A) = 4/9$ and $P(x = 1 | B) = 2/9$. Thus, from (19.1), we get the Bayesian estimate, $P(A | x = 1) = 0.75$. In contrast to the Bayesian estimate, an individual who uses the representativeness heuristic draws one of the following inferences.

- 1a. $P(A | x = 1) = 1$. In this interpretation, the (small) sample has exactly the same proportion of balls of each color as urn A. Hence, representativeness leads individuals to infer with certainty that the sample came from urn A.
- 1b. $P(B | x = 2) = 1$. As in 1, if the sample comprises of 2 red balls and 1 black ball, then representativeness guides the individual to infer with certainty that the balls came from urn B.
2. $0.75 < P(A | x = 1) < 1$. Here, the individual makes an inference that is intermediate between a Bayesian posterior and the extreme form of the representativeness heuristic in 1a.

19.2.2 Gambler's fallacy and hot-hands fallacy

Evidence for the law of small numbers comes from at least three kinds of tasks. In *production tasks*, individuals are asked to produce random sequences. In *recognition tasks*, subjects are asked to recognize if the given sequence is random. In *prediction tasks*, subjects are asked to predict the next element in a sequence.

INABILITY TO PRODUCE A SEQUENCE OF RANDOM NUMBERS

One implication of the law of small numbers is that many people are unable to generate a truly random sequence of events. Suppose that individual subjects are asked to generate a sequence of random tosses of a coin; denote heads by H , and tails by T . Most experimental studies find that subjects switch the outcomes too frequently, *as if* to preserve the population proportions (50–50), even in a small sample. So, for instance, the sequence H, H, H, H, H is very likely to be followed by T . In other words, the supposedly random sequences generated by individuals exhibit too much negative autocorrelation (Bar-Hillel and Wagenaar, 1991).⁷

GAMBLER'S FALLACY

If subjects produce *negative autocorrelation* when asked to produce a hypothetical random process, they are said to engage in the *gambler's fallacy*. In a well-known example, there is reduced betting on a winning number in subsequent lottery draws (Clotfelter and Cook, 1993); subjects using the law of small numbers mistakenly think that the winning number has a lower probability of coming up in subsequent draws, despite an unchanged objective probability.

Table 19.1 shows the results of a production task in Rapoport and Budescu (1997) reported by Rabin and Vayanos (2010). These subsequences are culled from an underlying hypothetical sequence of 150 tosses produced by subjects in experiments. The probabilities in the last column are the fraction of subjects who make a choice of H, conditional on the subsequence in the three leftmost columns. On average, the probability of H on the next flip, given that the last toss resulted in a H is $\frac{30+38+41.2+51.3}{4} = 40.1$. It follows, that conditional on the last toss resulting in a H, the average probability of T on the next toss is $100 - 40.1 = 59.9$; this is sometimes

⁷ For important empirical work that served to further establish this phenomenon, see Rapoport and Budescu (1992, 1997) and Budescu and Rapoport (1994).

Table 19.1 Assessed probability that the next flip of a coin will be heads (H), given a sequence of three flips; each flip is either heads (H) or tails (T). Based on Rapoport and Budescu (1997, Table 7, p. 613).

3rd-to-last	2nd-to-last	Very last	Prob. next will be H (%)
H	H	H	30.0
T	H	H	38.0
H	T	H	41.2
H	H	T	48.7
H	T	T	62.0
T	H	T	58.8
T	T	H	51.3
T	T	T	70.0

Source: Rabin and Vayanos (2010).

referred to as the *switching rate*. The difference between the probabilities of T and H is 19.8%, when the objective difference should be 0%. This negative autocorrelation suggests a gambler's fallacy.

There is negative autocorrelation not only with respect to the very last coin flip but also with respect to the "2nd-to-last" flip also. Changing the outcome from H to T on the "2nd-to-last" flip raises the average probability of an H on the next toss (last column in Table 19.1) from 43.9% to 56.1%, an increase of 12.2%. Similarly, changing the outcome from H to T in the "3rd-to-last" flip raises the average probability of an H in the next toss from 45.5% to 54.5%, an increase of 9%. Rabin and Vayanos (2010) propose a model that can account for the gambler's fallacy when the outcomes are not binary and are not iid.

A great deal of the evidence for the gambler's fallacy has come from the study of betting behavior. Clotfelter and Cook (1993) and Terell (1994) find that betting on the winning number falls in pick-3 and pick-4 games and takes several months to recover. Metzger (1984) found that people bet lower amounts on long-shot horses if another long-shot horse had won earlier in the day (the two horses can be different horses). Terrell and Farmer (1996) and Terrell (1998) show that bettors underbid on a previous race winning number in dog races, even if different dogs are involved in the two races.

Croson and Sundali (2005) confirmed the gambler's fallacy from real-world data obtained from casino players with real stakes. They considered betting on roulette wheels where one can either engage in *inside bets* on numbers 0–36 or on *outside bets* (e.g., on even/odd numbers, on numbers shaded red or black, on numbers 1–18, or 1–12, and so on). The method of identifying the gambler's fallacy is to observe a streak of winning numbers, then check if there is reduced betting on these numbers in subsequent bets. For instance, betting on Black following a streak of outcomes of color Red is an example of the gambler's fallacy. Streaks of length $l = 1, 2, \dots, 6$ are considered.

Figure 19.1 shows the results for streaks of lengths 1–6 for all outside bets. For instance, 531 bets were placed on one of the outside bets after observing one spin of the roulette wheel that produced an outside outcome; 255 of these bets, or 48%, of these bets satisfied the criterion of the gambler's fallacy. Following streaks of length 2, 258 outside bets were placed, of which the gambler's fallacy was observed in 51% of cases. Thus, for streaks of small length, the observed




Figure 19.1 Results for streaks of length 1–6 for all outside bets. The height of shaded bars is the percentage of outside outcomes and the height of blank bars is the percentage of outside bets. The horizontal axis measures streaks of length $l = 1, 2, \dots, 6$.

Source: With kind permission from Springer Science + Business Media: *Journal of Risk and Uncertainty* 30(3): 195–209. Croson, R. and Sundali, J. "The gambler's fallacy and the hot hand: empirical data from casinos." © 2005 Springer Science + Business Media.

percentages are no different from random. However, statistically significant effects of gambler's fallacy are found for streaks of length 5 and 6. For instance, for $l = 6$, 85% of the bets are consistent with the gambler's fallacy.

The gambler's fallacy has also been applied to finance. For instance, consider the *disposition effect*, which is the tendency to sell winning stocks and hold on to loss making stocks for too long (Shefrin and Statman, 1985; Odean, 1998). The idea is that people think that a stock that has risen in the past is now due for a fall and vice versa, which is consistent with the gambler's fallacy. Odean (1998) found that people sold winning stock at a 50% higher rate than losing stock. This is a robust phenomenon and documented widely for Finnish investors (Grinblatt and Keloharju, 2001), for Israeli investors (Shapira and Venezia, 2001), for Chinese investors (Feng and Seasholes, 2005), for employee stock options (Heath et al., 1999), and for trading volumes following IPOs (Kaustia, 2004; Brown et al., 2006). The robustness of the disposition effect can be gauged from the fact that even when attempts are made to debias the subjects by making the purchase price very salient in experiments at the time of the selling decision, the disposition effect is reduced by only 25% (Frydman and Rangel, 2014).

Furthermore, while the disposition effect is relatively stronger for individual investors, institutional investors are also subjected to the disposition effect (Frazzini, 2006; Chen et al., 2007; Choe and Eom, 2009). However, Barber et al. (2007) paint a slightly nuanced picture for Taiwanese data; corporate investors and dealers exhibit the disposition effect but mutual funds and foreign investors do not. There is some evidence that the disposition effect can be mitigated with experience (Feng and Seasholes, 2005; Seru et al., 2010), and with the level of wealth (Dhar and Zhu, 2006); and it is more prevalent for hard to value stocks (Kumar, 2009).

While prospect theory has been used to justify the disposition effect, the predictions of the relevant theoretical models are sensitive to the auxiliary assumptions used.⁸ It is more straightforward to explain the disposition effect using the gambler's fallacy. Stocks that are doing well in the past are believed to revert to lower returns, so individuals sell them too early. Stocks that are not doing well are believed to appreciate, so investors hold on to them for too long.

HOT HANDS FALLACY

In contrast to the gambler's fallacy, there are instances where subjects may come to believe that there exists *positive autocorrelation* in an inherently random process. The inference of such positively autocorrelated streaks, known as the *hot hands fallacy*, has been documented in many contexts, such as in the game of basketball.

When basketball players are identified to be on a hot streak, observers assign a high probability that they will be successful in making the next shot too. However, the data does not always justify the assignment of a hot streak, given the small sample at hand. Indeed, several pioneering studies showed that the observed sequence of successful basketball shots was statistically a random sequence (Gilovich et al., 1985; Tversky and Gilovich, 1989a,b). In contrast, the betting behavior of individuals reflects beliefs in such streaks (Camerer, 1989).

Recent, unpublished research, has raised the possibility of actual hot streaks in basketball and baseball, thus, there might be no hot hands fallacy in these contexts (Green and Zwiebel, 2013; Bocskocsky et al., 2014; Miller and Sanjurjo, 2015). The basic idea is that in games like basketball, a player on a hot streak is likely to be more tightly marked by the opposition, and forced to take more difficult shots, hence, reducing the observed impact of hot streaks. Controlling for these factors, a weak hot hands phenomenon has been found in some studies. In Bocskocsky et al. (2014), for instance, players who have hot hands are 1.2 to 2.4 percentage points more likely to make a successful shot in basketball. As the authors note, there are several potential issues with how the level of difficulty of a shot is described, as well as the construction of non-standard measures (such as "complex heat") used to infer the extent of hot hands.

By contrast, Miller and Sanjurjo (2015) found stronger evidence of hot hand effects in the neighborhood of 7 percentage points. It should be pointed out, lest it not be obvious, that a statistical demonstration of the hot hands phenomenon does not necessarily negate the hot hands fallacy. For the negation to occur, it would be necessary to demonstrate that individuals see through the statistical occurrence of the hot hands phenomenon. For instance, a control group of subjects could be tested on a randomly generated sequence of basketball shots with streaks in them. Can the subjects spot that the underlying process is purely random? To the best of my knowledge such controls have not been carried out yet, hence, claims that the hot hands fallacy is not real are premature.

However, there is good evidence from other contexts where the *hot hands phenomenon* does not exist, yet the *hot hands fallacy* exists. In a novel field experiment, Guryan and Kearney (2008) find that sales at lotto stores that have sold a winning ticket soar in the immediate weeks following the lotto win; Figure 19.2 plots the log sales at lotto stores in the 40 previous and 40 subsequent weeks following a winning ticket sold at the store. The effect does diminish over time, yet it lasts for an impressive 40 weeks following the win. One possible confounding factor is that a lottery win in a zip code makes the act of playing a lottery salient. In order to check for this, the sales figures of all lotto stores in the winning zip code area are also examined. There is some overall

⁸ For prospect theory explanations of the disposition effect, see Barberis and Xiong (2009), Hens and Vleck (2011), Henderson (2012), and Kaustia (2010).




Figure 19.2 Retail level data for Lotto Texas wins.

Source: Guryan and Kearney (2008) with permission from the American Economic Association.

increase in sales in the same zip code, suggesting an advertisement, or spillover, effect. However, relative to all stores within the zip code, or stores within a mile of the winning store, the sales increase in the winning store is significantly higher, suggesting a store-specific effect.

A natural explanation for the increased sales at winning stores is the hot hands fallacy. In this case, the hot hands fallacy does not arise from a winning streak, but rather from a single observation. However, bettors typically observe a streak of losses and a win at one store, while there is a continuing streak of losses at all other stores. This observation suggests a possible expansion in the scope of the hot hands fallacy. A competing explanation is that people believe in non-existent variation in luck across stores.

The effect on sales in the winning store is increasing in the size of the win. A week after the win, sales increased in the Texas Lotto stores by 37.7%. The corresponding increase in sales for smaller jackpot games is 21.5% (for the Texas Two Step) and 12.4% (for the Cash Five game). Furthermore, the effect is stronger in areas that have more high school dropouts, higher rates of poverty, and a greater proportion of elderly people, suggesting that some or all of these factors may enhance cognitive biases that give rise to the observed effect.

The hot hands fallacy is also found in Croson and Sundali (2005) who study gambling on roulette wheels. Here, 80% of the subjects quit playing after losing on a spin. In contrast, only 20% quit on winning. This is consistent with the idea that winning players wish to continue playing because they believe that luck is on their side (i.e., they are *hot*, or on a *roll*). The regression results confirm this finding. In this case, for individual-level data, there is a statistically significant effect of winning in one round and the number of bets placed in the next round after controlling for individual heterogeneity.

Jørgensen et al. (2011) find the presence of the gambler's fallacy and the hot hands fallacy using panel data from the Danish State Lottery. There are two types of individuals in their data. The first type picks the same numbers every week. The second type of individuals avoids the non-winning numbers they chose last time (gambler's fallacy) and are more likely to choose numbers that are perceived to be on a winning streak (hot hands fallacy).

Powdthavee and Riyanto (2015) show that individuals are willing to pay for an expert who they believe has a hot hand in predicting inherently random sequences of numbers. This theme is

explored in chapter 20 in Kahneman (2011) under the title “The Illusion of Validity.” He argues that the returns to professional funds are typically (although not always) the result of luck rather than skill.⁹ As evidence, he cites the fact that the temporal correlations in performance for any mutual fund are quite weak, and close to zero, which is consistent with the luck story. Yet, small investors are willing to pay significant amounts for expert financial advice. Malmendier and Tate (2009) find that superstar CEOs (based on previous performance) underperform in subsequent periods. This is also consistent with the story in which high outcomes are achieved, at least partially, by luck, but one eventually experiences regression to the mean (see Section 19.10.1). Some of these issues are taken up in greater detail in Section 21.7 in Chapter 21.

WHY DO THESE FALLACIES ARISE?

Tversky and Kahneman (1974) report that in successive tosses of an unbiased coin, the sequence H, T, H, T, T, H is considered more likely than the sequence H, H, H, T, T, T ; however, both sequences have identical statistical probabilities. One plausible explanation is that the second sequence is not considered to be *representative* of the population proportions. As the authors beautifully put it (p. 1125): “Chance is commonly viewed as a self-correcting process in which a deviation in one direction induces a deviation in the opposite direction to restore the equilibrium. In fact, deviations are not ‘corrected’ as a chance process unfolds, they are merely diluted.” This also explains the role of the representativeness heuristic in the gambler’s fallacy; gamblers do not bet on the previously winning numbers because they believe that deviations will be corrected, rather than diluted over a sufficiently long period of time.

Like the gambler’s fallacy, the seemingly opposite phenomenon of positive correlation when the underlying process is the hot hands fallacy, can also be potentially explained by the representativeness heuristic. In this case, following several identical signals, the sequence of observed signals appears more and more representative of a positively autocorrelated underlying random process. An observer who is subject to the hot hands fallacy then makes predictions based on the incorrectly perceived underlying random process.

Ayton and Fischer (2004) suggest that the domains of the gambler’s fallacy and the hot hands fallacy might be different. They propose that the hot hands fallacy is more likely to arise for human performance, as is the case with basketball and tennis players, rather than for inanimate processes such as throws of a dice or coin tosses. The reason is that a streak in sports may signal unobservable variables about a player to observers, such as the level of confidence, fatigue, and the quality of training prior to the game. These unobservable variables do not apply to inanimate processes for which subjects may switch too often in small samples to preserve the population proportions, as suggested by the representativeness heuristic. This also potentially explains emerging evidence of hot hands in basketball and baseball referenced above, yet in inanimate processes such as the outcome of lotteries, there is strong evidence of the hot hands fallacy (see above).

In one set of experiments, the authors presented the subjects with various sequences of outcomes that had varying degrees of positive and negative autocorrelation. Subjects were asked to guess the source of these sequences: human skilled performance (e.g., a professional basketball player’s scoring attempts, or a professional tennis player’s first serve in) or an inanimate process (e.g., a coin toss, or a binary outcome roulette wheel). Each sequence consisted of 21 binary

⁹ Kahneman writes (2011, p. 218): “The idea that large historical events are determined by luck is profoundly shocking, although it is demonstrably true.” He gives an interesting example. Before an embryo is fertilized there is a 50–50 chance that it will be a male or a female. Thus there was a 1/8 chance that the world could have survived the three dictators, responsible for a large number of twentieth century deaths—Hitler, Stalin and Mao Zedong.”




Figure 19.3 Is the observed sequence with a given alternation rate more likely to have come from basketball or a coin toss?

Source: With kind permission from Springer Science + Business Media: *Memory and Cognition* 32: 1369–78. Ayton, P. and Fisher, I. (2004). "The hot hand fallacy and the gambler's fallacy: two faces of subjective randomness?" © 2004 Springer Science + Business Media.

outcomes. The alternation rate is the number of times a sequence changes direction divided by 10.

The results for one of the comparisons is shown in Figure 19.3. Consistent with the authors' predictions, sequences with low rates of alternation were rated by subjects to have more likely come from basketball rather than coin tosses. Conversely, sequences with high rates of alternation were rated as more likely to have come from a coin toss. These results highlight when the gambler's fallacy is more likely to arise (high alternation rates arising from an inanimate process) and when the hot hands fallacy is more likely to arise (low alternation rates and involvement of human skills).

19.2.3 A formal model of the law of small numbers

Consider one possible method of formalizing the law of small numbers in a simple model with the following features (Rabin, 2002).¹⁰

1. Move by nature: Nature chooses a parameter $\theta \in \Theta$. Individuals have a prior distribution $\pi(\theta)$ over the set Θ .
2. The *actual* data generating process: Consider a Bernoulli trial with two possible outcomes: x (success) and y (failure) that occur with respective probabilities $\theta \in [0, 1]$ and $1 - \theta$. Suppose that we independently repeat the Bernoulli trial n times, with replacement. Denote by n_x , the number of times that the outcome equals x in the n trials. It is well known that the number of successes follows a binomial distribution; the probability that there are n_x successes out of n Bernoulli trials is given by

$$p(n_x) = \binom{n}{n_x} \theta^{n_x} (1 - \theta)^{n - n_x}; \quad n_x = 0, 1, \dots, n.$$

¹⁰ For a theoretical application to finance, see Section 21.6 below.

3. Misperception about the data generating process: The true underlying process is one with replacement. An individual who knows the true model of the world is known as a *Bayesian*. However, suppose that individuals believe, incorrectly, that the data is generated by the following process, without replacement.

Let time period t be an odd number. At time t , an urn has N balls; θN balls denote outcome x and the remaining, $(1 - \theta)N$, balls denote outcome y . A ball is drawn with replacement at time t . At time $t + 1$, an even time period, the individual believes (incorrectly) that there are only $N - 1$ balls left in the urn. In period $t + 2$, which is an odd time period, the urn is replenished again to N balls.

Thus, in an odd period, the probability of outcome x is $\frac{\theta N}{N}$. Now suppose that outcome x materializes in an odd period. Then, in the following period (an even period), the individual believes (incorrectly) that the probability of the outcome x is $\frac{\theta N - 1}{N - 1}$, while the probability of outcome y is $\frac{(1 - \theta)N}{N - 1}$. To ensure that the problem is not vacuous, we assume that there are at least two balls corresponding to each outcome in every possible urn. So $\theta N \geq 2$ and $(1 - \theta)N \geq 2$.

Such an individual, who believes in an incorrect model of the world, but is otherwise skilled in all aspects of classical statistical inference, is referred to as an *N-Freddy*. If N is very large, then *N-Freddy* behaves very much like a Bayesian but for small N he can be very biased.

CERTAINTY ABOUT THE TRUE θ AND THE GAMBLER'S FALLACY

The gambler's fallacy is a direct implication of the model. Consider the probability of outcome x in an "even period" following outcome x in the previous (odd) period, i.e., the conditional probability $P(x | x)$. For a Bayesian who uses the correct model, $P(x) = P(x | x) = \theta$. In contrast, for *N-Freddy*, who uses an incorrect model, $P(x) = \frac{\theta N}{N} = \theta$ and $P(x | x) = \frac{\theta N - 1}{N - 1} < \theta$. Hence, following an outcome x , *N-Freddy* reduces the probability that the next draw will produce the signal x , which is precisely the gambler's fallacy. In contrast, the Bayesian is not subject to the gambler's fallacy.

Suppose we say that Freddy gets even "Freddier" if N falls. What are the implications for the gambler's fallacy? For an *N-Freddy*, $\frac{\partial P(x|x)}{\partial N} = \frac{1 - \theta}{(N - 1)^2} > 0$. So, as N falls, $P(x | x)$ also falls and the gambler's fallacy becomes even worse.

Whilst the Bayesian decision maker and Freddy might revise beliefs in the same direction, there are systematic differences between their predictions. To see this, suppose that there are two values of θ : $\theta_L < \theta_H$. Starting in an odd period, consider a sample of two consecutive x outcomes. For any θ , the likelihood of this sample for Freddy is $L(x, x | \theta) = \frac{\theta N}{N} \frac{\theta N - 1}{N - 1} = \frac{\theta(\theta N - 1)}{N - 1}$. Thus for two different values, θ_L, θ_H , the likelihood ratio for Freddy is

$$\phi_F = \frac{L(x, x | \theta_H)}{L(x, x | \theta_L)} = \frac{\theta_H(\theta_H N - 1)}{\theta_L(\theta_L N - 1)}. \quad (19.3)$$

On the other hand, for a Bayesian, who uses the correct model of the world, $P(x) = P(x | x) = \theta$, so the relevant likelihood ratio is

$$\phi_B = \frac{L(x, x | \theta_H)}{L(x, x | \theta_L)} = \frac{\theta_H^2}{\theta_L^2}. \quad (19.4)$$

Comparing (19.3) with (19.4) we get that $\phi_F > \phi_B$. Thus, Freddy assigns a lower likelihood to the signal coming from $\theta = \theta_L$ relative to the Bayesian. In other words, observing two successive x outcomes, Freddy assigns a much greater likelihood of the urn containing a larger number of x outcomes, relative to a classical statistician.

UNCERTAINTY ABOUT THE TRUE θ AND OVERCONFIDENCE

Now suppose that the urn has 4 balls, $N = 4$ (so we have a 4-Freddy) and Freddy is uncertain about the true value of θ . Suppose that the outcome can either be *good*, x , or *bad*, y . One may imagine that the draw of a ball from an urn corresponds to the outcome of a production/financial task that is overseen by a manager. We make two assumptions.

1. The probability of the good outcome, θ , takes three equally likely values, $\theta_L = \frac{1}{4}$, $\theta_M = \frac{2}{4}$, $\theta_H = \frac{3}{4}$. Thus, θ indexes the competence of the manager; higher values denote greater competence.
2. The urn is now replenished every four periods and the starting period is odd as before.

Thus, a 4-Freddy believes that the outcomes for a type θ_L manager are drawn from an urn containing four balls such that only one ball corresponds to the good outcome x and three balls correspond to the bad outcome y , thus, $x = 1$, $y = 3$. Similarly, for types θ_M and θ_H , $x = 2$ and $x = 3$, respectively. On the other hand, a Bayesian knows that all draws are independently drawn from a Bernoulli distribution and the urn is replenished every period.

Suppose that starting in period 1, the performance of the manager is good for two consecutive periods. The Bayesian knows that the random draws are identical and independent, so for any θ , he assigns the conditional probability $P(x, x | \theta) = \theta^2$, and, in particular,

$$P(x, x | \theta_L) = \frac{1}{16}, P(x, x | \theta_M) = \frac{1}{4}, P(x, x | \theta_H) = \frac{9}{16} > \frac{1}{2}. \quad (19.5)$$

Given assumptions 1 and 2, the 4-Freddy instead assigns the conditional probabilities (or sample likelihoods) as follows.

$$P(x, x | \theta_L) = \frac{1}{4} \times \frac{0}{3} = 0, P(x, x | \theta_M) = \frac{2}{4} \times \frac{1}{3} = \frac{1}{6}, P(x, x | \theta_H) = \frac{3}{4} \times \frac{2}{3} = \frac{1}{2}. \quad (19.6)$$

For any θ , relative to a Bayesian, the 4-Freddy assigns a lower value of $P(x, x | \theta)$. The reason is that once an x outcome is obtained, Freddy believes that the urn contains one less x outcome, reducing the probability that a subsequent outcome will be x .

We are now interested in computing $P(\theta_H | x, x)$, i.e., the posterior belief that the analyst is of the highest type, following a sequence of two good performances. Using Bayes' rule we get that

$$P(\theta_H | x, x) = \frac{P(x, x | \theta_H)P(\theta_H)}{P(x, x | \theta_L)P(\theta_L) + P(x, x | \theta_M)P(\theta_M) + P(x, x | \theta_H)P(\theta_H)}. \quad (19.7)$$

Noting that $P(\theta_L) = P(\theta_M) = P(\theta_H) = \frac{1}{3}$ (because each type is equally likely) and substituting (19.5) and (19.6) sequentially into (19.7), Freddy computes $P(\theta_H | x, x) = \frac{21}{28}$, while a Bayesian computes $P(\theta_H | x, x) = \frac{18}{28}$. Thus, based on a sample of two observations, the 4-Freddy is more confident than the Bayesian that the fund manager is of the highest type. This is a form of over-

confidence, based on observing very small samples, which plays an important role in behavioral finance.

Overconfidence is considered in greater detail in Section 21.7 in Chapter 21. Here, we offer some brief comments. Suppose that in financial markets, an observer who behaves like Freddy observes the performances of several financial analysts. Then those analysts who, purely by luck, achieve successive good outcomes, say, x, x, x , will be believed by the observer to be more likely to be extremely competent ($\theta = \theta_H$), relative to a Bayesian. Analogously those who achieve a sequence of bad outcomes, say, y, y, y , again purely by luck, are believed to be more likely to have low competence ($\theta = \theta_L$), relative to a Bayesian. Indeed, the observer may be willing to pay a premium to use the services of the analyst that they mistakenly perceive to be more competent (Powdthavee and Riyanto, 2015).

This situation is not unrepresentative, even if it is artificially constructed. Kahneman (2011, ch. 20) recounts that he was once invited to speak to a group of financial advisors who provided their services to very wealthy clients. Examining the data on 25 of those advisors over eight years, he ranked the advisors by performance in each year and then computed the correlation coefficients of the ranks in each pair of the eight years. The average of the correlations was close to zero, 0.01. Thus, it appears that it is not skill but luck that governs the relative performance of the advisors and the firm was rewarding luck not skill.

Needless to say, exposure to these results did not go down very well with the analysts. As Kahneman (2011, p. 216) puts it: “Their own experience of exercising careful judgment on complex problems was far more compelling to them than an obscure statistical fact.” Since financial analysts are typically good at explaining past financial events with the benefit of hindsight (see hindsight bias in Section 19.8 below), an understanding of the past gives them the mistaken view that they can predict the future.¹¹ For a theoretical model of overconfidence in finance that is consistent with the views of Kahneman (2011) described above, see Gervais and Odean (2001). In this model, investors who do better than others, purely by luck, view their success as vindication of their skill rather than luck. This creates overconfidence that might nevertheless be tempered by subsequent experience.

19.3 Conjunction fallacy

Another illustration of the violation of a basic principle of probability theory is the *conjunction fallacy*. Tversky and Kahneman (1983) posed the following problem to Stanford students who were well-trained in decision sciences.

“Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.” Which of the following statements is most likely?

1. Linda is a bank teller (B).
2. Linda is active in the feminist movement (F).
3. Linda is a bank teller and is active in the feminist movement (BF).

¹¹ For more details, see Chapters 19 and 20 in Kahneman (2011) under the headings “Illusion of understanding” and “Illusion of validity.”

Since the event $BF \subseteq B$ and $BF \subseteq F$, it follows under classical statistics that $P(BF) \leq P(B)$ and $P(BF) \leq P(F)$. Any violation of these inequalities is a *conjunction fallacy*. The problem stated above has three choices: B , F , and BF . An earlier version of the problem was handed out with eight choices. Between 85–90% of the respondents rated the event BF as more likely relative to the separate events B and F . Thus, subjects exhibit the conjunction fallacy. In one version, 85% of the sample of doctoral students in the decision-science program at the Stanford graduate school of business committed the conjunction fallacy.

There is a very large literature on the conjunction fallacy. In addition to the literature reviewed below, the interested reader may consult the survey in Moro (2009) and Kahneman (2011, ch. 15). For a skeptical view, see Hertwig and Gigerenzer (1999). For a review of the literature on the causes of the conjunction fallacy, see Tentori et al. (2013).

The *Linda problem*, as it came to be known, was bitterly criticized and continues to be criticized, although, as we show below, no conclusive evidence of its rejection has been given. Indeed Tversky and Kahneman (1983) themselves addressed many potential criticisms, but their caveats were conveniently forgotten.

19.3.1 *Does a natural frequency format resolve the conjunction fallacy?*

Hertwig and Gigerenzer (1999) report that if the Linda problem is posed in terms of natural frequencies (e.g., 15 out of 100 rather than 15%), then subjects do not engage in the conjunction fallacy. However, Kahneman and Tversky (1996) find that presenting the information in natural frequencies does not eliminate the conjunction fallacy. They presented the following Linda problem to three groups of subjects (Groups 1, 2, 3) in a natural frequency format.

Linda is in her early thirties. She is single, outspoken, and very bright. As a student she majored in philosophy and was deeply concerned with issues of discrimination and social justice. Suppose there are 1000 women who fit this description. How many of them are

- (a) High school teachers? [Groups 1, 2, 3]
- (b) Bank tellers? [Groups 1, 2]
- (c) Bank tellers and active feminists? [Groups 1, 3]

In the between-subjects treatment, the median response for choice (c) was statistically larger than the median response for choice (b), confirming the conjunction fallacy. Subsequent evidence in Tentori et al. (2004) also supports the view that presenting the information in the form of natural frequencies does not eliminate the conjunction fallacy.

19.3.2 *Do monetary incentives and groupthink eliminate the conjunction fallacy?*

Charness et al. (2010) introduce incentives in the Linda problem and allow for consultation in groups of twos (pairs) and threes (trios); as a baseline, they also consider individual decisions (singles). The incentive for giving the correct answer is an extra \$2 over the show-up fee of \$2. The results are shown in Table 19.2; T&K, 1983 is Tversky and Kahneman (1983) and CLK, 2008 is the working paper that forms the basis of the reported results in Charness et al. (2010).

Tversky and Kahneman (1983) find that 85.2% of their subjects engage in the conjunction fallacy. The corresponding number in Charness et al. (2010), based on a test of proportions, is 58.1%

Table 19.2 Violations of the conjunction fallacy.

Study	Details	Incorrect answers/ total sample	Error rate (percent)
<i>Individuals</i>			
T&K, 1983	UBC undergrads, no incentives	121/142	85.2
CKL, 2008	UCSB students, singles, no incentives	50/86	58.1
CKL, 2008	UCSB students, singles, incentives	31/94	33.0
CKL, 2008	UCSB students, total singles	81/180	45.0
<i>Pairs</i>			
CKL, 2008	UCSB students, in pairs, no incentives	27/56	48.2
CKL, 2008	UCSB students, in pairs, incentives	5/38	13.2
CKL, 2008	UCSB students, total in pairs	32/94	34.0
<i>Trios</i>			
CKL, 2008	UCSB students, in trios, no incentives	10/39	25.6
CKL, 2008	UCSB students, in trios, incentives	5/48	10.4
CKL, 2008	UCSB students, total in trios	15/87	17.2

Source: Charness et al. (2010).

and it is statistically smaller. Furthermore, Table 19.2 shows that the violation of the conjunction rule falls significantly when singles are given incentives (from 58.1% to 33%).

The conjunction rule is violated even less when one moves from pairs to trios, as compared to moving from singles to pairs.¹² Thus, “three heads are better than one.” In the presence of incentives, trios violate the conjunction rule in only 10.4% of the cases. Since university students are likely to be better trained in probability theory, it will probably be worth replicating these results with a non-student population. Furthermore, the high rates of violations of the conjunction fallacy for singles leaves open the possibility that in real-world situations with singles, the conjunction fallacy may yet play an important role.

In contrast to these results, Bonini et al. (2004) show that monetary incentives do not crowd out the conjunction fallacy. They give subjects the option of allocating 7 euros among three options: (1) Option *B*. (2) Option $B \wedge F$ (where $\wedge$ is the logical operator “and”). (3) Option $B \wedge (\sim F)$, where $\sim F$ is the option “not *F*.” The third option controls for the objection that subjects may actually be interpreting (1) as (3), i.e., the option *B* as $B \wedge (\sim F)$. It is a dominant action to bet on Option *B*, so if subjects did not engage in the conjunction fallacy, we should observe all subjects to bet 7 euros on Option *B*.

Across all blocks of experiments, the average amount allocated to the first option was 1.99 euros. In contrast, the corresponding allocation to options 2 and 3, respectively, the dominated bets, was 3.15 euros and 1.86 euros. The difference between the allocations to options 1 and 2 is statistically significant at 5% significance, using a one tailed t-test. In the five blocks of experiments, the average percentage of subjects that engaged in the conjunction fallacy was 85.32, which is nearly identical to the figure of 85% reported in Tversky and Kahneman (1983). Furthermore,

¹² Similar advantages of group consultation in the reduction of stochastically dominated options is reported in Charness et al. (2007).

the experimental results in Stolarz-Fantino et al. (2003) further argue for the robustness of the conjunction fallacy to monetary incentives and framing effects.

19.3.3 *Why does the conjunction fallacy arise?*

One possible explanation of the conjunction fallacy is the use of the representativeness heuristic. The description of Linda as a feminist and a bank teller is relatively more representative of her described traits. Kahneman (2011, p. 158) writes of his initial excitement about the result: “we had pitted logic against representativeness, and representativeness had won!”

Conditional on Linda’s description, it is plausible that Linda is a member of the feminist movement. If this were substituted by a non-plausible event (e.g., Linda works for the local council), then there is no presumption that the conjunction fallacy would arise. Consider the following problem given in Kahneman (2011).

Which alternative is more probable?

Mark has hair.

Mark has blond hair.

In this problem, the given information does not make blond hair more plausible, hence, as expected, no conjunction fallacy was found.

Yet another potential explanation of the conjunction fallacy is based on interpreting the role of “and” in the Linda problem (Camerer, 1995, p. 598). The explanation centers on the validity of the interpretation of the word “and” as the logical operator $\wedge$ such that $A \wedge B$ means the intersection of the sets A and B . Hertwig et al. (2008) argue that in the statement A and B , “and” could be interpreted as (i) events A and B happening in temporal sequence as in “I went to the store and bought some whisky”, or (ii) as a causal relationship “Smile and the world smiles with you.” Hertwig et al. (2008) then go on to explore the precise sense in which their subjects in experiments interpret “and.” They find that subjects who interpret “and” as the logical operator are less likely to be engaged in the conjunction fallacy.

However, Tentori et al. (2004) show that although the conjunction fallacy is lessened by taking account of the literal differences between the conjunction “and” and the logical operator $\wedge$, it nevertheless remains an important fallacy. A fairly detailed discussion of these issues is provided in the careful study by Tentori and Crupi (2012). They concede that the word “and” may have different meanings in different sentences, but within the same sentence its meaning is often unambiguous. Since “and” is one of the most widely used words in the English language, ambiguity of its use within the same sentence would have been disastrous for communication, which has been at the heart of human evolution. They also point out that the multiple conjunctions considered in Hertwig et al. (2008) (e.g., in their experiment 3) have even lower probability than the single conjunctions in the Kahneman–Tversky results and these do not disprove the conjunction fallacy. They argue, in particular, that the theoretical foundations for the claims and the empirical inferences in Hertwig et al. (2008) are questionable. They perform two novel experiments and their results directly contradict the Hertwig et al. results.

Kahneman (2011, pp. 164–5) summarizes his recent feelings about the fallout from his research on the Linda problem: “The Linda problem attracted a great deal of attention, but it also became a magnet for critics of our approach . . . These arguments were sometimes extended to suggest that our entire enterprise was misguided: if one salient cognitive illusion could be weakened or explained away, others could be as well . . . The net effect of the Linda problem was an increase in the visibility of our work to the general public, and a small dent in the credibility of our approach among scholars in the field. This was not at all what we had expected.”

19.4 The availability heuristic

Tversky and Kahneman (1974) describe the availability heuristic as follows: “There are situations in which people assess the frequency of a class or the probability of an event by the ease with which instances or occurrences can be brought to mind. For example, one may assess the risk of heart attack among middle-aged people by recalling such occurrences among one’s acquaintances.”

As one might expect, this heuristic can lead to inferential mistakes. In one experiment, Tversky and Kahneman (1974) presented subjects with a mixed-gender list of names in which the men and women were equally represented. In treatment 1, the men in the list were famous but not the women. In treatment 2, the opposite was true. Subjects have to infer if there were more men named in the list than women. In treatment 1 subjects reported a higher number of men, while in treatment 2, they reported more women. Thus, when instances in a class are more easily retrieved, the class appears to be more numerous, even if the alternative classes are equally represented.

Salience of an event might also influence how easily, similar events are retrieved from memory. So, for instance, subjects assign higher probabilities to road accidents or house fires, after having just seen one. Tversky and Kahneman (1973) ask subjects the probability that a randomly picked three letter word from the dictionary begins (event *A*) or ends (event *B*) with the consonants *r* or *k*. Because it is easier to retrieve words beginning with the letters *r* or *k* rather than those that end with these consonants, subjects assigned greater probability to event *A*. However, the correct answer is event *B*.

If subjects cannot recall events from memory, then the ease of hypothetically constructing specific instances will determine which event is thought to be more likely. Tversky and Kahneman (1974) ask subjects in experiments to imagine, or hypothetically construct, all the possible contingencies, risks, or dangers that might arise from a risky expedition. If it is particularly easy to mentally construct these contingencies, then the expedition may be deemed too dangerous and abandoned. Conversely, if the contingencies are too hard to construct, then the opposite could be true. However, ease or difficulty of constructing these contingencies might not be perfectly related to the actual dangers involved in the expedition. Hence, in each case, individuals following such a heuristic could make potentially fatal errors.

More vivid memories may be more easily retrieved. Hence, individuals may assign higher subjective probabilities to a range of events such as airplane accidents, shark attacks, muggings, and snake bites. The media plays a potentially important role in creating vivid memories and selecting for us the subsets of events that can be retrieved. In particular, the media is often driven by considerations of the novelty, rarity, and poignancy of events. In the context of the availability heuristic, Eisenman (1993) shows that, based on media stories, individuals may believe that drug use has increased, when in fact it has decreased. Lichtenstein et al. (1978) find that people’s perception of risk is higher for events that are more extensively covered by the media. As Kahneman (2011, p. 138) puts it: “The world in our heads is not a precise replica of reality.”

The empirical study of Lichtenstein et al. (1978) has been very influential in establishing the availability heuristic.¹³ Their main contribution was to show a tight link between the frequency of death by various causes and two proxies for the availability of such events in human memory. They discovered that the first proxy, *direct experience* (e.g., death of a friend or close relative),

¹³ For a review of the empirical literature on the availability heuristic, see Betsch and Pohl (2002) and Kahneman (2011).

predicted better as compared to the second proxy, *indirect experience* (e.g., exposure to statistics via the media). This result is supported by Hertwig et al. (2005) who found that the best predictor of an individual's choice between the risks of two alternative causes of death is the knowledge of similar events in one's social network.

We consider now the empirical study by Pachur et al. (2012) that not only replicates the empirical work of Lichtenstein et al. (1978) but also draws attention to the relation between the availability heuristic and the *affect heuristic* that we consider below in Section 19.5. In the first set of experiments, Swiss university students were given three tasks. In the incentivized *choice task*, among the set of 24 cancers, subjects had to guess for each possible pair (276 pairs in total), the cancer that causes more deaths per year. In the *estimation task*, subjects had to guess the number of deaths caused by each type of cancer per year. In the *valuation task*, individuals were required to make *value of statistical life judgments* (VSL); this required them to place a monetary value, measured in Swiss francs, on one life saved from each type of cancer. For each type of cancer, individuals also reported the number of deaths that they could recall within their social network.

The results are shown in Table 19.3; *Mdn* denotes the median value. For the moment, we ignore column 5, the dread score; it is used later in our discussion of the affect heuristic. The main results are as follows.

1. In the binary *choice task*, there were 62.2% correct answers, i.e., subjects guessed correctly which of the two cancers caused greater fatalities.
2. In the *estimation task*, the correlation between the estimated and actual fatalities is 0.46 ($p = 0.2$).
3. In the *valuation task*, there is much variation in the monetary figures that subjects assign to saving a life from a particular kind of cancer. For instance, the VSL figure on breast cancer is four times higher than on rectal, bladder, or ovarian cancer. A fuller explanation of these results is given in the discussion on the affect heuristic below.
4. The correlation between the median estimated frequency of risk (column 3) and the mean number of recalled instances (column 6) was 0.62 and this is highly significant at the 1% level. The corresponding correlation reported in Lichtenstein et al. (1978) was 0.5 to 0.9.

In a second set of experiments, the authors consider the classic set of risks in Lichtenstein et al. (1978). These risks pertain to diverse causes of death such as diseases, homicides, asthma, suicides, murders, and natural hazards. The three tasks in this experiment (choice, estimation, and valuation) are identical to the first experiment. The main difference is that the last column in Table 19.3 is now split up into direct experiences (number of recalled instances in one's social network) and indirect experiences (the number of recalled instances in the media).

The results are as follows. In the binary choice task that compared the relative mortality rates from two different causes of death, subjects guessed the correct answer in 69% of the cases. The correlation between the estimated frequencies and direct experience is 0.50, which is significant at the 1% level. This suggests an important role of the availability heuristic. However, the correlation of indirect experiences with the estimated frequency is low.

Kuran and Sunstein (1999) introduce the idea of *availability cascades*. This refers to a self-fulfilling sequence of events that are typically sparked by an initial media report about an event that is not major, followed by an emotional, affective, reaction from the public. This gives rise to another heightened round of emotional, self-fulfilling events, and so on. Eventually, and typically, this leads to a policy response that is not justified on cost-benefit terms. The interested reader can consult the authors' explication of their ideas based on two case studies. The *Love Canal affair* and the *Alar scare*.

Table 19.3 Empirical evidence for the *availability* and *affect* heuristics.

	Annual mortality rate	<i>Mdti</i> frequency estimate	<i>Mdn</i> VSL	Affect (M dread score)	Availability (M number of instances in social network)
Penis cancer	10.2	40.0	10000	4.65	0.00
Testicular cancer	17.2	109.5	25000	4.55	0.03
Bone cancer	37.5	80.0	30000	4.88	0.12
Thyroid cancer	69.3	80.0	30000	4.44	0.03
Larynx cancer	94.2	50.0	20000	3.92	0.24
Cancer of the connective tissue	94.3	45.0	17500	4.40	0.00
Cancer of the gall bladder	196.5	30.0	11010	4.31	0.09
Malignant melanoma (skin cancer)	242.0	50.0	17102	3.59	0.24
Cervical cancer	295.8	95.0	20000	4.39	0.15
Renal cancer	339.2	50.0	20000	4.30	0.03
Cancer of the mouth and throat	351.0	60.0	12609	3.91	0.00
Esophageal cancer	384.5	50.0	15000	4.31	0.09
Rectal cancer	437.2	80.0	10000	4.14	0.00
Bladder cancer	450.5	30.0	10616	4.23	0.03
Ovarian cancer	453.2	100.0	20000	4.43	0.03
Cancer of the nervous system	455.0	90.0	30000	5.02	0.32
Hepatic cancer	513.0	110.0	10000	4.04	0.18
Stomach cancer	572.2	115.0	20000	4.22	0.18
Pancreatic cancer	897.8	60.0	17501	4.43	0.26
Colon cancer	1172.2	120.0	10000	4.23	0.15
Prostate cancer	1312.3	87.0	23115	4.53	0.32
Leukemia and lymphoma	1331.7	100.0	35000	4.91	0.38
Breast cancer	1347.3	200.0	40000	4.24	0.79
Lung cancer	2756.0	300.0	22500	4.00	0.50

Source: Pachur et al. 2012.

19.5 The affect heuristic

There is a correlation between the ease of retrieval of information about an event, which is the basis of the availability heuristic, and the emotional, or *affective*, reactions to the event. In making judgments, individuals tag the *affect* or the positive and negative quality/feeling *induced* by the representation of an event, signal, or stimulus. These tags, in turn, produce an *affect pool* that individuals draw on, in making judgments and decisions (Slovic et al., 2002). The affective response is the immediate response to a stimulus, as distinct from a deliberative response. We have already encountered the affect heuristic in the book. Two leading examples come from Part 6. These are

the *risk as feelings* hypothesis (Loewenstein et al., 2001) and the work of Antonio Damasio that studied the link between affective reasoning and rational judgment.

In an early contribution, Zajonc (1980) noted that we do not just see a car, we see a *sleek car*, a *cool car*, an *ugly car*, depending on the *affect* of the car on us. Slovic et al. (2002, p. 398) write: “We sometimes delude ourselves that we proceed in a rational manner and weight all the pros and cons of the various alternatives. But this is probably seldom the actual case. Quite often ‘I decided in favor of X’ is no more than ‘I liked X’ . . . We buy the cars we ‘like,’ choose the jobs and houses we find ‘attractive,’ and then justify these choices by various reasons.” This view is radically different from the textbook consumer theory taught to students of economics.

Kahneman (2011) notes that these affective shortcuts are an instance of *substitution* in which one substitutes the answer to a hard question (What do I think about it?) by the answer to an easier question (How do I feel about it?). This provides a cognitive shortcut to engaging in a deliberative response that is time consuming, particularly in the case of complex and difficult problems that must be decided within a time constraint. For this reason, we speak of a heuristic in these situations—the *affect heuristic*.

On the one hand, the affect heuristic offers a speedy solution to complicated problems. On the other hand, it might not lead to an optimal payoff for an individual. Firms routinely use diverse strategies to tap into the affect heuristic of consumers to increase their sales (see the problems at the end of Part 7). The affective system is not very good at tracking small changes, or changes that are remote in time. This explains the difficulty in resisting one more fag (Slovic et al., 2002), which is not optimal for an individual interested in kicking the habit.

The policy recommendation by Cass Sunstein, in several publications, is to base public policy on a strict cost–benefit basis in order to ensure maximum mileage for the taxpayer’s dollar. This is in contrast to the position taken by Paul Slovic, the leading advocate of the affect heuristic (see below), who argues that there are no objective risks. Since experts typically measure objective risks (e.g., lives saved or lost) whereas the public has a subjective notion of risk, basing policy on a cost–benefit basis ignores the will of the public, which is not democratic. For a further discussion on this fascinating issue, see Kahneman (2011).

19.5.1 *Empirical evidence for the affect heuristic*

Fischhoff et al. (1978) showed that the perception of risks of various hazards depended on the *dread* that these hazards invoked. The hazards included antibiotics, X-rays, vaccines, smoking, and alcohol. Furthermore, the judgment of risks of various hazards, and the subjective benefits that individuals assign to these risks, are negatively correlated. Thus, individuals assign lower benefits to risks that they dread. However, there need not be any relation between the degree of risk and the benefits associated with an event. These results show that human judgments are based not only on what they *think* but also, perhaps mainly, on what they *feel*.

Alhakami and Slovic (1994) found that the relation between perceived risks and benefits from activities was influenced by how *affectively charged* the activity was. Finucane et al. (2000) exploited these results to show that the affect of an event, substance, object, or technology (we use the collective term “hazard”) determines both the perceived risk as well as the perceived benefit. Exploiting an implication of this insight they found that (i) the perceived benefits of a hazard can be altered by revealing the degree of risk, and (ii) the perceived risk can be altered by revealing the benefits. For instance, in the case of nuclear power, subjects assigned low (high) risks when it was revealed to them that the benefits are high (low). These results were even stronger when subjects had to make their decisions under greater time pressure.

Slovic et al. (1999) replicated these results with experts who were members of the British Toxicology Association, demonstrating that the affect heuristic is probably hardwired and difficult to undo with training. The experts were given a list of 30 chemical items and asked to make a quick intuitive rating of the items on an *affect scale* that ranged from bad to good; the chemicals included second hand smoking, dioxins in food, benzene, and aspirin. Once the experts had made the affect rating, they were asked to make a risk judgment of exposure to each chemical when the exposure level was 1/100 of the level rated to be dangerous by an appropriate regulatory agency. Clearly, at these levels, the chemicals should carry little or no risk and there should probably be little variation in the assessed risks. However, the affect ranking and the risk ranking of the chemicals was strongly negatively correlated for 95 out of the 97 experts; chemicals that had a relatively bad affect were judged to be relatively more risky.

We now revisit the data in Table 19.3 and interpret the results in column 5 that reports the dread score based on 12 risk characteristics in Slovic et al. (1980). Experts tabulate deaths by objective measures such as the number of lives lost. However, for most individuals, there are *better deaths* and *worse deaths*. For instance, death in a ghastly accident, or a painful death, elicits a more affective response as compared to learning that someone died in their sleep. Thus, the perception of individuals, unlike those of experts, about the risks associated with various cancers, or the risk of death by various causes, are often based on an affective reaction.

For the first set of experiments on the set of 24 different cancers (reported in Table 19.3), the availability heuristic performs better than the dread score in predicting binary choices among fatalities arising from different types of cancers. However, the dread score had better explanatory power in explaining the VSL judgment. The correlation between the two is 0.44 with a p-value of 0.03. In the second set of experiments (classic set of risks in Lichtenstein et al., 1978), the availability heuristic does better than the affect heuristic when it comes to predicting fatalities in binary comparisons of a classic set of risks. However, in the VSL judgment, both heuristics do equally well.

19.5.2 *The relative effectiveness of the probability and frequency formats*

Slovic et al. (2000) explored the effect of frequency representations and equivalent probability representations of information on the affect heuristic. The subject pool was experienced forensic psychologists and psychiatrists. Subjects were asked to guess the likelihood that a patient discharged from a mental hospital would commit an act of violence within six months of being discharged. In the first treatment, subjects were given the relative frequency estimate of other experts of similar patients committing the violent act; e.g., 10 out of every 100 such patients commit a violent act. In the second treatment, the view of other experts was given as a percentage (e.g., 10% of similar patients commit a violent act). The frequency estimates of the opinion of other experts (treatment 1) produced significantly higher estimates as compared to the probability estimates (treatment 2). Thus, frequency estimates appear to be more *affective*.

These results are in conformity with the empirical results of Denes-Raj and Epstein (1994). Here, subjects were given a prize for drawing a red bean from two different urns whose composition could be seen in a picture. The majority of the subjects chose the urn with the higher number of red beans (e.g., 10 out of 100) rather than the urn with the higher percentage of red beans (e.g., 2 out of 10). Images of a higher number of red beans are more affective.

19.6 Anchoring and adjustment

It is an extremely robust and reliable empirical finding that the inferences made by people are heavily dependent on an *initial suggestion* or *anchor*. The anchor could be informative or entirely uninformative. This phenomenon is known as *anchoring*. Unlike some of the other judgment heuristics, we can numerically measure the precise degree of anchoring. Early studies of the anchoring effect first appeared in psychophysics, which studies the relation between stimulus and sensation. For instance, the perceived duration of a target sound depended on a long or short sound (anchor) that subjects had previously been exposed to; for a brief review of this literature, see LeBoeuf and Shafir (2006).

The field was revolutionized by the work of Tversky and Kahneman (1974), who departed from the practice in psychophysics by focusing on numeric anchors. Their method has since become the established norm for experiments on anchoring. In the typical experiment, subjects are given a pair of questions. The first question asks them to make a *comparative judgment*. Here, subjects indicate if a target quantity, x_T , is greater or lower than some anchor, x_0 . In a second question, subjects are asked for an *absolute judgment*, their best estimate of the target value, x_T . Anchoring arises if in the second question, the estimated target value is influenced by the anchor, x_0 .

19.6.1 Empirical evidence for anchoring

In a celebrated experiment, Tversky and Kahneman (1974) rigged a wheel of fortune with the numbers 1–100 to stop at either of the two numbers: 10 or 65. Subjects had to write down the number that the wheel stopped at. They were then asked the following two comparative and absolute judgment questions.

Comparative judgment question: Is the percentage of African nations among UN members larger or smaller than the number you just wrote?

Absolute judgment question: What is your best guess of the percentage of African nations among UN members?

The answers to the second question were found to be anchored too closely on the irrelevant number that came up on the wheel. Those who observed the number 10 (respectively, 65) in the comparative judgment question, answered 25% (respectively, 45%) in the absolute judgment question. Kahneman (2011, p. 120) writes: “We were not the first to observe the effects of anchors, but our experiment was the first demonstration of its absurdity.” These results immediately spawned great interest and the results were replicated in a large number of contexts and shed light on many disparate phenomena.

The anchoring effect can be numerically measured. Kahneman (2011) reports the following experiment. Subjects were asked the following two questions in a San Francisco Exploratorium.

1. Is the height of the tallest redwood more or less than x_0 feet?
2. What is your best guess about the height of the tallest redwood?

Two different values of the anchor, x_0 , measured in feet, were given to two different groups of subjects; a high anchor value, $x_0^H = 1200$, and a low anchor value, $x_0^L = 180$. The answers to the target values in the second question were, respectively, $x_T^H = 844$ and $x_T^L = 282$. Using a method proposed by Jacowitz and Kahneman (1995), we can calculate the *anchoring index* as $\frac{x_T^H - x_T^L}{x_0^H - x_0^L} \times 100 = 55\%$.

In Mussweiler and Strack (1999), in the comparative condition, subjects were asked: “Is the mean temperature in Antarctica in winter higher or lower than $-x_0$ °C?” The initial temperature

was intended to serve as an anchor. Two anchors were used, a high anchor, -17°C , and a low anchor, -243°C . This is followed by an elicitation of the absolute judgment of subjects of the mean winter temperature in Antarctica. The answer in the absolute judgment condition was closely correlated with the anchor values of $-x_0^{\circ}\text{C}$.

Anchors need not always be provided by the experimenter. They can be self-generated. As an example, consider the experiment in Tversky and Kahneman (1974) in which two groups of subjects was asked to take 5 seconds to find the answers to two calculations: $1 \times 2 \times \dots \times 8$ and $8 \times 7 \times \dots \times 1$. Due to the time pressure, the first few numbers are likely to serve as anchors. If this reasoning is correct, then the first calculation gives a low anchor and the second gives a high anchor. The correct answer to the identical calculations is 40,320. However, the median answers to the ascending and descending sequences are, respectively, 512 and 2250, which is consistent with an anchoring explanation.

The finding of Tversky and Kahneman (1974) that implausible anchors influence the outcomes has been replicated widely. Strack and Mussweiler (1997) asked subjects in the comparative judgment question if Mahatma Gandhi was younger than 140 years (high anchor) or older than 9 years (low anchor); both anchors are implausible. However, the respective answers by the two groups of subjects in the absolute judgment question (Mahatma Gandhi's age) were 67 years and 50 years, demonstrating a strong anchoring effect. Mussweiler and Strack (2000a) report similar findings in which subjects had to guess the length of the Mississippi river in the absolute judgment task. In the comparative judgment tasks, two anchors were given. A short but plausible anchor, 2000 miles, and a long, but implausible anchor, 30,000 miles. The answers to the absolute judgment question were significantly influenced by the anchors; mean estimates in the low and the high anchor conditions were 3,768 miles and 12,144 miles.

Field evidence also supports anchoring. Northcraft and Neale (1987) showed that the estimated value of properties by estate agents was influenced by the list price, which served as an anchor. Estate agents were given a 10 page booklet with all relevant information about a property. The booklet also contained a list price, which was either \$83,900 (high anchor) or \$65,900 (low anchor). Agents exposed to the high anchor priced the property significantly higher and the *anchoring effect* was 41%. Business school students were only slightly more susceptible to anchoring; the anchoring effect in this case was 48%. The estate agents insisted that the anchor had no effect on their calculations, yet the significant anchoring effect is incompatible with their insistence.

Anchoring may be related to entitlements. In Ariely et al. (2006), two groups of students attend a poetry reading session. The first group is asked if they *would pay* \$2 to attend the session. The second group is asked if they would attend if they *were paid* \$2. The anchor given to the first group suggested that the poetry reading session would give them utility, while the anchor to the second group is suggestive of disutility. Later, both groups were asked if they would attend for free. A third of the subjects in the first group agreed to attend but only 8% in the second group agreed to attend.

In the efficient markets hypothesis (see Chapter 21 for details), the past pattern of prices of securities should not reveal any new information that enables one to predict the future price. Mussweiler and Schneller (2003) use the anchoring heuristic to provide one test of theory. Their subject pool is Business and Economics students at a German university who have had experience of the stock market, ranging from a year to approximately three years. They provide students with background information about a company and the relevant economic conditions, drawn from an actual financial firm.

They also show students either the left or the right panel of Figure 19.4 that plots the stock price of the company over the previous 12 months. The left panel has a salient high price and the right




Figure 19.4 High price is salient (left panel) and low price is salient (right panel). Starting and ending prices are the same in both cases. The average profits over the 12 months are 20% in both cases.

Source: Mussweiler, T. and K. Schneller (2003). "What goes up must come down—how charts influence decisions to buy and sell stocks," *The Journal of Behavioral Finance* 4(3): 121–30. Reprinted by permission of Taylor & Francis Ltd., <http://www.tandfonline.com>

panel has a salient low price. Otherwise the starting and the ending prices as well as the average profits corresponding to both panels are identical. Subjects were asked to state price expectations of the stock in each case and also asked if they would invest in the asset.

The anchoring heuristic suggests that subjects who observe the high salient price will have relatively higher price expectations and are more likely to invest as compared to subjects who are shown the low salient price. These predictions are confirmed. The predicted mean prices when the high and low price, respectively, is salient are 132 and 112; the difference is statistically significant.

19.6.2 Legal and regulatory implications of anchoring

Experienced individuals also exhibit the anchoring effect. The subjects in Englich and Mussweiler (2001) are experienced trial judges with an average experience of 15 years. In one experiment, they were given the case of a fictitious shoplifter who is accused of stealing from a supermarket for the 12th time. Detailed case material was compiled and it included opinions from a psycho-legal expert, and testimonies by the defendant and a witness. A pre-test with experienced legal professions showed that the case material was complete and realistic. In the pre-test, the mean sentence was 5.62 months with a standard deviation of 2.57.

Subjects in the experiment were asked to determine an appropriate sentence. In two different treatments, the case material showed that the prosecutor had asked for a sentence of, respectively, 9 months (high anchor) and 3 months (low anchor). Subjects in the high anchor treatment chose a sentence of 8 months and those in the low anchor treatment chose a sentence of 5 months. The anchoring effect was $\frac{8-5}{9-3} \times 100 = 50.0\%$. Crucially, the authors find that experience does not mitigate the anchoring effect. Legal professionals and experienced judges exhibit a similar bias.

Anchoring has interesting regulatory implications. Consider, for instance, the problem of caps on damages in litigation. Saks et al. (1997) considered two treatments. In the first treatment, damages were capped at the relatively high level of \$250,000 for trial awards for minor injuries. In the second treatment, there were no caps. They found that the average trial award was \$3895 in the uncapped treatment and \$15,718 in the capped treatment; anchoring at a high cap level significantly increased the average trial award.

Babcock and Pogarsky (1999) considered the rate of pretrial settlements with and without caps. They conducted experiments in which subjects in their role as either plaintiffs or defendants received case material on a personal injury lawsuit. They then had to negotiate a settlement with

the other party within 20 minutes. The default award made by a judge, should the bilateral negotiation fail, was capped at \$250,000 in one treatment and left uncapped in another treatment. It was found that the rate of pretrial settlement was higher in the capped treatment.

Pogarsky and Babcock (2001) found that the rate of pretrial settlement is negatively influenced if the damages cap is set very high. They considered a cap of \$1,000,000 in the capped treatment and reduced the severity of the plaintiff's injuries in the case material, relative to Babcock and Pogarsky (1999). This ensured that the cap appeared to any reasonable observer as too high. Unknown to the subjects, an actual trial court judge had opined a settlement of \$325,000, based on the case material. The non-binding settlement cap reduced the settlement rate, as the authors had expected; 79% of the plaintiff–defendant pairs settled in the uncapped treatment, while 67% settled in the capped treatment. Thus, litigants appear to look ahead and anticipate the effect on the award of damages that is caused by high caps.

19.6.3 *Robustness of anchoring*

The anchoring phenomenon is remarkably robust to a wide range of domains, contexts, and frames. These include estimates of price (Mussweiler et al., 2000; Northcraft and Neale, 1987); the probability of a nuclear war (Plous, 1989); the evaluations of lotteries and gambles (Chapman and Johnson, 1994); issues of legal judgment (Chapman and Bornstein, 1996; Englich and Mussweiler, 2001); and first offers in price negotiation (Galinsky and Mussweiler, 2001). Anchoring has been used to explain a range of other judgment heuristics. These include the hindsight bias (Fischhoff, 1975); preference reversals (Lichtenstein and Slovic, 1971); and non-linear probability weighting (Tversky and Kahneman, 1974).

Anchoring is also robust to a wide variety of conditions. These include, clearly uninformative/random anchors (Mussweiler and Strack, 2000b; Tversky and Kahneman, 1974), or anchors known to be very extreme (Chapman and Johnson, 1994; Strack and Mussweiler, 1997). Anchoring appears not to be affected by forewarning the subjects or by the provision of incentives (Tversky and Kahneman, 1974; Wright and Anderson, 1989; Wilson et al., 1996; Chapman and Johnson, 2002). We have already noted that experienced subjects also exhibit the anchoring effect. In conjunction, these findings suggest that humans appear hardwired to use the anchoring heuristic.

19.6.4 *Explanations of anchoring*

In light of the robustness and reliability of the phenomenon, it is important to ask what causes anchoring. Tversky and Kahneman (1974) conjectured an explanation in terms of *insufficient adjustment* from an anchor to a proposed target. Quattrone (1982) and Quattrone et al. (1981) propose a bounded rationality explanation for insufficient adjustment. In this proposal, the individual begins with an *anchor*, x_0 , about some target value, x_T , that is believed to belong to an interval $[a, b]$, where a, b are real numbers. Suppose that $x_0 \notin [a, b]$. The individual then engages in a series of improvements of the anchor using a *satisficing heuristic* and stops whenever he hits one of the boundaries of the interval $[a, b]$. If the guess that $x_T \in [a, b]$ is objectively correct, the individual ends up at a distance $x - a$ or $b - x$ from the actual value, x . This is referred to as insufficient adjustment. However, if $x_0 \in [a, b]$, then no adjustment is needed.

Evidence for the insufficient adjustment hypothesis is given in LeBoeuf and Shafir (2006, 2009) who were inspired by the original experiments in the psychophysics of anchoring. In LeBoeuf and Shafir (2006), subjects were required to draw a line that was 3.5 inches (or 89 mm) long; this corresponds to x_T . They were given one of two anchors; see Figure 19.5. The short anchor, x_0^s , was a line starting from the bottom of the page; subjects were told that this line was shorter than 3.5



Figure 19.5 Anchoring and insufficient adjustment.

Source: *Journal of Behavioral Decision Making* 19(4): 393–406. LeBoeuf, R. and Shafir, E., “The long and short of it: physical anchoring effects.” Copyright © 2006 John Wiley & Sons, Ltd.

inches. Subjects were asked to extend it, so that, in their assessment, it was 3.5 inches (88.9 mm) in length. Another group of subjects were shown a line starting from the top of the page and told that it was longer than 3.5 inches; this is the long anchor, x_0^l . They had to shorten the long line to 3.5 inches. Subjects exposed to x_0^s produced, on average, a line length of 61.2 mm with a standard deviation of 17 mm, and subjects exposed to x_0^l produced a line that had an average length of 74.7 mm with a standard deviation of 15.4 mm.

A second explanation for anchoring is that anchors induce System 1 in the brain to search for anchor-consistent information. In this explanation, the anchor serves a *priming function*. Consider Tversky and Kahneman’s (1974) experiments on anchoring described above in which subjects are asked for the number of African countries in the UN. Suppose that the wheel of fortune comes up with a high number, which is the higher anchor. A high anchor enables one to access information that is consistent with the number of countries in Africa being high. For instance, that Africa has a large size, or Africa has a large coastline.

It is crucial to appreciate that System 1 attempts to evaluate the given information by trying to prove that it is true (Snyder and Swan, 1978; Crocker, 1982; Trope and Bassok, 1982; Klayman and Ha, 1987). As Kahneman (2011) argues, System 1 tries its best to construct a world in which the anchor is the true number. System 2 is deliberative, yet it deliberates on the menu of memories chosen by System 1. Hence, a biased set of memories, evoked by priming, have a ripple effect on the decision made by the individual in the absolute judgment task.

When should we evoke the insufficient adjustment explanation and when should we evoke the priming explanation? Epley and Gilovich (2001) give one possible resolution of this problem. They distinguish between two kinds of anchors. Those that are self-generated and those that are provided by the experimenter. They argue that self-generated anchors are known to be wrong from the beginning, so the individual does not need to engage in retrieving anchor-consistent information to judge its correctness. Rather, the problem is to adjust the initial anchor towards the target, as explained above in the case of insufficient adjustment. However, anchors provided by the experimenter need to be evaluated for correctness by evoking anchor-consistent information.

Consider two examples of self-generated anchors from Epley and Gilovich (2001). (1) Subjects are asked to estimate the freezing point of vodka. In thinking about the answer, participants typically start at 0°C, the freezing point of water, and adjust downwards, knowing that the freezing

point of alcohol is lower than water. (2) Which year was George Washington first elected president? In this case, most Americans seem to begin with the year of the declaration of independence, 1776, and adjust upwards from it.

Epley and Gilovich find that self-generated anchors trigger adjustment, while those provided by the experimenter do not. For instance, in the second question (election year of George Washington), 64% of the subjects were found to engage in anchoring and adjustment when the anchor was self-generated. However, when the anchor was provided by the experimenter in two other questions, mean length of a whale, and mean winter temperature in Antarctica, only 12% and 14%, respectively, engaged in anchoring and adjustment. This reasoning is confirmed in the experimental results of Epley and Gilovich (2006). Furthermore, they find that adjustment entails cognitive effort. When subjects are exposed to cognitive load, say, having to remember a set of numbers, or they are drunk, then they engage in insufficient adjustment.

19.7 Base rate neglect and conservatism

The use of Bayes' rule underpins all aspects of neoclassical economics. In this section, we examine the evidence on Bayes' rule. Suppose that we have n mutually exclusive random variables $X_1, X_2, \dots, X_n$ such that the sample space $X = X_1 \cup X_2 \cup \dots \cup X_n$. Then Bayes' rule implies the following computation for any event $A \subseteq X$ such that $P(A) > 0$,

$$P(X_i | A) = \frac{P(A | X_i)P(X_i)}{P(A)} = \frac{P(A | X_i)P(X_i)}{\sum_{j=1}^n P(A | X_j)P(X_j)}. \quad (19.8)$$

In (19.8), $P(X_i)$ and $P(A)$, $i = 1, \dots, n$, are referred to as the *base rates*.

19.7.1 Base rate neglect

An important empirical finding is that base rates have very low salience in calculations of $P(X_i | A)$. Meehl and Rosen (1955) is an early study that found underweighting of base rates by clinical psychologists. Consider the following well-known *cab problem* introduced by Kahneman and Tversky (1972); modified versions of the problem are considered in Bar-Hillel (1980) and Tversky and Kahneman (1980).

There are only two cab companies in the city, Green and Blue; 85% of the cabs are Green. There was an accident last night. A witness comes forward to testify that the cab involved in the accident was Blue. In similar conditions, the reliability of the witness is 80%, i.e., the probability that the witness gets it wrong is 20%. What is the probability that the actual cab involved in the accident was Blue?

Let X_1 = Blue cab, X_2 = Green cab, and let the event, "identification by the witness that the cab is Blue" be A . Then, using (19.8),

$$P(X_1 | A) = \frac{P(A | X_1)P(X_1)}{P(A | X_1)P(X_1) + P(A | X_2)P(X_2)}. \quad (19.9)$$

Using $P(A | X_1) = 0.8$ (probability of a correct identification), $P(A | X_2) = 0.2$ (probability of an incorrect identification), $P(X_1) = 0.15$ (base rate), and $P(X_2) = 0.85$ (base rate) we get

$$P(X_1 | A) = 0.414.$$

By contrast, the median and the modal response was $P(X_1 | A) = 0.8$. Notice that 0.8 is also the probability with which the witness correctly identifies a Blue cab as the one involved in the accident. In particular, individuals appear to underweight base rates, which is not consistent with Bayes' rule. They ignore the fact that only 15% of the taxis are actually Blue. The problem of underweighting of base rates diminishes but does not disappear when greater salience is given to base rates in the design of experiments (Ajzen, 1977; Bar-Hillel, 1980).¹⁴

Consider an adaptation of the problem that Casscells et al. (1978) gave to students at Harvard Medical School. Suppose that a police Breathalyzer test discovers false drunkenness in 5% of the cases when the driver is sober. However, the Breathalyzers are always able to detect a truly drunk person with complete certainty. In the general population, a fraction 1/1000 of drivers engage in driving while drunk. Suppose that a driver is checked at random, and takes a Breathalyzer test, which shows that the driver is drunk. How high is the probability he or she really is drunk? The modal response was 95%, the mean response was 56% and a sixth of the students gave the correct answer, which is approximately 2%. Thus, even very educated and clever subjects underweight the base rate in relatively simple problems.

In some cases, underweighting of base rates does not take place. Consider, for instance, the following variant of the cab problem from Kahneman (2011). "There are only two cab companies in the city, Green and Blue, who own an equal number of cabs. 85% of the cabs involved in accidents are Green." The rest of this problem is identical to the earlier cab problem following the instruction: "There was an accident last night" . . . This, and the earlier cab problem, are mathematically equivalent. Yet, one finds that subjects give attention to the base rate in the second problem. The median response is that the cab involved in the accident was blue is 0.60. Indeed this appears as an average of the correct answer, 0.41, and the reliability of the witness, 0.80.

What accounts for the relative salience of base rates in this case? The answer depends on the distinction made by Ajzen (1977) between *causal* and *incidental* base rates. A causal base rate conveys information on some causal factors relevant to the problem, while incidental base rates do not. For instance, in the original cab problem, the base rates are incidental. However, in the modified cab problem, the base rates are causal because they suggest to the decision maker that drivers of Green cabs are particularly reckless. This is because despite identical numbers of Blue and Green cabs, Green cabs are involved in 85% of the accidents. Hence, the individual assigns a lower probability to Blue cabs being involved in the accident, which curbs their tendency to underweight the base rate of Blue cabs.

Bar-Hillel (1980) showed that subjects may combine information on base rates with other non-causal information to mitigate base rate neglect. For instance, in one modification of the cab problem, subjects were told that the hit-and-run cab was equipped with an intercom and intercoms are installed on 80% of the Green cabs and 20% of the Blue cabs. The mathematical structure of this problem is identical to the original cab problem, yet the median response was 0.48, not far from the Bayesian answer of 0.41.

A variety of factors have been identified that reduce base rate neglect. Individuals appear to pay more attention to base rate information, when they are instructed to think like statisticians rather than clinical psychologists (Schwarz, Strack, Hilton, and Naderer, 1991). Individuals with higher measured intelligence also underweight base rates less (Stanovich and West, 2000). Girotto and

¹⁴ For more evidence on the neglect of base rates, see Grether (1980) who found that likelihoods are more salient than base rates, although the latter were not altogether ignored. See also Grether (1992) for an extension of his earlier work.

Gonzalez (2001) find that if the attention of subjects is drawn to the computations required in a Bayes' rule calculation, then base rate neglect falls. Sloman et al. (2003) find that presenting the information in a diagram that highlights set inclusion also reduces base rate neglect.

Empirically it has been observed that presenting the information in a frequency format rather than a probability format reduces base rate neglect (Brase, 2002b; Cosmides and Tooby, 1996; Gigerenzer and Hoffrage, 1995; Pinker, 1997). Consider the following problem of *base rate neglect* used originally in Eddy (1982) and then in the experiments of Gigerenzer and Hoffrage (1995). The problem is expressed in probability terms as follows.

The probability of breast cancer is 1% for a woman at age forty who participates in routine screening [base-rate]. If a woman has breast cancer, the probability is 80% that she will get a positive mammography [hit-rate]. If a woman does not have breast cancer, the probability is 9.6% that she will also get a positive mammography [false-alarm rate]. A woman in this age group had a positive mammography in a routine screening. What is the probability that she actually has breast cancer?

The Bayesian solution is readily found. Let X be the event "positive mammography" and Y the event "breast cancer" and let $\sim X, \sim Y$ be negations of these events. Then

$$P(Y | X) = \frac{P(X | Y)P(Y)}{P(X | Y)P(Y) + P(X | \sim Y)P(\sim Y)} = \frac{0.008}{0.008 + 0.096(0.99)} = 0.0776.$$

Thus, the Bayesian response is that conditional on getting a positive mammography, the probability that the individual has breast cancer is 7.76%. In Eddy's (1982) original study, less than 5% of the respondents gave answers that conformed to the Bayesian solution. Gigerenzer and Hoffrage (1995) tested the same problem in the probability format and the natural frequency format. In the probability format they found that the choices of 16% of the respondents conformed to the Bayesian solution, but in the natural frequency format, 46% of the responses conformed to Bayes' rule.

A sample of the studies that use the frequency format is shown in Table 19.4, drawn from Barbey and Sloman (2007). The second column shows that when the problem is presented in a probabil-

Table 19.4 Percentage of responses consistent with Bayes' rule in different empirical studies. Sample sizes in parenthesis.

Study	Information format and judgment domain	
	Probability	Frequency
Cascells et al. (1978)	18(60)	—
Cosmides & Tooby (1996; Exp. 2)	12 (25)	72 (25)
Eddy (1982)	5 (100)	—
Evans et al. (2000; Expl)	24 (42)	35 (43)
Gigerenzer (1996b)	10 (48)	46 (48)
Gigerenzer and Hoffrage (1995)	16 (30)	46 (30)
Macchi (2000)	6 (30)	40 (30)
Sloman et al. (2003; Exp 1)	20 (25)	51 (45)
Sloman et al. (2003; Exp 1b)	—	31 (48)

Source: Barbey and Sloman, 2007.

ity format, the percentage of subjects conforming to Bayes' rule is relatively small. However, there is much greater conformity with Bayes' rule when the information is presented in a frequency format (third column). The Cosmides and Tooby (1996) result is an outlier that Sloman et al. (2003) and Evans et al. (2000) could not replicate. Despite the rhetoric, in the presence of a frequency format, less than half of the subjects conform to Bayes' rule in this cross section of studies.

It is also likely that other factors could facilitate conformity with Bayesian reasoning when the natural frequency format is given (Cosmides and Tooby, 1996). These factors include, in particular, monetary incentives and controlling for the ability levels of students, e.g., hiring subject pools at top universities. It is argued that these factors lead to statistically significant differences in the facilitation role of natural frequencies relative to the probability format (Brase et al., 2006). But this also casts doubt on the evolutionary explanation, which lies at the heart of the facilitating role of natural frequency formats. At best, these results are suggestive of the relation between cognitive effort and cognitive ability with the facilitation role of natural frequencies in improving conformity with Bayes' rule.

When data is presented in natural frequencies, but the set inclusion relation is not made obvious (although it can be easily computed from the data), then the facilitation role of natural frequencies drops significantly (Giroto and Gonzalez, 2001). Furthermore, Brase (2002a) finds that the probability and the natural frequency format are both perceived to be equally very clear unless the probabilities are of a very low magnitude. This is difficult to reconcile with the claimed facilitating role of the natural frequency format.

Based on the evidence, the case for the natural frequency format in reversing base rate neglect appears overstated.

Why might problems expressed in terms of natural frequency facilitate Bayesian reasoning? Some of the competing theories considered by Barbey and Sloman (2007) are as follows. In the *mind as a Swiss army knife* view (Cosmides and Tooby, 1996; Gigerenzer and Selten, 2001), the brain comprises several modules. Each module is responsible for a specific set of functions and is not under the voluntary control of the individual. In this view, there is a specific module, with evolutionary origins, designed to process natural frequencies, but not probabilities. One natural advantage of the frequency format is that it preserves the sample size and makes transparent the subset relation between two sets (e.g., 1 out of every 100 women have breast cancer).

A second view is that the brain has a *natural frequency algorithm* (Gigerenzer and Hoffrage, 1995). This view makes no assumptions about the modularity of brain architecture. However, for all practical purposes, this is fairly similar to the first view. Barbey and Sloman's (2007) preferred explanation is based on a System 1, System 2 distinction. In this view, the frequency format is more consistent with the natural inputs required by System 2 to process information. As the peer commentary in the same issue of *Behavioral and Brain Sciences* makes clear, however, there is no consensus among psychologists on this issue.

19.7.2 *Conservatism*

Conservatism refers to the underweighting of the likelihood of a sample. We illustrate this using a problem in Camerer (1995), which is originally due to Edwards (1968). Suppose that there are two urns A and B, each of which contains only red and blue balls. Urn A has 7 red and 3 blue balls and urn B has 3 red and 7 blue balls. It is common knowledge that nature is equally likely to choose one of the two urns, so $P(A) = 0.5$ and $P(B) = 0.5$. A sample of 12 balls is drawn *with replacement* from one of the urns, but subjects do not know which one.

Suppose that the number of red balls in a sample of size 12 is $x \in \{0, 1, \dots, 12\}$. Consider the sample, 8 red and 4 blue balls, which does not exactly represent the proportions in any of the two urns. Hence, we cannot use the representativeness heuristic. Let p_A and p_B be the respective probabilities of drawing a red ball from urns A and B. Analogous to (19.2), the conditional probability of drawing x red balls in a sample of 12 balls from urn $i = A, B$ is

$$P(x | i) = \binom{12}{x} p_i^x (1 - p_i)^{12-x}; i = A, B \text{ and } x \in \{0, 1, \dots, 12\}. \quad (19.10)$$

The Bayesian posterior belief that the sample came from urn A is given by

$$P(A | x = 8) = \frac{P(x = 8 | A) P(A)}{P(x = 8 | A) P(A) + P(x = 8 | B) P(B)}. \quad (19.11)$$

We already know that $P(A) = 0.5$, $P(B) = 0.5$. Since $p_A = 0.7$, $p_B = 0.3$, we get from (19.10) that $P(x = 8 | A) = 0.231$, $P(x = 8 | B) = 0.008$. Substituting these numbers in (19.11), we get that $P(A | x = 8) = 0.967$. However, the typical response reported in Camerer (1995) is 0.7 to 0.8. Hence, subjects in experiments appear too conservative about sample likelihoods.

19.7.3 A reconciliation of conservatism and base rate neglect

It appears contradictory that in some cases the base rate is underweighted (as in the cab problem), while in others the sample information is underweighted (conservatism). We need a model that identifies circumstances when we are confident (underweighting base rates), and underconfident (conservatism). Griffin and Tversky (1992) show that these two phenomena can be reconciled by distinguishing between the *strength* and the *weight* of evidence. Consider two illustrative examples of these concepts.

1. Suppose that one is trying to evaluate a reference letter for a graduate student, written by an academic referee. The strength of the letter is the information that it provides about the student, for instance, how warm and positive the letter is, while the weight is the credibility/standing of the referee.
2. Suppose that we toss a coin n times to determine if it is biased. The strength of the evidence is the number of times it shows up as heads while the weight is the sample size.

While it is not always possible to neatly separate the strength and weight of evidence in every case, in many cases one may be able to do so. Griffin and Tversky (1992) argue that people give undue importance to strength as compared to the weight of the evidence, relative to what would be prescribed by standard statistical techniques. In particular, people begin by giving undue importance to the strength of the evidence and then adjusting insufficiently for the weight of the evidence. For instance, in the first example, we often see how warm the reference letter is and then adjust for the credibility of the writer. We feel underconfident when there is a moderately positive letter from a very credible referee. An overconfident judgment arises when strength is high relative to weight. Conversely, an underconfident judgment gives relatively more importance to the weight, particularly when the strength is low.

Consider the following experiment that Griffin and Tversky conducted with 35 students. A coin is tossed n times and the number of heads (H) and tails (T) is noted. Subjects are told that the bias of the coin is $3/5 = 0.6$, but they are not told the direction of the bias. Hence, subjects believe

that either $P(H) = 3/5$, or $P(T) = 3/5$. Denote the events, “head-biased coin” and “tail-biased coin,” respectively, by H_b and T_b . Prior probabilities of these mutually exclusive and exhaustive events are 50 : 50,

$$P(H_b) = P(T_b) = 0.5. \quad (19.12)$$

In the experiments, the sample size, n , varies from 3 to 33, while the number of heads varies from 2 to 19; see Table 19.5 for the results. Consider a sample of size n in which there are h_n heads and $n - h_n$ tails. Denote this sample by h_n . The posterior distribution $P(H_b \mid h_n)$ that the coin is head-biased, conditional on the sample h_n , is given by

$$P(H_b \mid h_n) = \frac{P(h_n \mid H_b)P(H_b)}{P(h_n \mid H_b)P(H_b) + P(h_n \mid T_b)P(T_b)}. \quad (19.13)$$

Using the formula for the binomial distribution, we can compute $P(h_n \mid H_b)$ and $P(h_n \mid T_b)$ and rewrite (19.13) as

$$P(H_b \mid h_n) = \frac{\binom{n}{h_n} \left(\frac{3}{5}\right)^{h_n} \left(\frac{2}{5}\right)^{n-h_n}}{\binom{n}{h_n} \left(\frac{3}{5}\right)^{h_n} \left(\frac{2}{5}\right)^{n-h_n} + \binom{n}{h_n} \left(\frac{2}{5}\right)^{h_n} \left(\frac{3}{5}\right)^{n-h_n}}. \quad (19.14)$$

The first three columns in Table 19.5 give the experimental data, and the final column gives the computation of posterior probabilities, based on (19.14).

In an analogous manner, we get

$$P(T_b \mid h_n) = \frac{P(h_n \mid T_b)P(T_b)}{P(h_n \mid T_b)P(T_b) + P(h_n \mid H_b)P(H_b)}. \quad (19.15)$$

Table 19.5 Probability that the coin is head-biased, conditional on the observations.

n	h_n	$n - h_n$	$P(H_b \mid h_n)$
3	2	1	0.60
3	3	0	0.77
5	3	2	0.60
5	4	1	0.77
5	5	0	0.88
9	5	4	0.60
9	6	3	0.77
9	7	2	0.88
17	9	8	0.60
17	10	7	0.77
17	11	6	0.88
33	19	14	0.88

Source: Griffin and Tversky (1992).

Dividing (19.13) by (19.15) and using (19.12), we get

$$\frac{P(H_b | h_n)}{P(T_b | h_n)} = \frac{P(h_n | H_b)}{P(h_n | T_b)}. \quad (19.16)$$

Recall that the sample h_n denotes h heads out of a sample of n iid draws from either a heads-biased coin ($P(H) = 0.6$) or a tail-biased coin ($P(T) = 0.6$). Since the successive tosses of the coin are independent, the likelihoods of the sample under, respectively, the heads-biased coin and the tail-biased coin, are given by

$$P(h_n | H_b) = (0.6)^{h_n} (0.4)^{n-h_n}; P(h_n | T_b) = (0.4)^{h_n} (0.6)^{n-h_n}. \quad (19.17)$$

Substituting (19.17) in (19.16) and denoting by $t_n = n - h_n$, the number of tails in the sample, we get

$$\frac{P(H_b | h_n)}{P(T_b | h_n)} = \left(\frac{0.6}{0.4} \right)^{h_n - t_n}.$$

Taking logs on both sides and multiplying and dividing the right hand side by n , we get the following form for Bayes' rule.

$$\log \left(\frac{P(H_b | h_n)}{P(T_b | h_n)} \right) = n \left(\frac{h_n - t_n}{n} \right) c; \text{ where } c = \log \left(\frac{0.6}{0.4} \right), \text{ a constant.} \quad (19.18)$$

We can now use (19.18) to further clarify the notions of strength and weight, using terms on the right hand side.

1. The first term, n , is the *weight* of the evidence.
2. The second term, the difference between the proportions of heads and tails in the sample, is the *strength* of the evidence for H_b and against T_b ; the experiment only considers the cases such that $h_n > t_n$.
3. The third term, c , is the *discriminability* of the evidence.

Taking logs on both sides of (19.18) we get

$$\log \log \left(\frac{P(H_b | h_n)}{1 - P(H_b | h_n)} \right) = \log n + \log \left(\frac{h_n - t_n}{n} \right) + \log c. \quad (19.19)$$

Griffin and Tversky (1992) estimate a regression of the following form:

$$y = \beta_0 + \beta_1 \log n + \beta_2 \log \left(\frac{h_n - t_n}{n} \right) + \beta_3 \log c + u \quad (19.20)$$

where $y = \log \log \left(\frac{P(H_b | h_n)}{P(T_b | h_n)} \right)$, u is a random error, and $\beta_0, \beta_1, \beta_2, \beta_3$ are the regression coefficients. To be consistent with Bayes' rule, (19.18), the predicted regression coefficients are: $\beta_1 = \beta_2 = 1$. However, for 30 out of 35 subjects, $\beta_2 > \beta_1$ and the difference is statistically significant. Across subjects, the median ratio $\frac{\beta_2}{\beta_1} = 2.2$, i.e., subjects weighted the strength more than twice as much as the weight.

Since the conditional probability function is also a probability function, $P(H_b | h_n) + P(T_b | h_n) = 1$. Fix $P(H_b | h_n)$ at some number in the open interval $(0, 1)$, so that the LHS of (19.19) equals a constant, $\bar{y}$. An example is the value of $P(H_b | h_n) = 0.77$ for $n = 3$ and 9 (see Table 19.5). Thus, (19.19) can be written as

$$\log \left(\frac{h_n - t_n}{n} \right) = (\bar{y} - \log c) - \log n. \quad (19.21)$$

We now ask, what combinations of strength, $\frac{h_n - t_n}{n}$, and weight, n , satisfy equation (19.21), given that $\bar{y} - \log c$ is a fixed number? A curve that expresses such combinations is known as an *equal-support line*. In $(\log n, \log \left(\frac{h_n - t_n}{n} \right))$ space, the equal-support line based on (19.21) has a slope -1 and an intercept that is increasing in $P(H_b | h_n)$ (because $\bar{y}$ is increasing in $P(H_b | h_n)$). So, for instance, if we fixed $P(H_b | h_n) = 0.77$, then, using Table 19.5, (19.21) is satisfied for the following two (strength, weight) pairs: $\left(\frac{h_n - t_n}{n}, n \right) = (1, 3), \left(\frac{1}{3}, 9 \right)$.

In Figure 19.6, we plot two equal-support lines that are consistent with Bayesian updating (i.e., $\beta_1 = \beta_2 = 1$). These are shown as the two dotted lines. The upper one corresponds to a higher value of $\bar{y} = y_H$, relative to the lower one, $\bar{y} = y_L < y_H$.

We can also derive the equal support lines from the estimated regression. Rewrite (19.20) as follows.

$$\log \left(\frac{h_n - t_n}{n} \right) = \frac{1}{\beta_2} (y - \beta_0 - \beta_3 \log c - u) - \frac{\beta_1}{\beta_2} \log n. \quad (19.22)$$




Figure 19.6 Determining the relative importance of the strength and the weight of evidence.

Source: Reprinted from *Cognitive Psychology* 24(3): 411–35. Griffin, D. and Tversky, A. (1992). “The weighing of evidence and the determinants of confidence.” Copyright © 1992, with permission from Elsevier.

Since, empirically, subjects place greater salience on strength, relative to weight, $\frac{\beta_1}{\beta_2} < 1$, so the slope of the estimated equal-support line is flatter. Two such estimated equal-support lines are drawn for $y = y_H$ and $y = y_L$; these are shown as the two continuous lines.

Consider the intersection points e_H and e_L , where the Bayesian line intersects the estimated line. To the left of these intersection points (for any level of weight, $\log n$) the estimated strength is lower than the correct statistical (or Bayesian) strength. Individuals are said to be *underconfident* in this region. Conversely, to the right of these intersection points, estimated strength is greater than the Bayesian strength. Individuals are said to be *overconfident* in this region. Conservatism is a sort of *underconfidence* in the sample likelihood. On the other hand, base rate neglect arises when individuals are *overconfident* of the strength of the evidence. Hence, points such as e_H , e_L divide the space into regions where conservatism and base rate neglect arise naturally.

Consider an individual engaged in the stock market. When he hears isolated bits of news, he views the news as having low strength, so he does not find this information very useful (although it might have high weight; for instance, the news might be a report of an assessment by a leading expert). On the other hand, when the individual hears about a small sequence of good news about a company selling some stock, he might assign much higher strength to it as compared to a Bayesian. The individual might overreact to a sequence of good news based on an overweighted strength. However, the news might have little weight in predicting the future performance of the company.

19.8 Hindsight bias

Suppose that an individual is asked at time t to predict the outcome of an uncertain event, X , that will take place at some time $t + j$, $j > 0$. Denote the information set at time t by I_t . Then the conditional prediction, denoted by $E[X | I_t]$, reflects the individual's *predictive judgment*. Suppose that at time $t + j$, the realization of the random variable is x . Posterior to the observation of x , the individual is asked about the remembered prediction or the *postdictive judgment*. This is given by $E[E[X | I_t] | I_{t+j}]$, where I_{t+j} is the date $t + j$ information set, and $I_t \subseteq I_{t+j}$. In neo-classical economics, the typical assumption is $E[X | I_t] = E[E[X | I_t] | I_{t+j}]$. An individual is said to suffer from *hindsight bias*, or *creeping determinism* (Fischhoff, 1975), if the predictive and postdictive judgments differ, i.e.,

$$E[X | I_t] \neq E[E[X | I_t] | I_{t+j}], \quad (19.23)$$

and the postdictive judgment is biased in favour of the actual outcome. For instance, Camerer et al. (1989) incorporate hindsight bias by the following simple rule

$$E[E[X | I_t] | I_{t+j}] = \alpha x + (1 - \alpha)E[X | I_t]; \alpha \in [0, 1]. \quad (19.24)$$

The individual has no hindsight bias if $\alpha = 0$ and is fully-biased if $\alpha = 1$. Hindsight bias is sometimes termed as the “I knew it all along” effect; this is particularly obvious when $\alpha = 1$. Hindsight bias is quite distinct from learning. In learning, one learns to make better predictions of future events. However, under hindsight bias, the relevant issue is remembered past predictions, i.e., postdictive judgments.

19.8.1 *Empirical evidence on the hindsight bias*

In his pioneering work, Fischhoff (1975) ran six experiments in a between-subjects design; each experiment has a Before and After group. In each experiment, the Before group read a brief 150 word account of a historical or clinical event and then had to rate the probabilities of four possible mutually exclusive outcomes of which only one outcome actually occurred; this captures predictive judgment. The After group then read the same 150 word account, and was told that one of the four outcomes had occurred; all four outcomes were used in this stage, but each experimental group was given only one outcome and told that this was the real outcome. Based on this information, the After group was required to make a postdictive judgment.

Each of the six experiments used four events, A, B, C, and D, to construct the 150 word accounts. The events were chosen so that subjects were unlikely to know the actual outcome. For instance, event A used a historical account that was based on the British–Gurkha struggle. It was presented as follows (p. 305):

For some years after the arrival of Hastings as governor-general of India, the consolidation of British power involved serious war. The first of these wars took place on the northern frontier of Bengal where the British were faced by the plundering raids of the Gurkhas of Nepal. Attempts had been made to stop the raids by an exchange of lands, but the Gurkhas would not give up their claims to country under British control, and Hastings decided to deal with them once and for all. The campaign began in November 1814. It was not glorious. The Gurkhas were only some 12000 strong; but they were brave fighters, fighting in territory well-suited to their raiding tactics. The older British commanders were used to war in the plains where the enemy ran away from a resolute attack. In the mountains of Nepal it was not easy even to find the enemy. The troops and transport animals suffered from the extremes of heat and cold, and the officers learned caution only after sharp reverses. Major-General Sir D Ochterlony was the one commander to escape from these minor defeats.

The possible outcomes offered were: (1) British victory, (2) Gurkha victory, (3) military stalemate with no peace settlement, and (4) military stalemate with a peace settlement. The factually correct outcome was (3) but the After subjects were informed of one of the factual outcomes (1)–(4) as the correct one.

The average probabilities for each of the outcomes, 1–4, in the four events, A–D, are shown in Table 19.6. For each event, the first row gives the predictive judgment of the Before group. The next four rows give the postdictive judgments of the After groups, for each of the outcomes, 1–4, when told that outcomes 1, 2, 3, and 4, respectively, were the real outcome.¹⁵ Hindsight bias arises if the before probabilities (first row) and the after probabilities (along the diagonal of the matrix) differ and, in particular, the after probabilities are higher. This was true in all cases, and in 22/24 of the cases, the after–before difference was statistically significant at the 2.5% level. The mean difference between the after and before probabilities was 10.8%. At the level of the individual subjects, 70% of the After subjects assigned higher probabilities relative to the mean before probabilities.

¹⁵ Just to make sure that there is no confusion about Table 19.6, consider Event A (British–Gurkha struggle). On average, 33.8% of the before group assigns outcome 1 as the most likely outcome (predictive judgment). However, 57.2% of the after group that has been told that the outcome was 1, believes that the before group must have chosen outcome 1 as the most likely outcome (postdictive judgment); only 14.3% in the same group believe that the before group chose outcome 2. Similarly, 21.3% of the before group chooses outcome 2. However, 38.4% of the after group that was given outcome 2 as the actual outcome, believed that outcome 2 was chosen by the before group (postdictive judgment).

Table 19.6 Mean probabilities assigned to each of four outcomes, 1–4, in the four events, A–D. The actual outcomes for events A, B, C, and D were, respectively, 1, 1, 4, and 2.

Experimental group	n	Outcome provided	Outcome evaluated			
			1	2	3	4
Event A: British–Gurka struggle (English-speaking subjects)						
Before	20	None	33.8	21.3	32.3	12.3
After	20	1	57.2	14.3	15.3	13.4
	20	2	30.3	38.4	20.4	10.5
	20	3	25.7	17.0	48.0	9.9
	20	4	33.0	15.8	24.3	27.0
Event B: Near-riot in Atlanta (subjects with knowledge of statistics)						
Before	20	None	11.2	30.8	43.8	14.2
After	20	11	30.6	25.8	23.3	20.3
	20	2	5.5	51.8	24.3	18.5
	20	3	2.9	23.9	50.8	21.4
	20	4	16.7	31.9	23.4	27.9
Event C: Mrs Dewar in therapy						
Before	19	None	26.6	15.8	23.4	34.4
After	13	11	43.1	13.9	17.3	25.8
	17	2	26.5	23.2	13.4	36.9
	16	3	30.6	14.1	34.1	21.3
	17	4	21.2	10.2	22.6	46.1
Event D: George in therapy						
Before	17	None	27.4	26.9	39.4	6.3
After	18	11	33.6	20.8	37.8	8.0
	18	2	22.4	41.8	28.9	7.1
	20	3	20.5	22.3	50.0	7.3
	17	4	30.6	19.5	37.7	12.3

Source: Fischhoff (1975).

In order to address potential concerns about the between-subjects design of the earlier study, Fischhoff and Beyth (1975) used a within-subjects design. They asked Israeli subjects to make predictions about a set of 15 possible outcomes on the eve of President Nixon's visit to China and the USSR. Prior to the visit, the subjects filled in a prediction memory questionnaire. Following the actual visit of the President, subjects were asked to make postdictive judgments. Significant hindsight bias was discovered in the postdictive judgments.

In the context of the 9/11 tragedy in the US, Kahneman (2011) cites public reaction when it was revealed that on July 2001 the CIA had come across information to suggest that al-Qaeda might be planning an attack on the US. In this case, the CIA did not pass on the information to the President but did pass it on to the National Security Advisor. A former executive editor of the *Washington Post* wrote (p. 294): "It seems elementary to me that if you've got the story that is going to dominate history you might as well go right up to the President." However, in July 2001 nobody knew that this was going to make history.

A similar case of 20–20 judgment in hindsight is the proliferation of claims by experts following the 2008 financial crises that they had predicted the crash. From time to time, experts predict that a current economic situation may lead to a recession. However, in 2006/7 there was no widespread agreement among experts about the scale of the impending crash or certainty about the failure of Lehman Brothers. Kahneman (2011) infers from this that the crash was not *knowable*. Yet the postdictive judgment of many experts is that they knew it all along; this is an illustration of the *illusion of understanding*.

The hindsight bias has been found for all age groups from 3 to 95 (Bernstein et al., 2011). It has been found in many different contexts such as medical diagnosis (Arkes et al., 1988), accounting decisions (Anderson et al., 1993), football (Roese and Maniar, 1997), legal liability (Hastie et al., 1999; Goodwill et al., 2010), terrorist attacks (Fischhoff et al., 2005), the verdict in the O.J. Simpson criminal trial (Bryant and Brockway, 1997), and labor disputes involving firefighters (Pennington, 1981).

For attempts to debias subjects who suffer from hindsight bias, see Arkes (1991), Fischhoff (1982), Larrick (2004), and Harley (2007). The meta-study by Guilbault et al. (2004) suggests that expertise does not shield decision makers from the hindsight bias. Expertise is likely to mitigate a bias if experts receive continual feedback (weather forecasters, chess players, insurance firms) as compared to cases where less frequent feedback is available, or it is ambiguous (judges, stockbrokers, political experts); see, for instance, Kahneman and Klein (2009), McKenzie et al. (2008), and Tetlock (2006).

19.8.2 *Hindsight bias in legal situations*

The issue of hindsight bias is particularly important in legal situations. Consider a defendant who can take several possible actions; each action leads to a possible harm for others with some probability. Suppose that the chosen action of the defendant did indeed cause harm to a plaintiff, who then sues the defendant. A judge or a jury must now form a postdictive judgment of the reasonableness of the action taken by the defendant. In the presence of hindsight bias, the postdictive judgment is biased, so the defendant is more likely to be found guilty. Thus, liability is likely to be assigned in a biased manner, even if the defendant took the most reasonable action.

If potential defendants are forward looking, then anticipating potential future liability, they may take excessive precautions and engage in excessively risk averse behavior. For instance, doctors may prescribe more tests than they otherwise would have, or avoid minor surgical procedures, or be willing to pay for even more comprehensive and expensive insurance, which feeds into higher treatment costs.

Rachlinski (1998) argues that the legal system recognizes well the presence of hindsight bias and takes reasonable steps to mitigate it, as far as possible. For instance, in medical malpractice lawsuits, reliable ex-ante judgment of reasonable care is available; in this case, courts simply use this judgment as its legal postdictive judgment. Corporate law protects directors of corporations from liability for negligent business decisions. This eliminates the possibility of hindsight bias in legal judgment in this domain, enabling corporations to operate in an unfettered environment; this has its pros (fewer restrictions on business investments) and cons (socially irresponsible but privately optimal decisions). However, this does not apply to individuals in their private capacity, hence, the law views the negligence of private individuals as a form of strict liability; in this case, the outcome is third-best.¹⁶

¹⁶ Our usage follows the standard sense of the term in economics: The first best obtains if no harm is done, the second best occurs if a judgment can be passed in the absence of hindsight bias and the third best occurs if judgment is passed with hindsight bias.

In Kamin and Rachlinski (1995), subjects make predictive and postdictive judgments about a municipality's actions to engage in flood precautions. The damages from the flood are set at \$10⁷ and the cost of the precautions at \$10⁵; subjects were also provided with meteorological and anecdotal data. One group of subjects, group A, were asked to predict the probability of a flood in the absence of precautions. They were told that if they find the probability greater than 10% they should decide to take the precaution. Another group of subjects, group B, was asked to make the postdictive judgment. They were told that the precautions had not been taken and a flood had occurred. They were told that if their postdictive judgment is that the probability of a flood was greater than 10% they should find the defendant guilty. In Group A, 24% of the subjects thought that the precautions were merited, while 57% in group B found the defendant guilty.

19.8.3 What causes hindsight bias?

What causes hindsight bias? Kahneman (2011, p. 202) opines: "A general limitation of the human mind is its imperfect ability to reconstruct past states of knowledge, or beliefs that have changed. Once you adopt a new view of the world (or of any part of it), you immediately lose much of your ability to recall what you used to believe before your mind changed." Other explanations are that (i) people prefer to see a stable and predictable world and the hindsight bias is simply a coping mechanism in the face of an unpredictable world, and (ii) by exaggerating their ability to predict, individuals attempt to signal greater competence to others (Rachlinski, 1998).

19.8.4 Hindsight bias and the underestimation of financial volatility

We now show that hindsight bias can lead decision makers to reduce volatility estimates in financial markets and lead to insufficient revision of incorrect models (Biais and Weber, 2009). Consider the two equations (19.23) and (19.24); to simplify notation, let $E[X | I_t] = \mu$ and let $E[E[X | I_t] | I_{t+j}] = \hat{\mu}$. Then we can write (19.24) as

$$\hat{\mu} = \alpha x + (1 - \alpha)\mu. \quad (19.25)$$

Suppose that an investor at time t forecasts the average returns at time $t + j$ on an asset, or a portfolio of assets. When making the prediction at time t , the investor believes that $X \sim N(\mu, \sigma_L^2)$ with probability λ and $X \sim N(\mu, \sigma_H^2)$ with probability $1 - \lambda$, where $\sigma_L^2 < \sigma_H^2$. Thus, the density of X is given by

$$f(x | \sigma_i^2, \mu) = \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp\left(\frac{-1}{2\sigma_i^2} (x - \mu)^2\right), i = L, H. \quad (19.26)$$

At the end of time period $t + j$, after observing the realization, x , of the random variable, X , the decision maker forms posterior beliefs about which model (i.e., variance σ_L^2 or σ_H^2) is likely to have generated x . However, at the end of time $t + j$, on account of hindsight bias, the remembered mean is $\hat{\mu}$, given in (19.25). Thus, the relevant density, ex-post, is $f(x | \sigma_i^2, \hat{\mu})$; this may be used to determine, ex-post, which of the two competing models is more likely to hold. Denote by $\hat{P}(a | b)$ and $P(a | b)$, respectively, the probability of a conditional on b in the presence of hindsight bias ($\alpha > 0$) and in the absence of hindsight bias ($\alpha = 0$).

Using Bayes' rule, we have

$$\hat{P}(\sigma_L^2 | x) = \frac{f(x | \sigma_L^2, \hat{\mu})\lambda}{f(x | \sigma_L^2, \hat{\mu})\lambda + f(x | \sigma_H^2, \hat{\mu})(1 - \lambda)}. \quad (19.27)$$

Using (19.26) to substitute the value of $f(x | \sigma_i^2, \hat{\mu})$, we can rewrite (19.27) as

$$\hat{P}(\sigma_L^2 | x) = \frac{1}{1 + \left(\frac{1-\lambda}{\lambda}\right) \frac{\sigma_L}{\sigma_H} \exp\left(\frac{-1}{2} \left(\frac{1}{\sigma_H^2} - \frac{1}{\sigma_L^2}\right) (x - \hat{\mu})^2\right)}. \quad (19.28)$$

Substituting $\hat{\mu}$ from (19.25) in (19.28), we get

$$\hat{P}(\sigma_L^2 | x) = \frac{1}{1 + \left(\frac{1-\lambda}{\lambda}\right) \frac{\sigma_L}{\sigma_H} \exp\left(\frac{-1}{2} \left(\frac{1}{\sigma_H^2} - \frac{1}{\sigma_L^2}\right) (1 - \alpha)^2 (x - \mu)^2\right)}. \quad (19.29)$$

In the absence of hindsight bias ($\alpha = 0$), we get the posterior probability, $P(\sigma_L^2 | x)$, as used in typical models in economics and finance

$$P(\sigma_L^2 | x) = \frac{1}{1 + \left(\frac{1-\lambda}{\lambda}\right) \frac{\sigma_L}{\sigma_H} \exp\left(\frac{-1}{2} \left(\frac{1}{\sigma_H^2} - \frac{1}{\sigma_L^2}\right) (x - \mu)^2\right)}. \quad (19.30)$$

Comparing (19.29) and (19.30), we derive the main implications of hindsight bias in this model. From (19.29), since $\frac{-1}{2} \left(\frac{1}{\sigma_H^2} - \frac{1}{\sigma_L^2}\right) > 0$, we have $\frac{\partial \hat{P}(\sigma_L^2 | x)}{\partial \alpha} > 0$, so as hindsight bias increases (α increases), $\hat{P}(\sigma_L^2 | x)$ increases. Thus,

$$\hat{P}(\sigma_L^2 | x) > P(\sigma_L^2 | x) \text{ and } \hat{P}(\sigma_H^2 | x) < P(\sigma_H^2 | x). \quad (19.31)$$

In finance, volatility is often defined as the standard deviation of returns. From (19.31), decision makers who are subject to hindsight bias infer, ex-post, that volatility is likely to be low rather than high. Furthermore, the degree of underestimation of volatility is increasing in the degree of hindsight bias, as captured by the size of α . Since hindsight-biased decision makers underestimate volatility, if they have to invest in the relevant asset/portfolio, they will invest more than they really should, particularly if they are risk averse. One may expect that the performance of hindsight-biased individuals will, therefore, be inferior to those who are not hindsight-biased.

Using a within-subjects design, Biais and Weber (2009) propose the following empirical measure of hindsight bias: $\phi = \frac{\hat{\mu} - \mu}{x - \mu}$. Using (19.25), we get that $\phi = \alpha$. They consider two treatments. In the first treatment, the experimenter informs subjects, ex-post, of their ex-ante prediction, hence, eliminating the hindsight bias. In the second treatment, the hindsight bias is not eliminated. Subjects were given the market price of five German stocks (BASF, IKB, EON, Postbank, Premiere), two commodities (oil and gold), and the euro-dollar exchange rate. They were then asked to predict the prices for one week ahead. Subjects were also asked to give the upper and lower bounds of prices, which were converted into standard deviations or *initial implied volatility estimates*. A week later, once the relevant prices had been realized, the first treatment group was simply reminded of their previous estimates of prices and volatility, while the second treat-

ment group was asked to make the postdictive judgments (i.e., give a best estimate of their earlier estimates).

The results are shown in Table 19.7. The first row shows the hindsight bias calculated in the second treatment (this is $\phi = \alpha$). There is a positive hindsight bias for all eight assets, which is significant at the 10% level for four of the assets. The second row calculates the ratio of their initial implied volatility estimates and the remembered volatility. For five of the eight assets, the median estimates exceed 1 and for four assets this is significant at the 10% level. For one asset, the ratio is slightly less than 1 but it is statistically insignificant.

Hindsight-biased decision makers appear not to be surprised because they act as if they knew it all along. To test this, the *true surprise*, based on the ex-ante measures is defined as $\left(\frac{x-\mu}{\sigma}\right)^2$, where σ is the initial implied volatility. Define *biased surprise* as $\left(\frac{x-\hat{\mu}}{\hat{\sigma}}\right)^2$, where $\hat{\mu}$ and $\hat{\sigma}$ are remembered mean and remembered volatility; this measure can only be computed in treatment 2. For treatment 2, the biased surprise is smaller than the true surprise; the difference is statistically significant for two cases—Gold and dollar.

At the second date, the subjects are asked, yet again, to forecast their volatility estimates for another week ahead. The upward revisions in volatility are stronger for treatment 1 subjects relative to treatment 2 subjects. This is consistent with hindsight-biased subjects not updating volatility estimates as much as they ideally should, in the absence of the bias.

In order to test the performance of hindsight-biased actual investors, data was used from the Frankfurt and London branches of a large investment bank. A between-subjects design was used because competitive traders might not wish to admit that they forecast the future in a biased manner. The questions asked were similar in nature to those that the traders encountered in their jobs. Traders in both offices were found to be almost equally hindsight-biased. The estimate of α in the London and Frankfurt offices was, respectively, 0.547 and 0.579. Traders are split into high, middle, and low earners. Those in the high earners category demonstrate a lower mean and median hindsight bias, relative to the other two categories. The median difference between the middle- and low-income categories is similar. For individual-level data, and controlling for age and experience, the very highest earners have the lowest hindsight bias.

19.9 Confirmation bias

Suppose that an individual holds some initial beliefs. If the individual interprets subsequent evidence in a manner that is biased towards supporting the initial beliefs, relative to a Bayesian, then the individual is said to exhibit *confirmation bias*. For instance, buyers of a new car find that the particular model they drive suddenly appears more numerous on the roads. Since this takes the form of a biased assimilation of information, another name for this phenomenon, often used in psychology, is *biased assimilation*; for a survey of psychological research on the phenomenon see Nickerson (1998), Lord and Taylor (2009), and Mercier and Sperber (2011).

Rabin and Schrag (1999) note that confirmation bias is more likely to arise when: (1) the evidence is ambiguous and bears alternative interpretations, (2) the collection and scrutinization of evidence is selective (*hypothesis-based filtering*), (3) people are asked to judge the correlation between temporally separated events, and (4) individuals need to aggregate information from various sources. Confirmation bias is not mitigated by greater cognitive abilities or open mindedness (Stanovich and West 2007, 2008a, b).

A natural implication of confirmation bias is that individuals may be too overconfident because they typically search for confirmatory evidence and ignore or underweight contradictory evidence.

Table 19.7 Median values across decision makers. p-values are given in parentheses.

	BASF	IKB	EON	Postbank	Premiere	Oil	Gold	Dollar
Hindsight bias second treatment	0.086 (0.344)	0.463 (0.009)	0.155 (0.166)	0.784 (0.000)	0.375 (0.626)	0.182 (0.074)	0.094 (0.052)	0.264 (0.108)
Initial implied σ /remembered	1.142	1.267	1.000	1.577	1.273	1.090	1.000	0.944
Implied σ second treatment	(0.082)	(0.002)	(0.822)	(0.000)	(0.000)	(0.087)	(0.517)	(0.525)
True surprise first treatment	2340	25	57004	737	12	3320	418167	1.5
True surprise second treatment	3026	22	34649	717	18	5867	226082	1.1
Rank sum test difference between both treatments	(0.969)	(0.787)	(0.150)	(0.435)	(0.817)	(0.777)	(0.164)	(0.472)
Biased surprise second treatment	2073	8	21007	215	38	5.504	187809	0.7
Rank sum test difference between true and biased surprise	(0.910)	(0.667)	(0.158)	(0.866)	(0.278)	(0.765)	(0.024)	(0.041)
Initial implied σ /subsequent implied σ first treatment	0.923	1.024	0.630	1.001	1.041	0.981	0.996	0.892
Initial implied σ /subsequent implied σ second treatment	0.978	1.373	0.858	1.854	1.407	1.003	0.846	0.910
Rank sum test for difference between the two treatments	(0.172)	(0.021)	(0.099)	(0.002)	(0.001)	(0.349)	(0.520)	(0.369)

Source: Biels and Weber (2009).

Confirmation bias may also lead to the formation and hardening of prejudices and stereotypes. In the presence of confirmation bias, learning may lead asymptotically to the incorrect model of the world with probability 1. Rabin and Schrag (1999) distinguish confirmatory bias from a related phenomenon in which once subjects form strong beliefs about one hypothesis over competing hypotheses, they start being *inattentive* to new evidence for or against their preferred hypothesis.

The nature of confirmatory bias is beautifully captured in the following quote from Lord et al. (1979, p. 2099):

Thus, there is considerable evidence that people tend to interpret subsequent evidence so as to maintain their initial beliefs. The biased assimilation processes underlying this effect may include a propensity to remember the strengths of confirming evidence but the weaknesses of disconfirming evidence, to judge confirming evidence as relevant and reliable but disconfirming evidence as irrelevant and unreliable, and to accept confirming evidence at face value while scrutinizing disconfirming evidence hypercritically. With confirming evidence, we suspect that both lay and professional scientists rapidly reduce the complexity of the information and remember only a few well-chosen supportive impressions. With disconfirming evidence, they continue to reflect upon any information that suggests less damaging “alternative interpretations.” Indeed, they may even come to regard the ambiguities and conceptual flaws in the data opposing their hypotheses as somehow suggestive of the fundamental correctness of those hypotheses. Thus, completely inconsistent or even random data—when “processed” in a suitably biased fashion—can maintain or even reinforce one’s preconceptions.

19.9.1 *Empirical evidence for confirmation bias*

The typical experiments in confirmatory bias proceed as follows. The initial preference of subjects, or prior beliefs, for one hypothesis over another is elicited. For instance, do you support capital punishment or not (Lord et al., 1979); do you think nuclear technology is safe or unsafe (Plous, 1991); which football team do you support (Hastorf and Cantril, 1954); which party do you support in the US Presidential elections (Munro et al., 2002). The subjects are then given some information or evidence that supports one hypothesis over another and asked to form posterior beliefs. Confirmation bias arises when, relative to a Bayesian, the posterior beliefs of subjects adjust insufficiently from their priors. Such a bias was found in all the studies mentioned in this paragraph.

The initial beliefs of subjects in experiments can be created in the lab by providing them with anchoring information. Darley and Gross (1983) create two different profiles of a girl in two different treatments in a between-subjects design. In the first treatment, she is the daughter of college graduates who hold white collar jobs; a video clip shows the girl playing in a prosperous, middle class neighborhood. In a second treatment, she is the daughter of high school graduates who hold blue collar jobs; a video clip shows the girl playing in a poor, inner-city neighborhood. The subjects are divided into two sets; set A is used to elicit prior beliefs and set B to elicit posterior beliefs. Based on the information provided, the average estimate of set A individuals about the reading abilities of the girl (prior beliefs), in the first treatment, is slightly higher (4.29) as compared to the second treatment (3.90).

Subjects in set B are provided some evidence of the girl’s scholastic abilities; this takes the form of a video clip in which the girl answers a series of questions with varying degrees of success. The average posterior beliefs of individuals in set B about the girl’s reading abilities, in both treatments, were significantly different from the prior beliefs of set A. For instance, set B subjects in the second treatment (inner-city) rated the girl’s average ability to be 3.71 (posterior belief), which

is below the prior estimate of set A individuals, 3.90. The corresponding average posterior estimate of set B individuals in the first treatment (middle class) was 4.67, which is statistically higher than the prior estimate of set A individuals, 4.29. Despite the fact that subjects viewed identical videos, the girl's background has a significant influence on her perceived scholastic abilities. Girls from poor inner-city neighborhoods are assigned even lower ability levels while girls from more prosperous neighborhoods are assigned even higher ability levels. Thus, subjects exhibit the confirmation bias.

19.9.2 A formal model of confirmation bias and overconfidence

Rabin and Schrag (1999) propose a stylized model of confirmation bias that nicely illustrates the concept of confirmation bias and shows how it might lead to overconfidence. We consider this model below. Suppose that there are two mutually exclusive and exhaustive states of the world, $x \in \{A, B\}$; A and B may be thought of as competing hypotheses concerning a particular issue/event that an individual wishes to form beliefs about. The prior probability is $p(A) = p(B) = 0.5$. Time is discrete and denoted by $t \in \{1, 2, 3, \dots\}$.

At time t , the individual receives a signal, $s_t \in \{a, b\}$, that is correlated with the true state, A or B , of the world. Let $p(a | A)$ be the probability of receiving signal a if hypothesis A is true. Similarly, $p(b | B)$ is the probability of signal b if hypothesis B is true. The signals are temporally *iid* and $p(a | A) = p(b | B) = \theta$ and $0.5 < \theta < 1$. If the underlying hypothesis A is true, then $p(a | A) = \theta$ and $p(b | A) = 1 - \theta$; likewise if B is true then $p(b | B) = \theta$ and $p(a | B) = 1 - \theta$. After observing the signal, the individual uses Bayes' rule to update the beliefs about the relative likelihood of $x = A$ and $x = B$.

Confirmatory bias takes the form of misinterpreting signals that are in conflict with one's current beliefs about which of the two hypotheses is likely to be true. The signals perceived by the individual are given by $\sigma_t \in \{\alpha, \beta\}$. When α (respectively, β) is the perceived signal, the individual believes that the received signal is a (respectively, b). Signals that confirm the initial beliefs are correctly encoded, while signals that conflict with the initial beliefs are incorrectly perceived with probability $q > 0$.

Suppose, for instance, that at time t , the individual believes A is more likely than B . If the signal is $s_t = a$, then this is correctly encoded, i.e., $\sigma_t = \alpha$. But if $s_t = b$, then this is incorrectly perceived as $\sigma_t = \alpha$ with probability q and $\sigma_t = \beta$ with probability $1 - q$. Similarly if B is initially perceived to be more likely, then the signal $s_t = b$ is correctly perceived ($\sigma_t = \beta$) but $s_t = a$ is incorrectly perceived as $\sigma_t = \beta$ with probability q .

Confirmatory bias implies that although s_t is temporally *iid*, σ_t is not temporally *iid*. Denote the history of signals at the beginning of time t by $s^t = \{s_1, \dots, s_{t-1}\}$ and the history of perceived signals at the beginning of t by $\sigma^t = \{\sigma_1, \dots, \sigma_{t-1}\}$. Let $p(A | \sigma^t)$ and $p(B | \sigma^t)$ denote, respectively, the beliefs of the individual at the beginning of time t about hypotheses A and B .

Let θ^F be the probability that the individual perceives a signal σ_t that confirms his hypothesis, when it is *false* (i.e., the opposite hypothesis is, in fact, true), then

$$\theta^F = \begin{cases} p(\sigma_t = \alpha | p(A | \sigma^t) > 0.5, B) \\ p(\sigma_t = \beta | p(B | \sigma^t) > 0.5, A) \end{cases}, \quad (19.32)$$

where the conditioning expression $p(A | \sigma^t) > 0.5, B$ means that, conditional on the history of perceived signals, hypothesis A is thought to be more likely, when, in fact, B is true. Similarly the conditioning expression $p(B | \sigma^t) > 0.5, A$ means that conditional on the history of perceived signals, hypothesis B is thought to be more likely, when, in fact, A is true.

Consider the first row in (19.32). Suppose that conditional on the history of perceived signals at time t , σ^t , the individual believes A to be more likely, i.e., $p(A \mid \sigma^t) > 0.5$. How can he perceive $\sigma_t = \alpha$ when B is actually the correct hypothesis? Since B is the correct hypothesis, $p(b \mid B) = \theta$ and $p(a \mid B) = 1 - \theta$. If $s_t = a$, which, in this case, occurs with probability $1 - \theta$, the individual's beliefs are confirmed, so he perceives it correctly. However, with probability θ , the signal is b , which contradicts the initial beliefs about model A , so it is incorrectly perceived as α with probability q , hence

$$\theta^F = (1 - \theta) + \theta q. \quad (19.33)$$

Since the model is symmetric, the second row of (19.32) also leads to the same calculation reported in (19.33).

Let θ^T be the probability that the individual perceives a signal σ_t that confirms his hypothesis, when it is *true* (i.e., the opposite hypothesis is false), then

$$\theta^T = \begin{cases} p(\sigma_t = \alpha \mid p(A \mid \sigma^t) > 0.5, A) \\ p(\sigma_t = \beta \mid p(B \mid \sigma^t) > 0.5, B) \end{cases}. \quad (19.34)$$

Consider the first row in (19.34). Since A is the correct model in this case, $p(a \mid A) = \theta$ and $p(b \mid A) = 1 - \theta$; furthermore, the individual believes in model A over B , $p(A \mid \sigma^t) > 0.5$. The perception that $\sigma_t = \alpha$ can come about in two ways. First, if the signal is a (a probability θ event) it is perceived without error. Second, if the signal is b (a probability $1 - \theta$ event), it is perceived mistakenly as α with probability q . Putting these together we get

$$\theta^T = \theta + (1 - \theta) q. \quad (19.35)$$

Symmetry ensures that the second row in (19.34) leads to an identical calculation.

Definition 19.1 (*Unbiased Bayesian statistician*): An unbiased Bayesian statistician has $q = 0$, so $\theta^F = (1 - \theta)$ and $\theta^T = \theta$.

Definition 19.2 (*Confirmation-biased individual*): An extreme confirmation-biased individual has $q = 1$, so $\theta^F = 1$ and $\theta^T = 1$. Such an individual always misreads any contradictory evidence, hence, takes all evidence as supportive of his initial beliefs. Partial confirmation bias arises when $0 < q < 1$. In general, the degree of confirmation bias is increasing in q . All individuals who have $q \neq 0$ are said to be confirmation-biased.

Let us superscript variables belonging to a confirmation-biased individual by C and a Bayesian individual by B . Suppose that the individual observes n signals; n_α are perceived to be α and n_β are perceived to be β ($n_\alpha + n_\beta = n$). What are the posterior beliefs of a confirmation-biased individual that the true model is A ? Using Bayes' rule,

$$p(A \mid n_\alpha, n_\beta) = \frac{p(n_\alpha, n_\beta \mid A) p(A)}{p(n_\alpha, n_\beta \mid A) p(A) + p(n_\alpha, n_\beta \mid B) p(B)}. \quad (19.36)$$

The prior probabilities are $p(A) = p(B) = 0.5$. When the true model is A , the probability of n_α outcomes that are α and n_β outcomes that are β is $p(n_\alpha, n_\beta \mid A) = \theta^{n_\alpha} (1 - \theta)^{n_\beta}$. Similarly, we can compute $p(n_\alpha, n_\beta \mid B) = \theta^{n_\beta} (1 - \theta)^{n_\alpha}$. Substituting these calculations in (19.36), we get

$$p(A | n_\alpha, n_\beta) = \frac{\theta^{n_\alpha} (1 - \theta)^{n_\beta}}{\theta^{n_\alpha} (1 - \theta)^{n_\beta} + \theta^{n_\beta} (1 - \theta)^{n_\alpha}}.$$

Similarly computing $p(B | n_\alpha, n_\beta)$, we can determine the likelihood ratio of a confirmation-biased individual,

$$L^C(n_\alpha, n_\beta) = \frac{p(A | n_\alpha, n_\beta)}{p(B | n_\alpha, n_\beta)} = \left(\frac{\theta}{1 - \theta} \right)^{n_\alpha - n_\beta}. \quad (19.37)$$

Thus, conditional on the history, n_α, n_β , the confirmation-biased individual believes that A is more likely if $L^C(n_\alpha, n_\beta) > 1$ and B is more likely if $L^C(n_\alpha, n_\beta) < 1$. Since we have assumed that $0.5 < \theta < 1$, whenever $n_\alpha > n_\beta$, a confirmation-biased individual will assign greater probability to A being true relative to B . For instance, suppose that $n = 5$ and $n_\alpha = 3, n_\beta = 2$. Using (19.37), we have $L^C(n_\alpha, n_\beta) = \frac{\theta}{1 - \theta} > 1$. Hence, a confirmation-biased individual believes relatively more in A .

Relative to a confirmation-biased individual, denote the likelihood ratio formed by an unbiased Bayesian statistician by $L^B(n_\alpha, n_\beta)$. Suppose that the hypothetical unbiased Bayesian statistician knows that an individual suffers from confirmation bias and also misperceives signals in the manner outlined above. In general, we will have $L^B(n_\alpha, n_\beta) \neq L^C(n_\alpha, n_\beta)$. Since the perceived signals σ_t are not *iid* (although the actual signals s_t are *iid*), the order of the perceived signals received by a confirmation-biased individual influence the posterior beliefs.

Definition 19.3 Suppose that $n_\alpha > n_\beta$. The confirmation-biased individual is:

- (i) Overconfident if $L^B(n_\alpha, n_\beta) < L^C(n_\alpha, n_\beta)$. Such an individual believes that hypothesis A is more likely than can be justified by the evidence.
- (ii) Underconfident if $L^B(n_\alpha, n_\beta) > L^C(n_\alpha, n_\beta)$.

Example 19.1 (Overconfidence): Suppose that the sequence of perceived signals of the confirmation-biased individual is $\alpha, \alpha, \alpha, \beta, \beta$. What likelihood should an unbiased Bayesian statistician assign to A versus B ? We assume that our hypothetical unbiased Bayesian statistician knows that the individual suffers from confirmation bias and also misperceives signals in the manner outlined above.

The Bayesian likelihood ratio is

$$L^B = \frac{p(A | \alpha, \alpha, \alpha, \beta, \beta)}{p(B | \alpha, \alpha, \alpha, \beta, \beta)}.$$

Using Bayes' rule

$$p(A | \alpha, \alpha, \alpha, \beta, \beta) = \frac{p(\alpha, \alpha, \alpha, \beta, \beta | A) p(A)}{p(\alpha, \alpha, \alpha, \beta, \beta | A) p(A) + p(\alpha, \alpha, \alpha, \beta, \beta | B) p(B)}.$$

Writing the analogous expression for $p(B | \alpha, \alpha, \alpha, \beta, \beta)$ and noting that $p(A) = p(B) = 0.5$ we have

$$L^B = \frac{p(\alpha, \alpha, \alpha, \beta, \beta | A)}{p(\alpha, \alpha, \alpha, \beta, \beta | B)}. \quad (19.38)$$

Let us first compute $p(\alpha, \alpha, \alpha, \beta, \beta | A)$, i.e., what is the likelihood of getting the five perceived signals in the order given, conditional on the true model being A . Given the signals, since

$n_\alpha > n_\beta$, the individual always believes the true model is more likely to be A. Conditional on the true model being A (or $x = A$), the probability of the first α signal is θ . The second period signal could have been either $s_t = a$ (with probability θ) and correctly perceived to be $\sigma_t = \alpha$ or it might have been $s_t = b$ (with probability $1 - \theta$) and incorrectly perceived to be $\sigma_t = \alpha$ with probability q . This calculation is $\theta^T = p(\alpha \mid p(A \mid \sigma^t) > 0.5, A)$, the first line in (19.34). The third period perceived signal is α and gives rise to an identical calculation as in period 2. In period 4, the perceived signal is β . Now the relevant calculation is $1 - \theta^T$. An identical calculation is repeated in period 5. Thus,

$$\begin{aligned} p(\alpha, \alpha, \alpha, \beta, \beta \mid A) &= \theta (\theta^T)^2 (1 - \theta^T)^2 \\ &= \theta (\theta + (1 - \theta)q)^2 (1 - \theta)^2 (1 - q)^2 \text{ (using (19.35))}. \end{aligned} \quad (19.39)$$

Proceeding analogously, $p(\alpha, \alpha, \alpha, \beta, \beta \mid B) = (1 - \theta) (\theta^F)^2 (1 - \theta^F)^2$. Substituting the value of θ^F from (19.33) we get

$$p(\alpha, \alpha, \alpha, \beta, \beta \mid B) = (1 - \theta) ((1 - \theta) + \theta q)^2 \theta^2 (1 - q)^2. \quad (19.40)$$

Substituting (19.39), (19.40) in (19.38), the Bayesian likelihood ratio is

$$\begin{aligned} L^B &= \frac{\theta (\theta + (1 - \theta)q)^2 (1 - \theta)^2 (1 - q)^2}{(1 - \theta) ((1 - \theta) + \theta q)^2 \theta^2 (1 - q)^2}, \\ &= \left(\frac{\theta + q(1 - \theta)}{1 - \theta + q\theta} \right)^2 \frac{1 - \theta}{\theta}. \end{aligned}$$

It is now straightforward to show that $\left(\frac{\theta + q(1 - \theta)}{1 - \theta + q\theta} \right)^2 \frac{1 - \theta}{\theta} < \frac{\theta}{1 - \theta}$, hence, $L^B < L^C$. Thus, the confirmation-biased individual is overconfident.

Example 19.2 Suppose that the sequence of perceived signals of the confirmation-biased individual is $\beta, \beta, \alpha, \alpha, \alpha$. In this case, it can be shown that

$$L^B = \frac{(1 - \theta) (1 - \theta + q\theta) \theta^3}{\theta (\theta + q(1 - \theta)) (1 - \theta)^3} > \frac{\theta}{1 - \theta} = L^C. \quad (19.41)$$

Hence, the individual is underconfident.

19.10 Other judgment heuristics

We now briefly note some other judgment heuristics in this section. The aim is to provide no more than a minimum discussion of each heuristic without surveying the literature in any detail.

19.10.1 Regression towards the mean

Galton (1886) discovered that taller than average parents produce children who, on average, are shorter than them. Similarly, shorter than average parents produce children whose average height is greater than theirs. This phenomenon is known as *regression towards the mean*. Regression towards the mean always holds between two random variables whose correlation is different from 1 in absolute value.

Suppose that we measure test scores, or some other human performance measures, at two distinct points in time; each score is partly determined by random influences. Suppose that the test scores are random variables, X_1 and X_2 , which are identically distributed with common mean μ . Let us select from the first test scores, individuals whose scores deviate from the mean by more than $k > 0$, i.e., scores from the set $\{x : x - \mu > k, x \in X_1\}$. Suppose that we now track the performance of the same individuals on the second test where the test score is determined by the random variable X_2 . Denote the sample mean of the test scores of these selected individuals in the second test by μ_s . Then regression to the mean implies that $\mu_s - \mu < k$. Similarly, if we chose the original set of individuals from distribution X_1 who satisfy $\{x : \mu - x > k, x \in X_1\}$, then we must have $\mu - \mu_s < k$.

Consider the following interesting example of this phenomenon from the Israeli Air Force (Tversky and Kahneman, 1974; Kahneman, 2011). Air instructors found that when flight cadets were praised for particularly good performance, say a clean execution of an aerobatic maneuver, then their subsequent performance dropped. On the other hand, those flight cadets who were punished for poor performance, improved their subsequent performance. An inference was made that rewarding good performance was bad and punishing bad performance was good. However, this fits perfectly with an inability to appreciate the statistical phenomenon of regression to the mean.

On average, extreme performances on the aerobatic maneuver arise purely from luck. Flight cadets who were extremely lucky performed much better than the average while others did worse than the average. There is no causality between rewards/punishments and subsequent performance. The fluctuations in performance between two tests is purely random. It is of course true that low skilled pilots will continue to perform at a lower level, however, regression to the mean captures the variability in performance that arises from purely random factors.

The regression to the mean has been documented in other contexts. For instance, golfers with high scores on day 1 of a tournament typically report lower scores on day 2, while those who struggle on day 1, improve their average performances on day 2. Kahneman (2011) conjectures that people give inadequate attention to regression to the mean because the brain wishes to construct causal relations while regression is a purely statistical phenomenon that has an explanation but not a cause.

19.10.2 *False consensus effect*

The *false consensus* heuristic arises when people believe that their own behavioral choices and judgments are appropriate, and relatively common among others. Individuals walking on a tightrope across skyscrapers or launching a revolution do know that they are in a minority. Yet they also view their behavior as less deviant, relative to the behavior of others who do not engage in these behaviors (Ross et al., 1977).

The false consensus effect has some similarity with *evidential reasoning* or *projection bias*, and this is recognized in Ross et al. (1977). Under evidential reasoning, subjects ascribe diagnostic significance to their own beliefs when asked to guess the beliefs of others; see Part 4 of the book for a formalization of evidential reasoning (al-Nowaihi and Dhimi, 2015), and the experimental evidence (Robbins and Krueger, 2005).

In one experiment, Ross et al. (1977) asked Stanford University students whether they would be willing to walk around the campus for 30 minutes wearing a sandwich board inscribed with a message. For one group of students, the message read “Eat at Joe’s” and for another group it read “Repent.” When the message was Repent, those who agreed, estimated, on average, that 63.50%

of their fellow students would also agree, and those who refused, estimated that 76.67% of their fellow students would also refuse. For both messages combined, those who agreed to wear the board estimated that 62.2% of other students would also agree to wear it. Those who refused to wear the message board estimated that 67% would also refuse.

For a meta-study of 115 studies on the false consensus effect, see Mullen et al. (1985). Solan et al. (2008) show that laymen and judges alike are subject to the false consensus effect when interpreting contractual documents in insurance that have in the past led to different judicial outcomes. Leviston et al. (2013) find that individuals believe that their own views on climate are more widespread than they actually are. Those who thought that climate change was not happening believed that 43% of others thought likewise; the actual figure was 6%. Similarly, those who said “don’t know” thought that 34.1% of others believed likewise; the actual figure was less than 5%. Roth and Voskort (2014) document a false consensus effect for finance professionals in Germany. In an artefactual field experiment, finance professionals are asked to guess the risk preferences of others. A significant false consensus effect is found and professionals unduly project their own risk preferences to others. Indeed, the measured effect is strongest for experienced professionals.

19.10.3 *Confusion between necessary and sufficient conditions*

The distinction between *necessary and sufficient conditions* is critical for most inference in economics, yet people can find it difficult to make the distinction. Consider the following classic problem in Wason (1968). There are four cards labeled E, K, 4, 7 that are placed on a table. It is known that each card has a number on one side and a letter on the other. Subjects are invited to pick all the cards that can test the following rule. “*Every card with a vowel on one side has an even number on the other side.*”

Most subjects report that card E must be turned over while many report that both E and 4 must be turned over to test the rule. The rule can be written as the following logical statement:

$$\text{Vowel on one side} \Rightarrow \text{even number on the other side.} \quad (19.42)$$

Thus, a vowel on one side is sufficient for there to be an even number on the other side. But an even number on one side is only a necessary, but not a sufficient condition for a vowel on the other side. An implication of (19.42) is that

$$\text{not an even number on one side} \Rightarrow \text{not a vowel on the other side.} \quad (19.43)$$

Statement (19.42) dictates choosing card E, and (19.43) dictates choosing card 7. Choosing card 4 can never confirm (or falsify) the rule because an even number is not a sufficient condition for a vowel on the other side, yet about 65% choose card 4.

19.10.4 *Attribute substitution*

Kahneman and Frederick (2005) draw attention to the heuristic of *attribute substitution*. This occurs when one attribute of an object is substituted by another, perhaps because it comes more readily to mind. For instance, “similarity” might be substituted for “probability” in solving a prob-

lem. The usefulness of attribute substitution is that it simplifies difficult judgments by substituting easier ones in their place. The downside is, of course, that it may lead to incorrect inferences.

Consider, for instance, the following two questions asked of subjects in experiments (Strack et al., 1988). How happy are you with your life in general? How many dates did you have last month? When asked in this order, the correlation in the answers was low. However, when the order was reversed, the correlation rose to 0.66. Kahneman and Frederick (2005) propose the following explanation: “We suggest that thinking about the dating question automatically evokes an affectively charged evaluation of one’s satisfaction in that domain of life, which lingers to become the heuristic attribute when the happiness question is subsequently encountered.” Schwarz, Strack, and Mai (1991) replicate these results when marital satisfaction replaces the number of dates (the respective correlations were 0.32 when the questions were asked in direct order, and 0.67 in the reversed order). For a review of findings from similar experiments, see Schwarz and Bohner (2001).

19.11 Dual process models and judgment heuristics

In this section, we briefly outline some of the explanations for judgment heuristics and biases. Although our treatment is relatively condensed, this is an important area of research in psychology that has attracted a great deal of scholarship.¹⁷ Various judgment heuristics can be explained by taking a *two modules view* of the brain (Kahneman, 2011). The two modules have several distinct names in the literature such as “dual process theories” and “intuition and reason.” Following Stanovich and West (2000, 2002), we refer to the two modules as System 1 and System 2; these might not be the most illuminating terms but they are easiest to use and are popular.

System 1 is fast, automatic, and little mental effort is involved in using it. System 2, on the other hand, is much slower, deliberative, and requires cognitive effort. System 1 forms impressions, feelings, and intuitions for System 2. If System 2 endorses these intuitions, they turn into beliefs. System 2 typically springs into action when System 1 encounters a situation that is counter to its worldview. Table 19.8 gives the different usages of these terms, which illustrate nicely the distinction between the two systems, and the various features associated with them.

Kahneman (2011, p. 25) offers a very illuminating view of the two systems:

The division of labor between System 1 and System 2 is highly efficient: it minimizes effort and optimizes performance . . . System 1 is generally very good at what it does: its models of familiar situations are accurate, its short term predictions are accurate as well, and its initial reactions to challenges are swift and generally appropriate. System 1 has biases, however, and it is prone to making systematic errors in specified circumstances . . . it answers easier questions than the one it was asked and has little understanding of logic and statistics. One further limitation of System 1 is that it cannot be turned off.

It is important to note that the two systems are purely fictitious characters that provide a way of organizing our understanding of the brain. Indeed “there is no part of the brain that either of the systems would call home” (Kahneman, 2011, p. 29).

One explanation for judgment heuristics is that individuals save on cognitive effort by using the relatively less cognitively taxing and spontaneous System 1, as compared to System 2. This

¹⁷ The distinction between Systems 1 and 2 that we outline in this section is used extensively in formal economic models; see Part 6 of the book.

Table 19.8 Different usages of the System 1–System 2 terms.

	System 1	System 2
Dual-Process Theories:		
Sloman (1996)	associative system	rule-based system
Evans (1984; 1989)	heuristic processing	analytic processing
Evans & Over (1996)	tacit thought processes	explicit thought processes
Reber (1993)	implicit cognition	explicit learning
Levinson (1995)	interactional intelligence	analytic intelligence
Epstein (1994)	experiential system	rational system
Pollock (1991)	quick and inflexible modules	intellection
Hammond (1996)	intuitive cognition	analytical cognition
Klein (1998)	recognition-primed decisions	rational choice strategy
Johnson-Laird (1983)	implicit inferences	explicit inferences
Shiffrin & Schneider (1977)	automatic processing	controlled processing
Posner & Snyder (1975)	automatic activation	conscious processing system
Properties:	associative	rule-based
	holistic	analytic
	automatic	controlled
	relatively undemanding of cognitive capacity	demanding of cognitive capacity
	relatively fast	relatively slow
	acquisition by biology, exposure, and personal experience	acquisition by cultural and formal tuition
Task Construal	highly contextualized	decontextualized
	personalized	depersonalized
	conversational and socialized	asocial
Type of Intelligence	interactional	analytic (psychometric IQ)
Indexed:	(conversation and implicature)	

Source: Keith E. Stanovich and Richard F. West (2000). "Individual differences in reasoning: implications for the rationality debate?" *Behavioral and Brain Sciences* 23: 645–665, with permission by Cambridge University Press.

is also known as the *cognitive miser* model. For instance, Tversky and Kahneman (1973) found that mathematical psychologists exhibited the law of small numbers. Despite being well trained in statistics, they used their fast and reactive System 1, thus neglecting basic statistical analysis.

System 1 is biased to look for confirmatory evidence that enables individuals to make sense of the world around them. Such a bias is also reflected in the desire of individuals to reduce or avoid *cognitive dissonance*. Cognitive dissonance arises from the psychological stress placed upon individuals when they hold different *mental models* of the world (see Section 19.13 on mental models). The individuals' response in such cases is to rationalize the world around them and reduce conflict between different mental models. The term "cognitive dissonance" was proposed by Festinger (1957) and a classic set of experiments can be found in Festinger and Carlsmith (1959).

The fact that the suppression of an incorrect response by System 1 requires cognitive effort on the part of System 2 can be illustrated by the *Stroop task*. The term Stroop task is now used for any generic task in which an automatic, highly practiced response, by System 1 is, in fact, incorrect. System 2 then attempts to control the incorrect response (successfully or unsuccessfully), which accounts for the hesitation in subject responses. In the classical example, subjects are asked to read names of different colors, but some of the names are printed in a different color, e.g., the word “green” may be printed in the color “red.” Subjects take longer to answer these questions and make more errors relative to the case where the names of the colors and the color in which they are printed, match.

System 1 is also associated with high levels of skill and achievement. For instance, expert chess players can quickly gauge the strength of their position without considering all possibilities. Indeed, considering all possibilities would be seriously taxing even for System 2 and deplete scarce cognitive resources. Kahneman and Frederick (2005) provide a list of factors that determine the relative usage of the two systems that we now outline. Unfamiliar tasks and those involving abstract concepts are likely to be dealt with by System 2. However, System 1 is likely to be used to a greater extent when the time available for deliberation is smaller, exposure to statistical training is lower, and some traditional measures of IQ are lower.

The relative use of the two systems might also depend on the moods of individuals. A representation of the data in terms of frequencies, rather than probabilities, can also influence the ability of System 2 to correct the errors of System 1. For instance, Kahneman and Frederick (2005) write: “the language of frequencies improves respondent’s ability to impose the logic of set inclusion on their considered judgments.” Indeed, we have reviewed a great deal of evidence above that in some cases, frequency information rather than probability information reduces some behavioral biases but the biases do remain, often for a majority of the subjects.

Higher cognitive ability might be associated with a greater and more efficient use of System 2. If this insight is correct, then we should see a lower prevalence of judgment errors among subjects in experiments whose cognitive ability is higher. Oechssler et al. (2009) test this idea. In order to measure cognitive ability, they use the *cognitive reflection test* (CRT) proposed in Frederick (2005). In this test, subjects are asked the following three questions.

1. A bat and a ball together cost 110 cents. The bat costs 100 cents more than the ball. How much does the ball cost?
2. If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets?
3. In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake?

If one only engages System 1, then the answers to the three questions are respectively 10 cents, 100 min, 24 days. However, these answers are incorrect. The correct answers based on engaging System 2 are, respectively, 5 cents, 5 min, 47 days. The average number of correct responses was 2.05.¹⁸ As in Frederick’s (2005) sample, a lower CRT score is associated with greater risk seeking and greater impatience (higher discount rate).

There is a statistically significant difference in the error rates of individuals with high and low CRT scores. For the conjunction fallacy (the Linda problem), 62.6% of the low CRT group engaged in the fallacy, while only 38.3% of the high CRT group did. Now consider the conservatism problem, or the underweighting of the likelihood of a sample. In a problem given

¹⁸ In Frederick’s (2005) sample this score lies between the scores for the MIT and Princeton students.

to the subjects, the probability consistent with Bayes' rule was 0.967; however, the low CRT group reported a probability of 0.58, while the high CRT group reported a probability of 0.69. Interestingly cognitive ability did not turn out to have any effect on the anchoring heuristic. This suggests that the relative extent to which the heuristics are hardwired in the brain, differs. Some heuristics can be reduced with greater cognitive effort/ability, while others cannot be.

Bergman et al. (2010) find that as cognitive ability increases, the anchoring effect reduces but does not vanish. One reason for the difference in their findings from Oechssler et al. (2009) is that they do not use the CRT but rather a cognitive ability test (CAT) that is based on answering 44 questions in 20 minutes. Dohmen et al. (2009) find, using data from the German population, that cognitive ability and the number of years of schooling are positively related to better performance in probability judgment tasks, and a reduced incidence of the gambler's fallacy. For the relation between IQ and the extent of judgment biases, see Stanovich (1999) and Stanovich and West (2002).

A higher IQ can also be a doubled edged sword. On the one hand, a higher IQ can induce a more efficient system 2 that corrects for the mistakes of system 1. However, on the other hand, a high IQ can induce individuals to focus on a more plausible error rather than a random error; see for instance, Kahneman (2000).

19.12 Coarse thinking and persuasion

Persuasion is central to determining the outcomes in several forms of human interaction. Activities associated with persuasion include advertising, political campaigns, religious preaching, selling activities, media campaigns, and military propaganda. Of these, advertising as a form of persuasion, is the one that economists have focused on the most. A dominant strand in industrial organization insists that the objective of advertising is the dissemination of information (Stigler, 1987; Grossman and Hart, 1980; Milgrom, 1981).

In contrast, in psychology, persuasion is understood as the transference of information from one context to another (Kahneman and Tversky, 1982; Zaltman, 1997). Consider the following three examples of successful persuasion, taken from Shleifer et al. (2008).

1. American Express travelers cheques featured Karl Malden, an actor who played a detective in a well-known television show. The advertisements suggested that, for safety, people should carry American Express travelers cheques in preference to cash.
2. After the crash of the Internet bubble, the brokerage firm Charles Schwab launched an advertising campaign that featured judges, ship captains, crossing guards, doctors, grandmothers, and other steady and reliable individuals, each standing next to a Schwab investment specialist. The headline of one advertisement, depicting a Schwab professional standing next to a pediatrician, said: "Both are seen as pillars of trust."
3. Consider Arnold Schwarzenegger's memorable speech at the 2004 Republican National Convention. In the best remembered part of his speech, Schwarzenegger defended free trade: "To those critics who are so pessimistic about our economy, I say: Don't be economic girlie men! . . . Now they say India and China are overtaking us. Don't you believe it. We may hit a few bumps—but America always moves ahead. That's what Americans do."

All these examples illustrate a kind of thinking, known as *associative, analogical, or metaphoric thinking* in psychology. Information that has great relevance in one context is used in another context where it has much less, or worse, no relevance at all. Based on Mullainathan (2002),

Shleifer et al. (2008) make progress in formalizing these ideas by focusing on the concept of *coarse thinking*.

To explain *coarse thinking*, suppose that individuals are able to categorize a problem into one of n mutually exclusive categories. Any two problems in the same category have identical methods of solution. So, for instance, if there are two distinct problems in category $j \in \{1, \dots, n\}$, then they are solved using an identical model. The polar opposite of coarse thinking is known as *sophisticated thinking* in which different models may be suitable for each problem. Coarse thinking can be exploited by others for their benefit. So, for instance, a beautiful model or a well-known sports personality posing next to a car gives no information about the car, per se. However, individuals might, using associative thinking, transfer attributes they associate with the individual posing with the car, to the car itself.

Suppose that for some good, service, or information, a buyer needs to assess the level of quality, $q \in (-\infty, \infty)$. The buyer receives two possible messages $m \in \{\underline{m}, \bar{m}\}$, where $\underline{m}$ and $\bar{m}$ have the respective connotations of *bad* and *good* messages. The individual knows that there are two mutually exclusive *categories* or *situations* $s \in \{\underline{s}, \bar{s}\}$. Messages are only informative when $s = \bar{s}$; when $s = \underline{s}$ messages are completely uninformative. The joint distribution of quality, q , messages, m , and situations, s , is given by $p(q, m, s)$.

Example 19.3 Suppose that putting “silk” in shampoo has no effect on hair. Yet a shampoo company advertises putting silk in shampoo in the hope that associative thinking will persuade consumers to buy the shampoo. In this case, the message $\bar{m} = \text{silky}$, while $\underline{m}$ is any neutral message. One situation $\bar{s}$ is that a person is looking for a shampoo, and in a second situation $\underline{s}$ the person is not looking for a shampoo but simply assessing the quality of the plastic used in shampoo bottles; so $\bar{s} = \text{shampoo}$ and $\underline{s} = \text{bottle}$. The message $\bar{m} = \text{silky}$ is uninformative in the situation $\underline{s}$ but informative in the situation $\bar{s}$. A Bayesian thinker will simply avoid the shampoo ad in the first situation. However, a coarse thinker (see formal definition below) who lumps situations together in categories will give at least some weight to the shampoo advertisement.

Example 19.4 Consider the first of three examples of successful persuasion given above on American Express travelers checks. Here the message $\bar{m} = \text{Karl Malden advertisement}$, while $\underline{m}$ is any neutral message. The situations are: $\bar{s} = \text{an individual with travel plans}$ and $\underline{s} = \text{an individual who has no travel plans}$. Clearly the message on travelers cheques is only relevant to individuals with travel plans.

The assumptions are formally stated as follows.

A1: The distribution of quality, q , is independent of the situation, i.e.,

$$p(q \mid \underline{s}) = p(q \mid \bar{s}) = p(q). \quad (19.44)$$

A2: The average quality, $E[q] = 0$; using (19.44) this implies

$$E[q \mid s] = 0. \quad (19.45)$$

A3: Any message is possible under any situation

$$p(m \mid s) > 0; m \in \{\underline{m}, \overline{m}\}, s \in \{\underline{s}, \bar{s}\}. \quad (19.46)$$

A4: Messages are informative only in situation $s = \bar{s}$. In situation $\bar{s}$, the message $\overline{m}$ is good news about average quality, while the message $\underline{m}$ is bad news about average quality. Thus,

$$E[q \mid \underline{m}, \bar{s}] < E[q \mid \overline{m}, \bar{s}]. \quad (19.47)$$

By contrast, in situation $\underline{s}$, the message is uninformative, i.e.,

$$p(q \mid \underline{m}, \underline{s}) = p(q \mid \overline{m}, \underline{s}) = p(q \mid \underline{s}) = p(q). \quad (19.48)$$

The last equality uses (19.44).

From the definition of conditional probability, we know that

$$p(q \mid \underline{m}, \bar{s}) = \frac{p(q, \underline{m}, \bar{s})}{p(\underline{m}, \bar{s})} = \frac{p(q, \underline{m}, \bar{s})}{p(\underline{m} \mid \bar{s}) p(\bar{s})}. \quad (19.49)$$

Analogous to (19.49), we can write

$$p(q \mid \overline{m}, \bar{s}) = \frac{p(q, \overline{m}, \bar{s})}{p(\overline{m}, \bar{s})} = \frac{p(q, \overline{m}, \bar{s})}{p(\overline{m} \mid \bar{s}) p(\bar{s})}. \quad (19.50)$$

From the definition of a marginal probability distribution, we get that

$$p(q, \underline{m}, \bar{s}) + p(q, \overline{m}, \bar{s}) = p(q, \bar{s}). \quad (19.51)$$

Substitute $p(q, \overline{m}, \bar{s})$ from (19.50) in (19.51)

$$p(q, \underline{m}, \bar{s}) + p(q \mid \overline{m}, \bar{s}) p(\overline{m} \mid \bar{s}) p(\bar{s}) = p(q, \bar{s}). \quad (19.52)$$

Solve (19.52) for $p(q, \underline{m}, \bar{s})$ and substitute in (19.49) to get

$$p(q \mid \underline{m}, \bar{s}) = \frac{p(q, \bar{s})}{p(\underline{m} \mid \bar{s}) p(\bar{s})} - \frac{p(q \mid \overline{m}, \bar{s}) p(\overline{m} \mid \bar{s})}{p(\underline{m} \mid \bar{s})}. \quad (19.53)$$

But $p(q \mid \bar{s}) = \frac{p(q, \bar{s})}{p(\bar{s})}$, hence, (19.53) can be written as

$$p(q \mid \underline{m}, \bar{s}) = \frac{p(q \mid \bar{s})}{p(\underline{m} \mid \bar{s})} - \frac{p(q \mid \overline{m}, \bar{s}) p(\overline{m} \mid \bar{s})}{p(\underline{m} \mid \bar{s})}. \quad (19.54)$$

Multiply both sides by q and integrate both sides with respect to q to get

$$E[q \mid \underline{m}, \bar{s}] = \frac{1}{p(\underline{m} \mid \bar{s})} E[q \mid \bar{s}] - \frac{p(\overline{m} \mid \bar{s})}{p(\underline{m} \mid \bar{s})} E[q \mid \overline{m}, \bar{s}]. \quad (19.55)$$

From (19.45), $E[q | \bar{s}] = 0$, hence, we can write (19.55) as:

$$E[q | \underline{m}, \bar{s}] = -\frac{p(\bar{m} | \bar{s})}{p(\underline{m} | \bar{s})} E[q | \bar{m}, \bar{s}]. \quad (19.56)$$

On the other hand, in category/situation $\underline{s}$, using (19.45), (19.48) we get,

$$E[q | \underline{m}, \underline{s}] = E[q | \bar{m}, \underline{s}] = E[q | \underline{s}] = 0. \quad (19.57)$$

From (19.56), (19.57), a Bayesian uses the correct model for each category. In category $\bar{s}$, where the signals are informative, the Bayesian forms a different (and correct) posterior relative to category $\underline{s}$. We now describe how a coarse thinker infers quality, conditional on a message $m = \{g, b\}$.

We use the subscript “c” for a coarse thinker. So, p_c refers to the probability assigned by a coarse thinker to some event and E_c refers to the expectation computed by a coarse thinker. In the absence of a subscript, as in p , E , we refer to the analogous computations of a Bayesian who does not use categorization. We make the following assumption about a coarse thinker.

A5: A coarse thinker includes both situations $\underline{s}$ and $\bar{s}$ into the same category, hence, for any signal, s ,

$$p_c(q | m, s) = p(q | m, \underline{s}) p(\underline{s} | m) + p(q | m, \bar{s}) p(\bar{s} | m). \quad (19.58)$$

In (19.58), although the coarse thinker is able to observe the situation, the same model is used for both situations. A coarse thinker does not evaluate the informativeness of the message for the particular situation, but rather by the average informativeness across all situations in the category, $s \in \{\underline{s}, \bar{s}\}$.

Multiply both sides of (19.58) by q and integrate with respect to q to get

$$E_c(q | m, s) = E(q | m, \underline{s}) p(\underline{s} | m) + E(q | m, \bar{s}) p(\bar{s} | m). \quad (19.59)$$

Since $p(\underline{s} | m) + p(\bar{s} | m) = 1$ and $p(\underline{s} | m) \leq 1$, $p(\bar{s} | m) \leq 1$, it follows that the expected quality inferred by the coarse thinker is a linear combination of what a Bayesian would compute. Since $0 = E(q | m, \underline{s}) < E(q | m, \bar{s})$ it follows from (19.59) that

$$E(q | m, \underline{s}) \leq E_c(q | m, s) \leq E(q | m, \bar{s}). \quad (19.60)$$

This leads to the following straightforward implication.

Proposition 19.1 *Relative to a Bayesian, the coarse thinker overreacts when the signal is $\underline{s}$ and underreacts when the signal is $\bar{s}$.*

The intuition for the result in Proposition 19.1 is that the coarse thinker acts as a Bayesian who is unable to distinguish between the two situations. When the situation is completely uninformative, he still puts some weight on it being an informative situation, so he overreacts. This can explain the third example above; voters seem to put some weight on Arnold Schwarzenegger’s optimism of America’s ability to succeed under free trade. They transfer Arnold’s ability to win over strong characters in films into another context where such ability has apparently little relevance. Conversely, when the situation is completely informative, the coarse thinker puts some

weight on the situation being completely uninformative, so he underweights the information content of the message.

We now examine the incentives of a persuader to alter the message observed by an individual who might be persuaded. We make the following assumptions.

- A6:** The objective of the persuader is to maximize the expected quality as perceived by the individual. Presumably, an increase in the quality induces a buyer to buy more, or a politician to get more votes.
- A7:** Suppose that the true message is given by $m \in \{\underline{m}, \bar{m}\}$. Denote the message relayed by the persuader to the individual by $m' \in \{\underline{m}, \bar{m}\}$; we allow for $m' \neq m$. The persuader can distort the true message at some cost $c > 0$. The individual is a *face value individual* in the sense that any message given to the individual by the persuader is believed.¹⁹

In the light of A6, A7, the payoff of the persuader is given by

$$\begin{cases} E(q | m', s) & \text{if } m' = m \\ E(q | m', s) - c & \text{if } m' \neq m \end{cases}$$

We now examine the incentives of the persuader to distort the message.

Proposition 19.2 (*Bayesian individual*): When $s = \underline{s}$, the persuader always reveals the true message to a face value Bayesian and never incurs a cost of deception, c .

Proof: From (19.57),

$$E[q | \underline{m}, \underline{s}] = E[q | \bar{m}, \underline{s}] = 0,$$

hence, by altering a message there is no improvement in average quality. However, the persuader pays a cost, c , of altering the message. Hence, it is never optimal to do so. ■

Proposition 19.3 (*Coarse thinkers*): Suppose that $s = \underline{s}$. The persuader might not reveal the true message to a face value coarse thinker. In particular, if c is less than some critical value c^* , then the persuader will replace a bad message with a good message.

Proof: Using (19.57), the first term on the RHS of (19.59) is zero. Furthermore, using Bayes' rule we get that $p(\bar{s} | m) = \frac{p(m|\bar{s})p(\bar{s})}{p(m)}$, hence, from (19.58) we get that

$$E_c(q | m, s) = E(q | m, \bar{s}) \frac{p(m | \bar{s})p(\bar{s})}{p(m)}.$$

Using (19.56) we can derive the following expressions for the expected quality for a coarse thinker under the two possible messages $\bar{m}$ and $\underline{m}$.

$$E_c(q | \bar{m}, s) = E(q | \bar{m}, \bar{s}) \frac{p(\bar{m} | \bar{s})p(\bar{s})}{p(\bar{m})}. \quad (19.61)$$

¹⁹ One could alter this assumption to allow for strategic inference by the individual that takes into account the incentives of the persuader to distort the message. This is allowed for by Shleifer et al. (2008), but it does not alter the essential insights.

$$E_c(q | \underline{m}, s) = -E[q | \bar{m}, \bar{s}] \frac{p(\bar{m} | \bar{s}) p(\bar{s})}{p(\underline{m})}. \quad (19.62)$$

For the persuader, replacing the true message $\underline{m}$ with $m' = \bar{m}$ improves profits if

$$\begin{aligned} E_c(q | \bar{m}, s) - c &> E_c(q | \underline{m}, s). \\ \Leftrightarrow c &< E_c(q | \bar{m}, s) - E_c(q | \underline{m}, s). \end{aligned} \quad (19.63)$$

Substitute (19.61), (19.62) in (19.63) we get

$$c < c^* \equiv E[q | \bar{m}, \bar{s}] p(\bar{m} | \bar{s}) p(\bar{s}) \left(\frac{1}{p(\bar{m}) p(\underline{m})} \right), \quad (19.64)$$

which completes the proof. ■

We now give some comparative static results.

Proposition 19.4 *Suppose that c^* is as defined in Proposition 19.3. Then c^* is increasing in (1) the probability with which $s = \bar{s}$ occurs, i.e., $p(\bar{s})$, and (2) the informativeness of message $\bar{m}$ when $s = \bar{s}$, i.e., $E[q | \bar{m}, \bar{s}]$.*

Proof: Directly differentiate both sides of (19.64) successively with respect to $p(\bar{s})$ and $E[q | \bar{m}, \bar{s}]$. ■

From Propositions 19.3, 19.4, an increase in $p(\bar{s})$, $E[q | \bar{m}, \bar{s}]$ increases c^* , increasing the likelihood that the persuader lies by giving out good news ($\bar{m}$) when the news is actually bad. The intuition is as follows. From (19.61), (19.62), an increase in $p(\bar{s})$, $E[q | \bar{m}, \bar{s}]$ raises $E_c(q | \bar{m}, s)$ and reduces $E_c(q | \underline{m}, s)$. Thus, the marginal incentive of the persuader to lie and give out the message $\bar{m}$ increases.

DellaVigna and Gentzkow (2010) survey the non-experimental evidence on persuasion; there is much evidence that is supportive of persuasion. They consider four groups of individuals that are typically the target of persuasion: consumers, voters, donors, and investors. They consider questions such as: Are consumers affected by advertising? Is the voting decision of voters influenced by political campaigns? Do newspaper endorsements and endorsements by influential citizens persuade voters to vote differently? Does door to door campaigning for funds persuade donors to give more? Do analyst recommendations influence investment in stocks? The set of questions that falls within the domain of persuasion theories is wide and important. This continues to be an active area of research.

We do not address here the important issues of how and why these categories are formed and how/why they might change over time. These questions lie beyond the scope of the current model, but their resolution will ultimately determine the importance of this class of models. These observations are similar to those we made for *analogy-based equilibria* in Part 4. The issue of optimal categorization is considered in Mohlin (2014), Peski (2011), and Fryer and Jackson (2008). Mengel (2012) considers the evolution of coarse categorization. We close with an example of categorization, which speaks to the issue of racism in the workplace that arises despite the absence of social identity.

Example 19.5 (Fryer and Jackson, 2008) *Suppose that employees are either black or white. In a population of 100 workers, 90 are white and 10 are black (minority group). Workers' types*

are either high ability (H) or low ability (L). Half the workers in each race are of each type. Denote a black worker by 0, a white worker by 1, a low ability worker by 0 and a high ability worker by 1. Thus, there are four categories of workers: (0, 0), (0, 1), (1, 0), (1, 1). Suppose that the employer forms only three categories. Thus, at least one of the categories will contain either two different races, or two different types of workers, or both.

Suppose that the employer has previously interacted with a large number of workers in each of the categories, in proportion to their presence in the population. Since there are 10 black workers out of every 100 and half of those are low ability, so the employer must have interacted with 5 workers of type (0,0). Calculating in the same manner, the employer must have interacted with 5 workers of type (0,1), 45 workers of type (1,0) and 45 workers of type (1,1). The mean of any two categories that are grouped together is a population proportions weighted mean. So the mean of the categorization {(0,1), (1,1)} is (0.9, 1) while the mean of {(1,0), (1,1)} is (1, 0.5). The mean of any single category is the category itself; e.g., the mean of {(0,1)} is just (0,1). Let the objective of the employer be to minimize the “sum across categories of the total variation of the mean from each category”; clearly several other objectives are also possible. Fryer and Jackson show that the optimal categorization under such an objective is {(1,0), (1,1), {(0,0), (0,1)}}.

Thus, the optimal categorization groups together the two types of black workers in the same category. In this model, blacks are more coarsely sorted not because of racism, or malevolence, but simply because of the underlying objective function and the proportions of blacks in the population. Since the two types of black workers are put into the same category, the representative member of this category is viewed by the employer as $0.5(0,0) + 0.5(0,1) = (0,0.5)$. This is good news for a low ability black worker but bad news for a high ability black worker. Indeed if the decision to acquire human capital is endogenous, this is likely to reduce the incentive of black workers to acquire high ability.

19.13 Mental models

Under coarse thinking, individuals form categories such that all elements in a category are treated with an identical model. This economizes on the number of mental models required to analyze a large number of events, hence, it economizes on scarce cognitive resources. A useful working definition of *mental models*, based on Denzau and North (1994), is given in WDR (2015, p. 11): “When people think, they generally do not draw on concepts that they have invented themselves. Instead, they use concepts, categories, identities, prototypes, stereotypes, causal narratives, and worldviews drawn from their communities.” Chapter 3 in WDR (2015) gives an excellent introduction to mental models. The WDR (2015, p. 11) further writes: “There are mental models for how much to talk to children, what risks to insure, what to save for, what the climate is like, and what causes disease. Many mental models are useful; others are not as useful and may even contribute to the intergenerational transmission of poverty.”

Mental models are essential for people to make decisions by economizing on cognitive costs and can be passed down from generation to generation. Mental models can persist even when they are no longer useful, hence, while they economize on cognitive costs, they can lead to suboptimal decisions. An example of harmful mental models is stereotyping of other individuals and racial discrimination based on social group identities; see Chapter 7 and Example 19.5 above. Datta and Mullainathan (2014) document several examples of incorrect mental models that are detrimental to self-interest; let us consider some of these examples.

1. In many developing countries, there is a low take-up of the oral rehydration solution (ORS), a relatively cheap and effective mix of essential salts to treat diarrhea. However, 35% of poor women surveyed in India believed that the best response to diarrhea is to reduce, not increase, fluid intake; they possibly draw an incorrect analogy to a leaky bucket. Since ORS saves lives but does not alleviate the symptoms of diarrhea, it allows many people to hold on to an incorrect, and potentially fatal, mental model.
2. Farmers might overuse nitrogenous fertilizers because they believe that such fertilizers promote greener plants. This model works well for plants like spinach but not as well for grains, because it over-encourages green healthy leafs at the cost of lower crop yields.
3. Parents in Morocco and Madagascar pull their children out of school too early. In their incorrect mental model, education only pays if children go all the way up to secondary school. However, evidence shows that every extra year of schooling is equally valuable. Thus, while parents could afford to send their children to school for longer, doing so is inconsistent with their mental models (Banerjee and Duflo, 2011).

Insofar as undesirable outcomes are not simply the product of cognitive limitations but also of defective mental models, policy must address factors conducive for the formation of appropriate mental models.

Why do “bad” mental models persist? There is more than one potential cause. A range of judgment heuristics that we have already considered in this chapter may provide potential clues. Mental models may serve the purpose of *anchors* and the *availability heuristic* could ensure that existing mental models are the ones most likely to come to mind when taking decisions. *Confirmation bias* could ensure that individuals ignore information that contradicts their existing mental models. Social identity (see Chapter 7) too could play a role. By conforming to preexisting and shared mental models in a social group, individuals may wish to demonstrate that they belong to the group of insiders.

Existing mental models may also foreclose the choices that people could make to test the models (WDR, 2015). For instance, inadequate social exchange between communities that mistrust each other can allow mutual stereotypes to persist. Female genital mutilation may persist in parts of the world because contact with the non-mutilated genitalia may be perceived to be harmful. In the face of this incorrect model, individuals may be unwilling to experiment with alternative mental models, due to the high perceived costs of doing so (Mackie, 1996; Gollaher, 2000).

Although individuals can exhibit heterogeneity in mental models, shared mental models are essential to developing institutions. Indeed, there is a two-directional feedback between mental models and institutions. For instance, in Sierra Leone, the degree of competition for the role of paramount chiefs, who are elected for a lifetime, also determines the quality of governance (Acemoglu et al., 2013). In each area, the degree of competition is fixed because paramount chiefs can only come from families that were originally recognized by the British colonial authorities. Greater competition is found to be conducive to better governance. However, the persistence of the degree of competition in each area has influenced the mental models that people hold for respect towards authority. In areas with lower political competition, people are more respectful of authority and less likely to question the actions of leaders. This, in turn, perpetuates the existing degree of competition.

Some mental models are hardwired in humans due to their evolutionary history (e.g., fear of snakes and spiders; see Chapter 14 for more details). Historical accidents can also explain some mental models. For instance, the current low levels of trust in parts of Africa can be traced back

to the slave trade (Nunn, 2008; Nunn and Wantchekon, 2011). Black Africans required guns to protect themselves against capture by the slave traders. However, because they could not afford to buy guns, they kidnapped and sold other black Africans to slavery, in exchange for money that, in turn, financed guns. This eventually led to mistrust of strangers that still persists in some parts of Africa. In another example, the physical superiority of males gives them a comparative advantage in the use of agricultural implements, such as ploughs. Hence, gender inequities are more pronounced in societies with an agricultural past (Alesina et al., 2013).

Mental models are distinct from norms in that they are not enforced by social norms. Rather, mental models reflect “broad ideas about how the world works and one’s place in it” (WDR, 2015).

How does one get people to change incorrect mental models? Datta and Mullainathan (2014) argue that simply telling people that they hold the wrong mental model is unlikely to work. A more promising strategy is to provide evidence that falsifies the core beliefs in an incorrect mental model. For instance, a particularly potent method of getting people to pay attention to reducing a wasteful activity is by comparison with a reference social group or peers. For instance, when individuals are given information about a comparison of their electricity consumption relative to neighbors, they reduce consumption by as much as that caused by an 11–20% increase in the price of electricity (Allcott and Mullainathan 2010; Allcott and Rogers, 2014); a similar effect is documented for water consumption (Ferraro and Price 2013).

Villagers in poor countries are more likely to educate girls if greater information is provided on the availability of jobs for girls with high school degrees (Jensen, 2012). Long-running soap operas in Brazil that emphasized small family size led to reduced fertility, particularly for those of a similar age to the role models in the soap operas (La Ferrara et al., 2012). Exposure to cable TV in India led to a reduction in fertility, increase in women’s status, reduction in domestic violence, and an increase in children’s school enrollment, suggesting a change in mental models (Jensen and Oster, 2009). Legal or regulatory changes can also induce changes in mental models (Sunstein, 1996); this might have been the case in changes in attitudes towards slavery, wearing helmets, smoking, racism, and casteism.

19.14 Herbert Simon’s approach to bounded rationality

In neoclassical economics, individuals use subjective expected utility (SEU) to evaluate risky/uncertain situations and employ Bayes’ rule to update information. No cognitive limitations are placed on the individual, nor is the application of SEU dependent on the environment that an individual faces. The heuristics and biases approach shows that this view of human nature is often rejected. This has led some to argue that mainstream economic theory is based on an unrealistic picture of human decision making (Selten, 2001).

Rationality has several possible interpretations. A common view in economics is that individual behavior is rational if it can be interpreted as the solution to some optimization problem. Of course, the optimization problem must be a plausible one. Otherwise, it is possible to rationalize any behavior by assigning the utility function to take a value 1 for the action chosen by an individual and zero otherwise. The notion of the optimality of behavior, or the quality of the decision, has been termed as *substantive rationality* by Simon (1955). Classical economic theory focuses mainly on substantive rationality and does not give much attention to *procedural rationality*, or the quality of the process of decision.

Substantive rationality did not figure directly in the writings of classical economists. Simon (2000) interprets Voltaire's well-known saying "the best is the enemy of the good" as implying that "if you are too preoccupied with attaining the optimum, you won't even get an acceptable result." Indeed, much of the discourse in economics until well into the middle of the twentieth century used the notion of a *reasonable person* rather than that of a *utility maximizing rational person*.

Is the rationality hypothesis compelling on grounds of empirical evidence? The evidence is mixed at best, and success is mixed with well-documented and prominent failures. Börgers (1996) concludes that "Unfortunately, there are no striking examples of empirical regularities, which are well explained by the rationality hypothesis." Substantive rationality, as used in neoclassical economics, prescribes choosing on the basis of marginalist conditions; for instance, firms produce at the point where marginal revenues equal marginal costs. In his Nobel lecture, Simon (1978, p. 347) is very explicit on this issue: "But there are no direct observations that individuals or firms do actually equate marginal costs and revenues."

Furthermore, Selten (2001) writes. "However, there is overwhelming experimental evidence for substantial deviations for Bayesian rationality (Kahneman et al. 1982): people do not obey Bayes' rule, their probability judgements fail to satisfy basic requirements like monotonicity with respect to set inclusion, and they do not have consistent preferences, even in situations involving no risk or uncertainty." A more positive view of rationality is taken by Gintis (2009, p. 2) who argues for maintaining at least a minimum degree of rationality based on *preference consistency*; he argues that violations of preference consistency can be accounted for by changes in context, frames, and reference points.

The impact of the rationality assumption on other social sciences has not been profound. In evaluating the impact of such work, with particular reference to the collected works of Becker (1981), Simon (2000) says the following. "Perhaps its greatest failing was in not producing major insights into political and social processes that were not already well-known, if in somewhat less qualitative and less formalized statements."

In response to these concerns Herbert Simon proposed *bounded rationality* as an alternative approach. The term bounded rationality has become fashionable in a variety of contexts in economics and its usage and meaning have become variegated. To be specific, we follow the original usage of the term in Simon (1956, 1957). Simon used the metaphor of the two blades of a pair of scissors. One blade represents *cognitive limitations* of humans, while the second blade emphasizes the *structure of the environment*. To cut successfully, the scissors must use both blades, hence, for a proper approach to bounded rationality, one cannot choose to omit either cognitive limitations or the structure of the environment. By contrast, the practice in most of economics is to assume unbounded cognitive powers and underemphasize the structure of the environment, such as framing effects, context, and emotions.

Bounded rationality does not imply that individuals are irrational. Hence, bounded rationality is not synonymous with erratic human behavior. Boundedly rational individuals often solve some well-specified optimization problem. However, whether individuals should be modeled as optimizing agents is an open question. For arguments on both sides, see the "Forum on the Role of Bounded Rationality versus Behavioral Optimization in Economic Models" published in the June 2013 issue of the *Journal of Economic Literature*. The main difference between bounded rationality and substantive rationality (as used in neoclassical economics) is that, in the former, the optimization problem respects the cognitive bounds of the relevant economic actors and the structure of the environment. Bounded rationality is, thus, often concerned with *procedural rationality*. It tries to open up the black box of human decision making, and the steps that humans take to formulate and solve problems in different environments.

Bounded rationality has sometimes been interpreted as *optimization under constraints* because the information necessary for taking actions involves the costs of search. In one view, if one introduces constraints that capture the costs and benefits of search, then one could take account of bounded rationality (Stigler, 1961). However, the problem with this argument is that to estimate the costs and benefits of search, one would have to first estimate the costs and benefits of estimating the costs and benefits of search. One can construct an infinite regress of such search, which can be even more informationally demanding than the original problem. This compounds the very problem that these additional constraints were designed to address and makes the model even more psychologically implausible.

To see this slightly more formally, we can use the following intuitive construction in Selten (2001). He distinguishes between *familiar* and *unfamiliar* problems. For familiar problems, the decision maker *knows* the set of alternatives (or trivially knows where to find them) and also knows the *method* for choosing among these alternatives. For unfamiliar problems, the set of alternatives and methods for choosing among them are both unknown. Thus, in this case, the decision making process can be understood as being composed of the following two levels.

1. (Level 1) Find the set of alternatives to be chosen.
2. (Level 2) Find a method of making a choice among the set of alternatives.

A decision maker who is ignorant of what to do at Level 1, will also not know what to do at Level 2. Thus, we get:

3. (Level 3) Find an optimal method of solving the problem at Level 2.

Hence we get an infinite regress, with level k being:

- k . (Level k) Find the optimal method for solving Level $k - 1$.

While this is not a formal proof, it is strongly suggestive of the impossibility of classical optimization in limited time. In the words of Selten (2001) “Trying to optimize in such situations is like trying to build a computer that must be used in order to determine its own design.”

Simon (1978) in his Noble lecture is quite skeptical about theories that explicitly incorporate costly search and information. He offers the following sobering views (pages 358–9): “In none of these theories . . . is the assumption of perfect maximization abandoned. . . . Hence, the new theories do nothing to alleviate the computational complexities facing the decision maker . . . but simply magnify and multiply them. . . . Hence, to some extent, the impression that these new theories deal with the hitherto ignored phenomenon of uncertainty and information transmission is illusory. For many economists, however, the illusion has been pervasive.”

19.14.1 *Aspiration adaptation theory*

In an important class of bounded rationality models, the individual, or an institutional entity such as a firm, has a goal or a reference level, referred to as an *aspiration level* (Simon, 1957). Aspiration levels may refer, for instance, to the target profit of a firm, a target income for a consumer, or even a level of employment/growth rate for a government. A decision alternative is *satisfactory* if it meets or exceeds the aspiration level.

Decision alternatives are not exogenous but must be searched by the individual. To avoid an infinite regress problem, the decision maker can use a simple stopping rule. For instance, stop when an alternative (or the outcome from an alternative) exceeds the aspiration level. Such an

approach has also been called *satisficing* to distinguish it from *optimization*. The aspiration levels are not fixed, but evolve over time.

One simple method of updating aspiration levels is to adjust them downwards if successive alternatives fall below the aspiration level, much as one would reduce the sale price of a house that has been on the market for too long. As Simon (1978) puts it: “In a benign environment that provides many good alternatives, aspirations rise; in a harsher environment, they fall.” The individual has a well-defined optimization problem. Unlike the models in neoclassical economics that are based on substantive rationality, this model incorporates: (1) Limited cognitive abilities that are reflected in a simple stopping rule, and (2) the structure of the environment, as embodied in the formation of aspirations.

An early attempt at a formal model of *aspiration adaptation* was by Saurmann and Selten in 1962; an English version was published as Selten (1998). Consider the role played by aspirations in a symmetric two-player normal form stage game, denoted by $G(2, \{C, D\}, \{C, D\}, u_1, u_2)$, that is repeated over discrete time, $t = 0, 1, 2, \dots$ (Cho and Matsui, 2005). The two players are indexed by $i = 1, 2$, and player i has the strategy set $S_i = \{C, D\}$; think of these actions as, respectively, *cooperation* (C) and *defection* (D).

Denote the action of player i in period t by $s_{it} \in S_i = \{C, D\}$ and denote the action profile in any period t by $s_t = (s_{1t}, s_{2t})$. The respective payoffs of players 1, 2 in the stage game are denoted by u_1 and u_2 such that $u_i : \{C, D\} \times \{C, D\} \rightarrow \mathbb{R}$. We assume that players are sufficiently patient so that if both players were rational in the classical sense, then mutual cooperation (C, C) can be sustained as an equilibrium of an infinitely repeated game. However, we assume that players are boundedly rational.

In the spirit of Herbert Simon’s work, suppose that players are bounded rational and form aspiration levels. Since players do not have infinite cognitive abilities, it is assumed that they summarize the history of the game into a single number, the *aspiration level*, which is simply an average of past payoffs of the player; the precise sequence of outcomes and actions in the past is not relevant. Hence, the aspiration level of player i at time t , a_{it} , is given by

$$a_{it} = \frac{1}{t} \sum_{\tau=1}^t u_i(s_{\tau}), t = 1, 2, \dots \quad (19.65)$$

Rewriting (19.65), we get

$$a_{it} = a_{it-1} + \frac{1}{t} [u_i(s_t) - a_{it-1}]. \quad (19.66)$$

From (19.66), the aspiration level evolves in a manner that is reminiscent of *adaptive expectations*. Current aspirations are equal to lagged aspirations adjusted for the difference between the current actual payoff and lagged aspirations. Initial aspiration levels of the two players are exogenously given by $a_0 = (a_{10}, a_{20})$, and initial actions are also exogenously given by $s_0 = (s_{10}, s_{20})$.

The significance of the aspiration level a_{it} is that it is the payoff which the individual expects from the action s_{it} in the current period, conditional on the past payoffs. So, if past payoffs were high (low), then current aspirations are also high (low). Since player $i = 1, 2$ has only two actions, we denote the chosen action at time t by $s_{it}^+ \in S_i$ and the unchosen action by $s_{it}^- \in S_i$; for example, if the chosen action is C, then $s_{it}^+ = C$ and $s_{it}^- = D$. The individual uses the following *behavioral rule* to choose actions

	C	D
C	1, 1	-1, -1
D	-1, -1	0, 0

Figure 19.7 A coordination game to study the role of aspiration adaptation.

Source: Reprinted from *Journal of Economic Theory* 124(2): 171–201. Cho, In-Koo, and Akihiko Matsui (2005). "Learning aspiration in repeated games." Copyright © 2005, with permission from Elsevier.

$$s_{it+1} = \begin{cases} s_{it}^+ & \text{if } u_i \geq a_{it} \\ ps_{it}^+ + (1-p)s_{it}^- & \text{if } u_i < a_{it} \end{cases}, \quad (19.67)$$

where $0 < p < 1$ and $ps_{it}^+ + (1-p)s_{it}^-$ is a mixed strategy that plays s_{it}^+ with probability p and s_{it}^- with probability $1-p$. The idea behind the behavioral rule is simple. If the current payoff from an action is at least as good as the current aspiration level ($u_i \geq a_{it}$), then continue playing the current action, s_{it}^+ . Otherwise, switch to the other action, s_{it}^- , with some probability $1-p$. Thus, p represents a measure of *inertia*. If $p = 0$, then the individual completely switches to the other action (no inertia) and if $p = 1$, then the individual never switches (full inertia). In order to make the problem interesting, we rule out both possibilities by assuming $0 < p < 1$. One possible cause of inertia is potential switching costs that we ignore for pedagogical reasons.

Given the initial aspiration levels, a_0 , and initial actions, s_0 , we can use the definition of aspirations in (19.66) and the behavioral rule in (19.67) to generate a time profile of aspiration levels and actions $\{a_\tau, s_\tau\}_{\tau=1}^\infty$.

Consider the coordination game in Figure 19.7. If both players coordinate at (C, C), they get the highest payoff of 1 each. Lack of coordination pushes them to the lowest possible payoff tuple $(-1, -1)$. However, if they both do not cooperate, then both get a zero payoff, corresponding to the strategy profile (D, D). Assume that the initial aspiration levels are given by

$$a_{10} = a_{20} = -0.1. \quad (19.68)$$

1. Suppose that the initial strategy profile is $s_0 = (C, C)$. Then, the payoffs of the two players at time $t = 0$ are given by

$$u_1(C, C) = u_2(C, C) = 1. \quad (19.69)$$

Since $u_i(C, C) > a_{i0}$, (19.67) implies that the strategy profile at $t = 1$ is $s_1 = (C, C)$. From (19.66), the aspiration levels in period $t = 1$ are given by $a_{i1} = a_{i0} + [u_i(C, C) - a_{i0}]$. Using (19.68), (19.69), $a_{i1} = 1$. Since $u_i(C, C) = a_{i1}$ in period 1, it follows from (19.67) that in $t = 2$, the equilibrium action profile is again (C, C). The aspiration level in period 2 is $a_{i2} = a_{i1} + \frac{1}{2}[u_i(C, C) - a_{i1}] = 1$. Repeating the usual steps, at $t = 3$, the action profile is again (C, C). Extending this argument, the prediction of the model is that if both players start by coordinating, they continue coordinating. The reason is that the payoff from (C, C) is always at least as good as the aspiration level, hence, the players never feel the need to experiment and choose a different action.

2. Suppose that the action profile (C, D) is played in the initial period; the game is symmetric, so the results with the profile (D, C) are identical. The payoffs of the two players at time $t = 0$ are given by

$$u_1(C, D) = u_2(C, D) = -1. \quad (19.70)$$

From (19.68), (19.70), $u_i(C, D) < a_{i0}$, so (19.67) implies that at time $t = 1$, each player switches to the alternative action with probability $1 - p$. Thus, the actions chosen by the two players at $t = 1$ are given by

$$s_{11} = pC + (1 - p)D. \quad (19.71)$$

$$s_{21} = pD + (1 - p)C. \quad (19.72)$$

We can now compute the payoffs of the two players in period 1 using the mixed strategies of the two players given in (19.71), (19.72). These can be calculated to be

$$u_1(s_{11}, s_{21}) = u_2(s_{11}, s_{21}) = (1 - p)(2p - 1) - p^2. \quad (19.73)$$

The payoffs of both players are equal because of the symmetry of the game.

Consider the polar case of no inertia, i.e., $p = 0$. Then at $t = 1$, using (19.71), (19.72), both players switch their actions, so that the time $t = 1$ action profile is (D, C) . Evaluating (19.73) at $p = 0$, we have $u_1 = u_2 = -1$. When $p = 0$, the aspiration levels of the two players at date $t = 1$ are given by $a_{11} = a_{10} + [u_1(D, C) - a_{10}]$ and $a_{21} = a_{20} + [u_2(D, C) - a_{20}]$. Using (19.68), (19.70),

$$a_{11} = a_{21} = -1. \quad (19.74)$$

Since $u_1 = a_{11}$ and $u_2 = a_{21}$, the behavioral rule in (19.67) implies that the actions at $t = 2$ are the same as those at $t = 1$, so the action profile remains (D, C) . A simple calculation shows that $a_{12} = a_{22} = -1$, so aspirations are not revised in the next period and payoffs do not change either. Thus, given an initial action profile (C, D) , if there is no inertia, the game settles at the profile (D, C) for all time periods, starting at $t = 1$. Analogously, one can show that, with $p = 0$, starting from (D, C) , the game settles at (C, D) from $t = 1$ onwards. Now consider the case when there is inertia, $p > 0$. In this case, from (19.66) and (19.73),

$$a_{11} = u_1 > -1, a_{21} = u_2 > -1 \quad (19.75)$$

for all $p \in (0, 1)$. The behavioral rule (19.67) implies that actions are not altered in the next period ($t = 2$), so s_{12} and s_{22} are given, respectively, by the mixed strategy profile in (19.71), (19.72). Utilities continue to be given by (19.73) in the second period also. The aspiration levels in the second period for $i = 1, 2$ are $a_{i2} = a_{i1} + \frac{1}{2}(u_i - a_{i1})$, where a_{i1} is given by (19.75), thus, $a_{i2} = a_{i1}$. Clearly, the mixed strategies in (19.71), (19.72) are played every period from $t = 1$ onwards. Identical results follow in the case where the initial action profile is (D, C) .

3. Suppose that the initial action profile is (D, D) . The first period payoffs are then $(0, 0)$. Since $0 > a_{i0} = -0.1$, both players' aspiration levels are met and they stick to the profile (D, D)

in period 1. One can use this argument repeatedly to show that the action profile (D, D) is followed in every period. The aspiration level improves in every period and converges eventually to zero, from below.

Based on this analysis, the following observations can be made.

The equilibrium of the game is heavily dependent on the initial state of play. Furthermore, the cooperative outcome can be sustained in an equilibrium of the aspirations adaptation game. For the parameter values chosen in the example, this outcome arises if players begin by cooperating in the initial period. Non-cooperation or defection, however, begets further rounds of non-cooperation. Finally, players might coordinate on the defection strategy and never cooperate again.

In a more general model, Cho and Matsui (2005) introduce other extensions. They introduce a small probability of players playing a different action; one interpretation is that players make a mistake with some positive, but small, probability. It turns out that the inefficient coordination outcome (D, D) is not robust to this perturbation and the only equilibrium which survives is the “good” coordination equilibrium.

It should be quite intuitive that this framework can be used to support the cooperative outcome in a range of games (Colman et al., 2010). For instance, in the prisoner’s dilemma game, if aspiration levels are sufficiently high, then the non-cooperative outcome results in payoffs that are lower than the aspiration level. Individuals then experiment with other actions (such as cooperation). If both players happen to cooperate at the same time, then the payoffs exceed the aspiration level, in which case there is no further need to experiment with other actions. These ideas can be formalized in different ways. Börgers and Sarin (2000) and Karandikar et al. (1998) are two such attempts. Both papers contain the idea that aspirations are updated slowly in the direction of the eventual cooperative payoffs. Karandikar et al. (1998), in addition, introduce a small probability of trembles in actions and demonstrate in 2×2 games that most players will cooperate most of the time.

Aspirations in these models are based on plausible adaptive processes. However, ultimately empirical evidence will determine the precise nature of the formation of aspirations, which might well be context and situation dependent. In one interesting experiment, Bernard and Taffesse (2014) randomly selected households in Ethiopia to watch inspirational videos based on the actual life stories of people who had successfully followed goals and achieved a higher socio-economic position. Six months later, these individuals had higher savings rates and higher aspirations, particularly about their childrens’ educational future. It is also well known that poverty depresses aspirations and engenders a feeling of greater hopelessness (Ray, 2006; Haushofer and Fehr, 2014). Similar results are found in richer countries where poorer students have relatively lower aspirations towards education and employment (Guyon and Huillery, 2014).

In the models above, there is a single goal, feedback is perfect, and the underlying economic environment is commonly understood. In actual practice, however, individuals, institutions, and governments, often try to achieve complex long-term goals. For instance, (i) individuals might be interested in long-term happiness, (ii) firms in long-term profits, and (iii) governments in long-term economic prosperity. Yet it might not be clear how to achieve these complex goals. On the other hand, there could be a set of short-term goals that the individual believes are correlated with long-term goals. Corresponding to the three long-term goals mentioned above, examples of relevant short-term goals could be: the individual practices daily meditation, the firm invests in short-term improvements in customer relations, and the government appoints a particularly capable central banker.

Selten et al. (2012) outline a formal model of the decisions taken by a monopolist, which has these features; an approach that they call *goal systems*. In this framework, economic entities face a complex problem whose solution is unlikely to be obtained by purely rational analysis, although a simulated solution is available by using specialized optimization software. Economic entities then choose among a set of goal variables that are essentially short-term feedback variables. For instance, the monopolist chooses short-term quality of the product and advertisement expenditure that contributes towards raising demand for a product in that period. For each goal, they set an aspiration level that they can adapt upwards or downwards. If no upward adjustments in aspiration levels are feasible, then no further attempts at adjusting aspiration levels is made. Their experiments confirm some of the basic predictions of the aspiration adaptation model.

19.14.2 *Fast and frugal heuristics*

Humans do not possess unlimited cognitive abilities, nor do they have unlimited time to solve problems. Thus, as Herbert Simon suggested, modeling *procedural rationality*, i.e., the cognitive process and the decision environment in which decisions are made is important.

Rubinstein (1998) provides several examples of procedural rationality. In chapter 11 of his book, titled *Final Thoughts* he provides a summary of his correspondence with Herbert Simon regarding his book. Simon objects to the lack of empirical support for the rationality procedures introduced by Rubinstein. Simon is quoted as saying “Facts do not come from the armchair, but from careful observation and experimentation” Rubinstein (1998, p. 188). Rubinstein’s own response is detailed in his chapter 11, but the following remark illustrates well his position. “The crowning point of making microeconomic models is the discovery of simple and striking connections between concepts (and assertions) that initially appear remote” (Rubinstein, 1988, p. 191). Clearly, Rubinstein’s approach, while important and timely in terms of its proposed formalism, is not in the same tradition as Herbert Simon’s. Indeed, there is a basic disagreement about the method and purpose of economic models; this book’s view has already been laid out in the introductory chapter.

An approach that is closer to Herbert Simon’s own views can be found in the work of Reinhard Selten and Gerd Gigerenzer and colleagues at the ABC Research Group, based at the Max Planck Institute of Human Development. Central to the work of this group is the intriguing metaphor that humans draw upon an *adaptive toolbox of heuristics*, conditional on the environment that they find themselves in. Two excellent summaries of this approach are Gigerenzer et al. (1999) and Gigerenzer and Selten (2001). This approach differs fundamentally from the heuristics and biases approach of Kahneman and Tversky.

The view taken by Gigerenzer and colleagues is, in many respects, similar to the *Enlightenment view*, associated with the names of Laplace, Boole, and Piaget; see Gigerenzer et al. (1999) for a representative position. In this view, the normative and descriptive models of human inference are the same as the relevant statistical models. So, for instance, in updating probabilities, the expectation is that humans will use the Bayes’ rule.

On the other hand, the work of Kahneman, Tversky, and others (see above) has shown a divorce between the normative view of humans and actual behavior (Kahneman and Tversky, 2000). So, for instance, the behavior of individuals might not conform to the predictions of subjective expected utility theory (SEU), or to the predictions of classical statistics. In his Nobel lecture, Simon (1978, p. 362) is very explicit about the implications: “On the basis of these and other pieces of evidence, the conclusion seems unavoidable that the SEU theory does not provide a good prediction—not even a good approximation—of actual behavior.”

Gigerenzer and colleagues take the position that human heuristics are not a *bad* but a *good*. They argue that these heuristics are the appropriate response by humans, with limited cognition and time, to their environment. This is reflected in the title of one of their books: *Simple Heuristics That Make Us Smart*. They argue that the heuristics used by humans can outperform other statistical approaches to solving problems. Gigerenzer and colleagues also challenge the Kahneman–Tversky inspired biases on several other grounds. For instance, they argue that if the problems are presented to subjects in a natural frequency format (as opposed to probabilities) and in a knowledge domain familiar to them, then the biases are significantly reduced or eliminated. We have already examined the evidence on some of these positions in the sections above; we postpone a further discussion of these issues to Section 19.15.

TWO EXAMPLES OF FAST AND FRUGAL HEURISTICS

We start with two examples that illustrate fast and frugal heuristics. The perceptive reader will note that there is no necessary conflict in the insights from these problems with the Kahneman–Tversky approach, an issue that we return to in Section 19.15.

Example 19.6 (*The problem of catching a ball*): Consider, as in Figure 19.8, the problem of catching a ball in cricket, baseball, or football (by a fielder in the first two cases and by the goalkeeper in the third). Suppose that the origin of the ball is at point A. It travels along the thick parabola and lands at point B. A catcher is stationed at point C. How does he catch the ball? The optimizing, fully rational, approach in economics would suggest the following algorithm.

Use the two-parameter equation of a parabola. Take the point of origin of the ball and the initial velocity as the initial conditions of the problem. In order to determine the two parameters of the parabola, take account of factors such as the wind speed, humidity, and spin on the ball. Also account for the fact that these conditions could create perturbations from the shape of an ideal parabola. Then design a machine (unless one is handy) that takes account of all this information in order to compute point B. Since the shortest distance between two points is a straight line, the prediction of the rational model is that the catcher runs in a straight line from C to B to successfully catch the ball.

Of course, a weaker position, and the one typically taken in neoclassical economics, is that the calculation process need not be taken literally. Individuals are expected to behave “as if” they can perform the appropriate calculations.

In either case, it is very easy to test this prediction. When the runs of successful catchers are observed, it turns out that they do not run in a straight line from C to B (McLeod and Dienes, 1996). Rather, they start running as soon as the ball leaves point A and take the curved, dotted, path shown in Figure 19.8. Hence, the prediction of the standard model is easily rejected in the strict and the weak interpretations.

McLeod et al. (2006, p. 139) describe why catchers behave in this manner: “the path of a fielder running to catch a ball is determined by the attempt to satisfy 2 independent constraints. The 1st is to keep the angle of elevation of gaze to the ball increasing at a decreasing rate. The 2nd is to control the rate of horizontal rotation necessary to maintain fixation on the ball.” The important feature of this heuristic is that it does not use all the available information. Hence, it is frugal in the use of information. Because it uses only two simple constraints, the heuristic is incredibly fast (the catcher starts running immediately).

Example 19.7 (*Dealing with heart attack patients*): Suppose that the emergency unit of a hospital receives a patient with chest pains and a suspected heart attack. How best should




Figure 19.8 The problem of catching a ball.




Figure 19.9 Classifying the degree of risk of suspected heart attack patients.

Source: Peter M. Todd and Gerd Gigerenzer (2000). "Précis of simple heuristics that make us smart." *Behavioral and Brain Sciences* 23: 727–41.

they deal with the patient? One possibility is to measure all the possible indicators of a heart attack. Then use some statistical procedure to aggregate this information into a number; when this number exceeds some critical value it is almost sure to correctly diagnose a heart attack. However, time, information and cognitive abilities are limited.

Breiman et al. (1993) suggest a simple decision tree, which is able to classify patients by the degree of risk that they face; see Figure 19.9. Only three variables are used. If the systolic blood pressure is below 91, the patient is immediately classified as high-risk and no further information is needed. If not, then patients older than 62.5 years and suffering from sinus tachycardia are classified as high risk. Otherwise they are low risk.

The distinguishing feature of this simple decision tree, which really is a heuristic to identify high-risk patients, is that it does not try to use all possible information that could have been used. Furthermore, it is simple in the following two respects. It uses qualitative information with binary answers to decision criteria. So, for instance, the absolute difference in one's actual age from 62.5 years is irrelevant; only the cutoff value of 62.5 years is important. The heuristic does not attempt to combine the three variables into a single index. Hence, this heuristic is also fast and frugal. But does it work? Breiman et al. (1993) show that it is actually more accurate

than some other complex statistical procedures that are currently used. Hence, fast and frugal heuristics are not necessarily inaccurate or inefficient, but they can be in other cases.

The main difference between *satisficing* in the sense of Herbert Simon, and *fast and frugal* heuristics is this: satisficing need not always be a fast and frugal heuristic. Some forms of satisficing might even be computationally challenging, while fast and frugal heuristics never are.

A distinguishing feature of the work on fast and frugal heuristics has been the attempt to (1) pit the predictions of such models against the data, and (2) test the efficiency of such heuristics in comparison with standard statistical procedures. As an illustration of this approach, we first consider some examples from Gigerenzer and Goldstein (1996).

THE RECOGNITION HEURISTIC AND THE TAKE-THE-BEST HEURISTIC

In this section we describe a prominent example of a fast and frugal heuristic. Gigerenzer and Goldstein (1996) ask how individuals might solve the problem of comparing the relative size of different German cities? For instance, which city has a higher population? (a) Hamburg (b) Cologne. They suggest the following procedure that uses limited knowledge and limited cognitive abilities to make a relatively fast decision.²⁰

Consider a *reference class* R of cities; this is just a list of cities. For $a, b \in R$ we ask the individual: which is bigger? In answering this question, individuals tap into some *knowledge base*. Essentially, a knowledge base, C , consists of *probability cues* $C_0, C_1, C_2, \dots, C_n$. We focus, here, on cues that take a binary value. For instance, the cue C_i could be: does the city have a football team? The answer is either yes, no, or don't know; we denote these three cases, respectively, by $+$, $-$, $?$. As an example, suppose that we have four reference cities and five cues, so $R = \{a, b, c, d\}$ and $C = \{C_0, C_1, \dots, C_5\}$. Let the first cue, C_0 , be the *recognition cue*, i.e., does one recognize a city? The individual can either recognize ($+$), not recognize ($-$), or does not know if he can recognize ($?$). The cue values for each of the four cities are shown in Table 19.9.

The table illustrates two features. First, it uses *limited information* (in actual practice many other cues could have been used). Second, it shows *limited knowledge*. This is shown by the large number of $?$ signs, which indicate that the cue value is uninformative. Notice that for city d , the recognition cue indicates that it has never been heard of by the individual. For this reason, for all other possible cues, d takes a value $?$. Let $C_i(x)$ denote the cue value taken by the i th cue when the city is $x \in R = \{a, b, c, d\}$. Let $t(x)$ be the question asked from the individual or the *target variable*; in this case this is the population of city x . Suppose that for any two cities $x, y \in R$ and

Table 19.9 The cue values for each of the four cities.

	a	b	c	d
C_0 (Recognition)	+	+	+	—
C_1	+	—	?	?
C_2	?	+	—	?
C_3	—	+	?	?
C_4	?	—	—	?
C_5	?	?	—	?

²⁰ This draws on earlier work on probabilistic mental models; for instance, Gigerenzer (1993), Gigerenzer et al. (1991).

using cue C_i we get that $C_i(x) = +$ and $C_i(y) = -$. Then we say that cue C_i is able to predict correctly which city has the larger population, i.e.,

$$C_i(x) = +, C_i(y) = - \Rightarrow t(x) > t(y). \quad (19.76)$$

Now suppose that over all $x, y \in R$ we compute the number of cases, N , where (19.76) holds. Let M be the number of cases where it does not hold, i.e., where, for instance, $C_i(x) = C_i(y) = +$ and $C_i(x) = C_i(y) = -$. Ignore the cases where a cue takes a value ? (“don’t know”) because no comparison can be made.

Two features of cues, *ecological validity* and *discrimination rates*, are important. These are defined below.

Definition 19.4 (*Ecological validity*): The ecological validity v_i of cue C_i is the frequency with which it predicts the target from among the choices in the reference set, i.e.,

$$v_i = \frac{N}{N + M}.$$

So, for instance, let R be the set of German cities with a population of at least 100,000. Let cue C_i check if the city has a football team. In 87% of the cases, the city with the football team has greater population. Thus, cue C_i has an ecological validity of 0.87. Applying similar calculations to all cues we can rank the cues in term of their ecological validity.

Ecological validity reflects the phenomenon of limited information but not limited knowledge; it does not take account of the ? values. Now suppose that there are a total of $\bar{N}$ pairs of cities $x, y \in R$ such that x, y are distinct. Also let $\hat{N}$ be the number of cases (among all $x, y \in R$) such that the cue is able to discriminate between two options, i.e., either $C_i(x) = +, C_i(y) = -$ or $C_i(x) = -, C_i(y) = +$.

Definition 19.5 (*Discrimination rates*): The discrimination rate, d_i , associated with some cue C_i is the relative frequency with which a cue discriminates among two objects in the reference class. Hence,

$$d_i = \frac{\hat{N}}{\bar{N}}.$$

As d_i increases, the *inferential usefulness* of the cue, C_i , increases because it is able to distinguish among a larger number of cases.

Reverting back to the data provided in Table 19.9, how will an individual make a choice? One possibility is for individuals to use the take-the-best algorithm. In this algorithm, the cues are ranked in terms of their ecological validity first.²¹ Following the recognition cue, the remaining cues are arranged in the following order. The cue with the highest validity (C_1) at the top (also called the best cue) and the one with the lowest validity (C_5) at the bottom. The basic principle of this algorithm is to take-the-best and “ignore the rest.” The steps followed in the take-the-best algorithm are shown in Figure 19.10.

²¹ Another heuristic, the *minimalist heuristic*, ranks cues in a completely random order.




Figure 19.10 Description of the take-the-best algorithm.

Step I. The first cue that is used is the recognition heuristic. In making a choice among two objects, if one of the objects is recognized and the other is not (+−), then choose the recognized object. If both objects are unrecognized (−−), then guess. If both are recognized (++), then go to step II.

Step II. Use the cue with the highest ecological validity. If this cue discriminates among the two options $x, y \in R$ (i.e., $C_i(x) = +, C_i(y) = -$ or $C_i(x) = -, C_i(y) = +$ or $C_i(x) = +, C_i(y) = ?$ or $C_i(x) = ?, C_i(y) = +$), then choose the option that has a value + for the cue. Otherwise, if the cue is unable to discriminate, examine the cue with the next highest ecological validity and repeat step II.

Step III. If one runs out of cues and none is able to discriminate between the two options, then simply guess.

These steps, the definitions of ecological validity, and Table 19.9 can be used to illustrate several examples.

Example 19.8 Which of the two cities, a, d , has more inhabitants? Using Step I, the recognition cue, city a is recognized while city d is not, hence, the choice is made at the first step of the algorithm: city a is chosen.

Example 19.9 Which of the two cities, a, b , has more inhabitants? Using Step I, the recognition cue is unable to discriminate between the two cities. Using Step II, we now move on to the cue with the next highest validity, i.e., cue C_1 .²² It immediately distinguishes between cities, a, b because $C_1(a) = +, C_1(b) = -$. Hence, city a is chosen.

Example 19.10 Which of the two cities, b, c , has more inhabitants? In Step I, the recognition cue is unable to distinguish between the two (both are recognized). Cue C_1 , with the next highest validity, is also unable to distinguish because the cue value for city c is unknown ($C_1(b) = -, C_1(c) = ?$). The cue with the next highest validity, C_2 , is able to distinguish ($C_2(a) = +, C_2(b) = -$), and on that basis, city b is chosen.

These examples share some common features. First, only a limited part of the total stock of knowledge is searched. Second, the stopping rule is simple: “stop when the first discriminatory

²² Cues C_1, C_2 , and C_3 have identical ecological validity in this example.

cue is found.” Third, no attempt is made to integrate all the available information, using classical statistical procedures. These are all distinguishing features of fast and frugal heuristics. The take-the-best algorithm is sometimes said to be *non-compensatory* because the discriminating cue with the highest validity determines the final choice; no combination of the other cues can overturn this choice.

The take-the-best algorithm has been used successfully in several contexts. These include: the *confidence frequency effect*, the *less is more effect*, and why the confirmation bias explanation of the overconfidence effect may not be supported by the evidence.²³

In order to test the take-the-best algorithm in the context of city size, Gigerenzer and Goldstein (1996) consider 83 German cities with at least 100,000 inhabitants. Population is the target variable; 9 binary ecological cues are used, so there are 9×83 cue values. These cues included the following: Is the city a national capital? Does the city have a football team in the major league? Is the city on the intercity train route? Is the city a state capital? Is the city home to a university? Is the city in the industrial belt? The ecological validity of the cues is decreasing in the order written above.

The point of the exercise is to compare the take-the-best heuristic with some other statistical procedures, all of which use more information, more time, and more computation power. The statistical procedures used include tallying, weighted tallying,²⁴ unit weight linear model, weighted linear model,²⁵ and multiple regression. The alternatives against which these heuristics are compared are chosen by the experimenter. There need not be universal agreement that these

²³ For the *confidence frequency effect*, see Gigerenzer et al. (1991); the original take-the-best algorithm is also presented in this paper. For the *confirmation bias explanation of the overconfidence effect*, see Gigerenzer et al. (1991).

²⁴ Under tallying, first arrange the cues in decreasing order of validity. The recognition cue is not used. The idea is to take account of the (positive) weight of the evidence for a decision. So, for instance, in comparing cities a , b , tally the positive cue values for each city across all cues. Then choose the city with the largest number of positive cue values. Thus, city a is chosen over city b if

$$\sum_{i=1}^n \phi_i(a) > \sum_{i=1}^n \phi_i(b).$$

where for any city $x \in R$,

$$\phi_i(x) = \begin{cases} 1 & \text{if } C_i(x) = + \\ 0 & \text{if } C_i(x) = - \\ 0 & \text{if } C_i(x) = ? \end{cases}$$

In the event that both numbers are the same, one simply guesses.

²⁵ In the weighted linear model, city a is chosen over city b if

$$\sum_{i=1}^n v_i \phi_i(a) > \sum_{i=1}^n v_i \phi_i(b),$$

where v_i is ecological validity, and for any city $x \in R$,

$$\phi_i(x) = \begin{cases} 1 & \text{if } C_i(x) = + \\ -1 & \text{if } C_i(x) = - \\ 0 & \text{if } C_i(x) = ? \end{cases}$$

In the event that both numbers are the same, one simply guesses.

are the best possible alternatives. This is an important drawback in the claimed superiority of fast and frugal heuristics over other possible solutions.

The main empirical results in Gigerenzer and Goldstein (1996) are as follows. Suppose that 100% of the cue values are known. In this case, the take-the-best algorithm generated as many correct inferences as any of the alternative algorithms under consideration. Now suppose that we consider the following percentages of known cue values: 10, 20, 50, 75, 100, and take the average number of correct inferences across all of them. Again, the take-the-best algorithm performs as well as any of the alternative algorithms used. This result is strengthened if we note that these figures ignore the costs of using the alternative algorithms. These costs include the relatively greater costs of acquiring and using additional information, and the greater computational time required relative to the fast and frugal take-the-best algorithm.

If only the recognition heuristic is used (i.e., choose the recognized object over an unrecognized object), then the accuracy across all pairs of choices is 65%, which is significantly above purely random choice. Finally, the authors find an intriguing *more is less effect*. Giving the decision maker more information (i.e., increasing the proportion of recognized objects) beyond some intermediate level reduces the accuracy of decision making.

Oppenheimer (2003) offers one criticism of the use of the recognition heuristic in the context of judging city size. He notes a problem with the original studies; they appeared to have conflated the size of the city with its recognition. For instance, Berlin is recognized and known to be large. The recognition heuristic predicts that in a comparison of the city size in which one city is recognized and known to be a small one, and the other city is unrecognized, the former should be chosen. By introducing names of fictitious cities so that many would not be recognized, Oppenheimer (2003) finds the opposite result: Namely, that the unrecognized cities are chosen over “recognized and known to be small cities”. This suggests a more careful consideration of the robustness of the method and the claims.

The work of Gigerenzer and Goldstein (1996) is related to the work of Payne et al. (1993) who tested the performance of several heuristics in a choice between various risky gambles. They mainly looked at the following two one-decision heuristics that are very similar to the take-the-best heuristic. (1) The LEX heuristic in which cue order is determined not necessarily by ecological validity, unlike in the take-the-best heuristic, but by some measure of importance. (2) The LEXSEMI heuristic, similar to LEX except that the difference between cue values must exceed a certain threshold for a cue to be able to discriminate.

They found that different heuristics performed best in different environments, hence, the use of heuristics is strongly dependent on the environment, as Herbert Simon had proposed. Both LEX and LEXSEMI performed quite well in comparison with the weighted additive model (similar to the weighted linear model described above). Under severe time pressure, the LEX normally outperformed the others because it requires fewer computational steps.

The explanation of advertising has puzzled economists for a long time; none of the existing explanations in neoclassical economics is completely persuasive. However, the recognition heuristic might provide one important clue to an explanation. If consumers use the recognition heuristic in choosing among alternative products, then it is worthwhile for companies to spend significant sums of money on, often uninformative, but recognition-enhancing, advertising. The high-profile publicity campaigns of politicians, cities, or travel destinations can perhaps also be explained in a similar manner.

Gigerenzer and Goldstein (1996) take the cues as given; indeed these are often experimenter provided. In many contexts, the cues must be found in the first place. How do individuals search

for cues. One possibility is that individuals use fast and frugal heuristics to undertake the search in the first instance; see Todd (2001) for examples.²⁶

In another intriguing study in the context of stock markets, Borges et al. (1999) pit the simple recognition heuristic against highly trained fund managers. Over the period 1996–1997, individuals were asked to invest in a variety of American and German stocks by using only the recognition heuristic. Namely, between two stocks, if one is recognized and the other is not, then choose the recognized stock; two of the four authors used their own money to fund this stock market experiment. In competition with the recognition heuristic were pitted several classical benchmarks for stock market selection such as mutual funds, market indices, and chance (or dartboard portfolios).

The authors found that the average level of stock market return on a randomly chosen portfolio was consistently below the return arising from the use of the recognition heuristic. Portfolios of highly recognized stocks outperformed those that were unrecognized. The performance of the recognition heuristic was also better (in six out of eight tests) relative to two major managed funds, the American Fidelity Blue Chip Growth Fund and the German Hypobank Investment Capital Fund. This is suggestive of a relationship between recognition, market share, and profitability of companies; Borges et al. (1999) also review some empirical evidence in this regard. These interesting results are only suggestive, and further empirical evidence is required. For instance, the results were obtained in a strong bull market and the time horizon for the performance of the stocks was only one year. These design features need to be relaxed in future work, certainly in a manner that most experimental economists will find persuasive.

Czerlinski et al. (1999) examine the performance of the take-the-best heuristic and the *minimalist heuristic* against statistical benchmarks such as the linear regression model. The minimalist heuristic is identical to the take-the-best heuristic except that at the first stage, the cues are not arranged in any particular order (such as ecological validity) but chosen randomly. They consider 20 different environments. One of these environments is the relative population sizes of two cities in a binary comparison (this is as in Gigerenzer and Goldstein, 1996). However, other environments include (all using US data): high school dropout rates, rate of homelessness, mortality rate, house prices, rents, and obesity.

In each of these environments, individuals are given sensible cues. For instance, when individuals were asked to predict the high school dropout rates in 57 Chicago public high schools, they were given the following cues: percentage of low-income students, percentage of nonwhite students, and average SAT scores. Take-the-best often outperformed multiple regression while employing only a third of the cues used by multiple regression. Even the minimalist heuristic came close to the accuracy of the multiple regression model. The relatively good performance of these simple heuristics seems to be the result of the excellent correlation of the cues with the target variable. However, the cues are experimenter determined and provided, which compromises the bite in the results.

Heuristics often appear to outperform multiple regression in a variety of contexts, and in particular, when one looks at out-of-sample prediction (Gigerenzer et al., 1999). The intuition is that multiple regression, while taking into account the salient features of the problem, also fits the noise in the sample. Consider an extreme case in which, outside the sample, such noise is absent or the nature of such noise is very different. Then there might be large prediction errors when an

²⁶ The example of a new colleague who joined our university from the US is instructive. He was interested in a broadband, TV, and telephone connection. He asked a few selected colleagues what they had chosen; the answers allowed him to immediately narrow down his search process quite substantially.

attempt is made to fit the multiple regression to out-of-sample points. On the other hand, simple heuristics do not respond to noise. By relying on the salient, they might be able to predict better outside the sample. However, this could also work to the detriment of the heuristics when there is no noise in the sample but much greater noise out of the sample.

In his “Mais lecture,” Mervyn King (2005), the then Governor of the Bank of England, makes a case for using fast and frugal rules in central banking. His argument is that fully optimizing rules are too complex to fit in with the evidence on limited human cognition. Furthermore, as the state of the economy and the state of knowledge in economics changes, often at a rapid rate, problems of limited cognition and limited time become even more severe. Hence, an appropriate response is for both actors, the central bank and individual forecasters in the economy, to employ useful heuristics.

The central bank could adopt the following heuristic, depending on the environment. In normal times, set interest rates such that the expected inflation in two years time is equal to an inflation target. However, when the economy has been hit with a large shock, then the heuristic is to bring inflation back to target over a longer period. Hence, in this approach, the inflation target provides a nominal anchor around which the understanding of such a heuristic might become common knowledge. King also suggests heuristics that might be ideal (from the central bank’s point of view) for the individual forecasters to use. The ideal heuristic is for individuals to expect inflation to equal its target. An example of a bad heuristic is that if inflation is different from the target, then it is expected to drift further away. These heuristics are fast and frugal, and expected to be relatively conducive to economic stability.

There is also a literature on using heuristics to make choices among objects with multiple attributes (Payne et al., 1993). Payne and Bettman (2001) consider the problem of choosing among several brands of cars with different attributes (e.g., reliability, price, safety, horsepower). A key part of the problem is that there is conflict among the various attributes, so none of the cars dominates the others on all attributes. Furthermore, decision makers could have several goals such as maximizing accuracy, ease of justification of the decision to others, minimizing cognitive effort, and negative decision-related emotions. Furthermore, decision makers seem to use different goals in different settings, and at different times. How should a decision maker choose?

Several strategies are considered in the heuristics literature. One possibility is the *weighted adding strategy* (WADD) in which the score of each car is added over all attributes using some subjective weights (Edwards and Tversky, 1967). Another heuristic that is computationally simpler is the *equal weighting strategy* (EQW) that assigns unit weights to all attributes. The LEX strategy (see above) simply chooses the most important attribute, say, reliability, and then chooses the car with the highest value on this attribute.

Simon’s (1955) *satisficing strategy* (SAT) can also be used to make a choice in this case. Start with a base vector of attribute values that summarizes the aspiration levels for each attribute. Then pick the first car that does better than the base vector on all attributes. The *elimination by aspects* (EBA) strategy combines the LEX and SAT strategies. In this case, first choose the most important attribute and set an aspiration level for it. Then eliminate all cars that fall below this attribute value. Repeat this step for the next most important attribute, and so on, until a unique choice remains; if the choice is not unique, then some tie breaking rule is used (Tversky, 1972).

Payne and Bettman (2001) find that the heuristic which performs best is heavily contingent on the environment. The WADD strategy is the most robust in terms of accuracy across the various environments. But there are environments where the LEX achieves about 90% of the accuracy of the WADD with only about 40% of the effort. For more examples of fast and frugal heuristics and their performance, relative to statistical benchmarks, see Gigerenzer and Gaissmaier (2011).

We now offer a critical assessment of this approach.

1. The evidence provided by Gigerenzer and colleagues does not establish the conditions under which the heuristics proposed, such as take-the-best algorithm, will outperform the relevant statistical procedures or optimization approaches. They argue that the heuristics used by humans can outperform other statistical approaches to solving problems. This may often occur when the optimal benchmark is unclear. No heuristic can outperform an optimal solution when the optimal solution is unambiguously known. For instance, when the demand and cost curves are known, no heuristic can do better than the optimal solution: marginal revenues equal marginal costs. However, things are less clear if these curves are unknown or when one is interested in predicting stock market prices; in these cases the optimal solution with which to compare a heuristic is unclear. Furthermore, these results neither include the costs of finding the cues or the additional costs of information collecting, integrating, and computing associated with statistical procedures that are not fast and frugal; the latter issue is often considered in the theoretical computer science literature (Martignon, 2001).
2. In most studies of fast and frugal heuristics, the heuristics are proposed by the researchers or experimenters themselves. However, it needs to be established if individuals actually use these heuristics in actual practice. Some progress was made in this direction by Goldstein and Gigerenzer (1999). They gave individuals additional information that contradicted the decision that would have been made on the grounds of the recognition heuristic. They found that, nevertheless, 92% of the choices made by the participants still agreed with the recognition heuristic. It remains an open question whether individuals do follow the other fast and frugal heuristics considered above. In contrast to this, the thrust of the Kahneman–Tversky approach was to discover the “actual” heuristics that people follow.
3. How do individuals decide on the cues in a particular environment? One possibility is that the cues are searched using some simple heuristics. Perhaps we engage in simple learning that associates past cues with the outcomes. This is the basis of models of reinforcement learning that we consider in Part 5 (Erev and Roth, 2001). But it is also possible that we have, over the course of human evolution, acquired the ability to generate simple cues and heuristics. The problem of catching a ball (see Example 19.6) is a nice example of this phenomenon. However, it is less clear how these insights can be used to solve economic problems such as the design of contracts, choice of pension plans, choice of a mortgage, and choice of a monetary policy, for which evolutionary experience may provide only limited help. It is not immediately clear which fast and frugal heuristics to use in these cases. Unless these issues are directly addressed, they will limit the scope, and practical usefulness, of this work for economists.

19.15 The great rationality debate

One objection raised by some critics of Kahneman and Tversky’s heuristics and biases program (henceforth, KT) is that it overly focuses on human biases and paints a pessimistic picture of human inference and prediction. By contrast, it is argued, the main thrust of the approach by Gerd Gigerenzer and colleagues (henceforth, G&C) is more positive because it shows that heuristics outperform more sophisticated forms of statistical reasoning. Kahneman and Tversky’s approach came under great criticism only a few years after publication; see, Cohen (1981),

Einhorn and Hogarth (1981), Lopes (1991), Gigerenzer (1991, 1993, 1994), and Gigerenzer and Hoffrage (1995).

This debate between KT and G&C has come to be known in academic circles in psychology as the *Great Rationality Debate*; for surveys, see Stanovich and West (2000) and Stanovich (2012). Within psychology, the debate has been extremely acrimonious. Consider the following list of comments by psychologists on Kahneman and Tversky's work, taken from Stanovich and West (2000, p. 649):

Many critics have insisted that in fact it is Kahneman and Tversky, not their subjects, who have failed to grasp the logic of the problem (Margolis 1987, p. 158).

If a "fallacy" is involved, it is probably more attributable to the researchers than to the subjects (Messer and Griggs 1993, p. 195).

When ordinary people reject the answers given by normative theories, they may do so out of ignorance and lack of expertise, or they may be signaling the fact that the normative theory is inadequate (Lopes 1981, p. 344).

In the examples of alleged base rate fallacy considered by Kahneman and Tversky, they, and not their experimental subjects, commit the fallacies (Levi 1983, p. 502).

What Wason and his successors judged to be the wrong response is in fact correct (Wetherick 1993, p. 107).

Perhaps the only people who suffer any illusion in relation to cognitive illusions are cognitive psychologists (Ayton and Hardman (1997, p. 45).

These comments are sharp, some even vituperative. The discussion in Stanovich and West (2000) and the commentary following their excellent article illustrates just how muddled the debate really is. Part of the deadlock on these issues in psychology arises from disagreement over (1) the precise transmission channels in the brain through which judgment heuristics manifest (e.g., whether it is dual process theories or something else?), (2) the normatively correct model (as opposed to the rational benchmark in economics). These issues are not critical in economics so we omit them; interested readers can consult Stanovich and West (2000) and Stanovich (2012). Instead, we consider the remaining sources of disagreement below.

KT and G&C address very different questions. KT ask if people behave in the manner prescribed by the rational benchmark in economics. The rational benchmark implies compliance with subjective expected utility theory, and the classical rules of statistics and probability theory. The main finding in KT, that we extensively review in this chapter, is that a significant percentage of people do not act in a manner that is consistent with the rational benchmark.

G&C ask if there exist fast and frugal heuristics that still do better than "some statistical benchmarks" that rely on much greater information and computational requirements? They find the answer in the affirmative. However, typically little evidence is provided in this literature that people do indeed employ the experimenter-provided fast and frugal heuristic in real life. It is also not always clear what the proper statistical benchmark is against which a particular fast and frugal heuristic is compared.

Consider, for instance, the work that pits the recognition heuristic against several statistical benchmarks (see above) to determine which method earns greater stock market profits (Borges et al., 1999). Given that the efficient markets hypothesis is often rejected, the relevant benchmark that best predicts actual stock market prices is not clear. Indeed, economists and financial firms will not necessarily agree with the benchmarks employed in this literature. Hence, declaring victory for the recognition heuristic in this context may be premature. The only sensible inference that can be drawn is that for the set of benchmarks chosen by the experimenter and for

the chosen year and stock portfolio, the recognition heuristic did better. On the other hand, the full-optimization statistical benchmarks, even if they are known to exist, are likely to be computationally complex and informationally demanding. This is not a problem for financial firms who possess the computational firepower. However, for individual stock market participants with limited cognition and resources, the G&C approach poses appropriate questions, hence, imparting great usefulness to the approach.

Given the very different nature of the core issues dealt with in the KT and G&C frameworks, it is staggering that so many of the leading researchers could take such strong positions and fish in very muddy waters. Indeed, it is very likely that both KT and a less rhetorical G&C view of the world are, in fact, correct. The following two quotes from Kahneman (2000, p. 682) show just how close the two sides in the debate really are: (1) “Contrary to a common perception, researchers working in the heuristics and biases (HB) mode are less interested in demonstrating human irrationality than in understanding the psychology of intuitive judgment and choice.” (2) “All heuristics make us smart, more often than not ...”

There are, however, other, non-core issues that have to do with the nature of the evidence, where a debate is possible. A good source for this debate is Kahneman and Tversky (1996) and Gigerenzer (1996a). Some of the arguments against KT can be dismissed easily.

1. It has been argued that experimental evidence involving probabilities based on *singular events* is suspect because the frequentist approach requires a large number of repetitions (Gigerenzer et al., 1989). For instance, whether Linda is a bank teller or a feminist is a singular event. Whether a particular candidate will win the next Presidential elections in the US is also a single event. In this case, strictly speaking, one can speak only about subjective probabilities in a Bayesian sense (Gigerenzer et al., 1989). The argument is made that since only subjective probabilities are involved, any subjective probability will do, so the experimental results are uninformative about the violation of a rationality-based model. Kahneman and Tversky (1996) note that this issue only applies to 2 out of the 12 biases that they highlight. Furthermore, the responses show systematic biases in one direction, which weakens the case for arbitrary subjective probabilities.
2. It has been alleged that KT ignore the context, content, and framing effects in their demonstration of biases (Gigerenzer et al., 1988; Gigerenzer, 1993). However, this appears to be a misperception. Kahneman and Tversky (1996, p. 583) write: “the assumption that heuristics are independent of content, task and representation is alien to our position, as is the idea that different representations of a problem will be approached in the same way.” Consider, for instance, the distinction between causal and non-causal base rates and their effect on the violations of base rate (see Section 19.7.1) that is explicitly recognized in Tversky and Kahneman (1980).
3. The argument has been made that subjects may simply be making errors, could be inattentive, or may be suffering from temporary lapses of judgment. Some of the suggestions that have been made to cast doubt on the KT program include temporary insanity of the subjects, difficult childhood of the subjects, and entrapment by the experimenter (Kahneman, 1981, p. 340). The pattern of biases discovered by KT does not describe random mistakes, but systematic mistakes, hence, these objections cannot be taken seriously. Furthermore, there is typically a great deal of heterogeneity among subjects who exhibit biases, which contradicts conformity with a single rational response. A more detailed discussion can be found in Stanovich and West (2000).
4. A popular criticism that Gilovich et al. (2002) describe as the “*we cannot be that dumb critique*” argues that stellar human achievements, e.g., discovering the structure of DNA

and space flights, are not consistent with the idea that we might be using simple judgement heuristics. This critique arises from a confusion about the nature and process of scientific discoveries. First, scientists might well be using simple heuristics to build initial intuition about their problems. Second, scientific progress is not judged against the standards of individual human judgment heuristics. Rather, it is delegated to an established scientific process that prescribes stringent empirical testing of theories. Hence, there is no essential contradiction between scientific progress and the use of judgment heuristics.

We now consider other criticisms of KT that require either a more detailed or subtle answer.

1. *Frequency versus probability format*: A prominent claim is that the biases identified by KT are eliminated to a large extent when data is presented in frequency format rather than probability format (Gigerenzer et al., 1988, Gigerenzer, 1991, 1994, 1996b). However, the distinction between the frequency and probability formats was first noted by Tversky and Kahneman (1983) themselves in the context of the conjunction fallacy. They found that the frequency format leads to very few conjunction errors; see Section 19.3. However, the conjunction bias is found to be sufficiently high in a between-subjects treatment, relative to a within-subjects treatment, even when data is presented in frequency format (Kahneman and Tversky, 1996).

The anchoring heuristic was also tested by KT in a frequency format; see, e.g., the example on the number of African countries in the United Nations due to Tversky and Kahneman (1973, 1974) in Section 19.6. Indeed, strong evidence of anchoring was found despite the data being presented in terms of natural frequencies.

In the context of base rate neglect, Table 19.4 in Section 19.7.1 shows that while the frequency format reduces base rate neglect, on average about 60% of the subjects in a cross section of studies nevertheless exhibit base rate neglect. Evans et al. (2000) report that the frequency format has been found, depending on the experiments, to worsen, improve, or leave unchanged, the quality of the judgments made.

Thus, the weight of the evidence shows that while the frequency format improves compliance with the rules of classical statistics in some problems (e.g., in the conjunction rule), high non-compliance remains for most problems. Although humans might have evolved to cope better with natural frequencies, casual empiricism suggests that crucial real-world economic data is presented in percentage terms. For instance, interest rates on mortgage borrowing are quoted in percentage terms; the forecasts of inflation and unemployment by central banks are typically made in percentage terms; various insurable risks are often presented to potential insurees in percentage terms; grade information on open days in educational institutions is often given out in percentage terms; and league tables are typically based on inputs that are stated in percentage terms. Thus, it is important to study human judgment and perception in probability tasks.

2. *Formal definitions and statements of heuristics*: Another criticism of KT is that because the heuristics are not stated formally, it is not clear what constitutes a refutation (Gigerenzer, 1991, 1996a; Gigerenzer and Gaissmaier, 2011). For instance, talking of representativeness and availability, Gigerenzer (1991, p. 102) writes that these are: “largely undefined concepts and can be post hoc used to explain almost everything.” Kahneman and Tversky’s (1996) response is that since representativeness can be elicited experimentally (see Section 19.2.1) there is no need to define it a priori.

However, an even better response is that we now have several models that formally define the exact bias in question and derive the formal implications. We have considered several

models in this growing class of models, above. These include: Rabin's (2002) model of inference by believers in the law of small numbers and Rabin and Vayanos's (2010) model of gambler's and hot-hands fallacies (see Section 19.2.3); the model by Biais and Weber (2009) on hindsight bias (see Section 19.8); and the model of Rabin and Schrag (1999) on confirmation bias (see Section 19.9). To this list must be added models of shrouded attributes and limited attention that consider explicit formalizations; see Sections 19.16, 19.17. We anticipate further progress in this direction.

3. *Narrow norms:* Gigerenzer (1996a) questions the appropriate statistical norms that apply to the problems in the experiments in KT; see also Koehler (1996). His argument is that the underlying model is not fully specified, which leaves the appropriate statistical prediction unclear. He writes (p. 592): "A convenient statistical principle, such as the conjunction rule or Bayes' rule, is chosen as normative, and some real-world content is filled in afterward, on the assumption that only structure matters. The content of the problem is not analyzed in building a normative model, nor are the specific assumptions people make about the situation."

This statement should surprise most economists. The analogy in economics is the following. Either one can test the axioms of expected utility theory, e.g., the independence axiom, separately and without reference to an economic model; this often takes the form of observing choice in pairwise lotteries. Or, one can test the predictions of a model that has at its core, the model of expected utility. The first option is scientifically equivalent to the second option because the axioms of rationality are necessary and sufficient for an expected utility representation of preferences. However, the first option is the easier of the two because contrary evidence achieved in the second option is a joint rejection of expected utility and other auxiliary assumptions in the model. Indeed the first option has been massively employed in the testing of expected utility in economics. The airframe of a proposed new airplane might be tested in a wind tunnel, without adding-on all the other components in the airplane. This is a scientific, acceptable and sensible procedure.

Thus, Gigerenzer, in this case, criticizes an appealing feature of the KT tests, i.e., clean unambiguous predictions. Most reasonable economists would be comfortable with the predictions ascribed to the relevant statistical models in the KT tests; although they might contest issues of experimental design such as the appropriateness of subject pools and the appropriate incentivization of the subjects.

The conclusion reached by Kahneman and Tversky (1996, p. 584) is blunt: "The position described by Gigerenzer is indeed easy to refute but it bears little resemblance to ours. It is useful to remember that the refutation of a caricature can be no more than a caricature of a refutation." However, on another count, G&C need to be commended for asking fundamentally important questions about the usefulness of heuristics and making important headway in this direction. KT are likely to agree with this position, although they would be dismayed at how muddled and contentious the debate between the two sides has become.

Both sides also advocate and vigorously debate models of the brain that may lead to biases relative to statistical benchmarks. Kahneman (2011) devotes the first one-third of his book to developing such models of the brain, in this case a model of Systems 1 and 2, that facilitates a deeper understanding of the biases.²⁷ For instance, the quick, reactive, and automatic System 1 is responsible for most errors. System 2 has been likened to a *lazy controller* in Kahneman (2011)

²⁷ However, there is no universal agreement among psychologists about the appropriate models of the brain in this context (Stanovich and West, 2000).

and when it is called upon to intervene in an unusual situation, the agenda (e.g., affective emotions, recalled memory, associations) is chosen by System 1. On p. 21 he writes: “Although System 2 believes itself to be where the action is, the automatic System 1 is the hero of the book.” System 1 tries to make sense of a situation even when the events may have been generated purely randomly. As Stanovich (1999) points out, System 1 has a tendency toward automatic contextualization of problems. This gives rise, for instance, to the law of small numbers, the gambler’s and hot hands fallacies, and overconfidence in the ability to explain the future from a hindsight-biased explanation of past events that were purely random.

There should be no presumption that evolution must have favored rationality, as defined in neoclassical economics. Natural selection chooses among a population phenotype and a mutant phenotype, on the basis of which of the two adapts better to the environment. Indeed traits like loss aversion, prosociality, hyperbolic discounting, and level- k reasoning might be favored in this selection process. There is no presumption that evolutionary mechanisms impart greater rationality to individuals, understood as subjective expected utility maximization, or conformity with the rules of probability theory (Stanovich, 2012).

19.16 Shrouded attributes

Attributes are *shrouded* when some characteristics of the goods that consumers buy are not directly observable. These hidden attributes can typically be discovered by the consumers at some cost, which can be a monetary cost or a purely cognitive search cost. For instance, while opening a bank account, an individual might not know the minimum balance fees or bank surcharges on using its cash machines, or bounced check fees. Neither does the bank make an attempt to easily provide this information, even on its website, although it is tucked somewhere in the fine print, or in some unexpected corner of the website. Manufacturers of printers shroud information on the costs of their patented ink cartridges that can, over the printer’s lifetime, cost several times as much as the cost of the printer.

In classical theory, the costs of *unshrouding* the shrouded attributes are either considered so small that all attributes are assumed to be visible to the consumer, or competition among firms leads to the unshrouding of attributes. Consumers who are rational will believe that shrouded attributes signal bad news and attempt to find other sellers. In response, firms will have an incentive to unshroud information (Jovanovic, 1982; Milgrom, 1981).

19.16.1 *Empirical evidence*

In actual practice, one observes that individual decisions are often visibly influenced by shrouded attributes, even when the cost of unshrouding them appears to be quite small. Consider, for instance, two state taxes on alcohol in the US. The *excise tax* is included in the price of alcohol, while the *sales tax* is not included in the tagged price but it is added at the point of purchase. In classical theory, both taxes should have identical effects on the demand for alcohol. However, the sales tax is shrouded, while the excise tax is unshrouded. Figure 19.11 shows that the elasticity of demand of alcohol with respect to the two taxes is very different. Indeed, as the reader might expect, demand is relatively less responsive to changes in the shrouded tax, relative to the more visible tax.

Consider another example. Hossain and Morgan (2006) conduct a field experiment on eBay auctions in the US in which they put CDs and Xbox games up for sale for different levels of




Figure 19.11 Variations in demand due to a visible excise tax (left panel) and a shrouded sales tax (right panel).

Source: Chetty et al. (2009) with permission from the American Economic Association.

shipping charges and initial bids. The expectation is that shipping charges are shrouded attributes. eBay auctions are an example of ascending bid English auctions and standard auction theory predicts an efficient outcome. Suppose that a bidder has a maximum willingness to pay of $\$x$ for an item. If shipping and handling were free, the bidder should be willing to pay up to $\$x$. Suppose that the minimum opening bid is $\$m$, and shipping and handling costs are $\$s$, and $x \geq m + s$, where $m + s$ is known as the *effective reserve price*. Then, the bidder would be willing to bid $\$b \geq \$m$ for the object, up to $b + s = x$. In particular, so long as these inequalities hold, the buyer should be willing to part with a total of $\$x$ for the object, even as m and s vary.

Two treatments were run. In one treatment, the effective reserve price, $m + s$, was low, and in another it was high, but in each case, $x \geq m + s$. In each treatment, different levels of m and s , keeping fixed $m + s$, were tried out. The predicted revenues of the seller are identical. In contrast to the predicted results, the main empirical finding is that in both treatments, low m and high s led to more bidders, earlier bidding, and higher revenues for both CDs and Xbox games. Thus, it appears that an increase in shipping and handling charges, keeping the effective reserve price constant, increases bidding and seller revenues. This is consistent with the view that many buyers treat shipping and handling charges as shrouded attributes. These results are not unique to eBay auctions in the US and have been replicated in leading online auction platforms in Taiwan and Ireland (Brown et al., 2010).

19.16.2 A theoretical framework

Gabaix and Laibson (2006) develop a general framework to model market competition when some attributes of a good may be shrouded. We have noted above that in the neoclassical case with rational consumers, no shrouding is predicted to take place in equilibrium. However, when not all consumers are rational in the neoclassical sense, we show that shrouding of attributes can take place in equilibrium. Furthermore, firms do not have an incentive, in equilibrium, to educate consumers about the shrouding practices of their rivals.

Consider the following example that nicely illustrates these claims. Suppose that there are two components of the price of a hotel room. The room price or the base price, p_b , and the price of add-ons, p_a , such as costs of parking, telephones, internet connections, and room service. The total price paid by the consumer is $p = p_b + p_a$. The hotel does not keep p_a secret from the public, but makes it sufficiently inaccessible that one needs to expend non-trivial search costs to find it.

The hotel sector is assumed to be perfectly competitive. Suppose that a fraction $\alpha \in [0, 1]$ of the customers is *myopic* (myopes) and the remaining fraction $1 - \alpha$ is *sophisticated*. Myopic consumers cannot observe shrouded information.²⁸ However, myopes exhibit heterogeneity in the presence of unshrouded information. Under unshrouded information, a fraction $\lambda \in (0, 1]$ of the myopes observes the information (informed myopes) and the remaining fraction $1 - \lambda$ (uninformed myopes) pays no attention to unshrouded information about add-ons. Sophisticates comprise a fraction $1 - \alpha$ of the consumers and take account of add-ons when they are unshrouded.

Suppose that one of the hotels is the Hilton and it costs \$100 for the hotel to supply the room. Assume that add-ons have zero costs to the Hilton but consumers typically spend \$20 worth of add-ons when they have to pay for them. In competitive markets, the Hilton will set $p_b = 80$ and $p_a = 20$ and make no profits. Thus, add-ons have the advantage that they subsidize the base price. Consider a competitor hotel, called the Transparent, that supplies free add-ons and also advertises/educates consumers about the add-on prices at its competitors. We assume throughout that education of consumers is costless to the firm. Since the Transparent operates in competitive markets, it must set $p_b = 100$ and $p_a = 0$, thus, the add-ons price cannot subsidize the base price. Let us consider the behavior of each kind of consumer.

1. Uninformed myopes: These consumers find the Hilton base price cheaper. Since they ignore the price of add-ons, they choose the Hilton.
2. Informed myopes and Sophisticates: These consumers benefit from the information that is provided by the Transparent. So they realize that they will spend an amount \$20 on add-ons at the Hilton. They might be able to arrange for these add-ons from the outside. For instance, take a taxi to the hotel to avoid paying parking and take a cell phone to avoid hotel call charges. Suppose that this substitution effort costs them $e = \$10$. These consumers will find it profitable to book a room at the Hilton and pay the base price of \$80 plus the substitution effort cost of \$10 for a total of \$90. This works out cheaper than the price of \$100 for booking a room in the Transparent hotel.

Thus, in equilibrium, the Transparent does not benefit from its policy of advertising about add-ons. By contrast, if the Transparent shrouds add-ons, uninformed and informed myopes may book rooms with it. Unlike the classical framework with purely rational consumers, the presence of myopic consumers allows some attributes to be shrouded in equilibrium.

Consider a formal version of the two-period problem given above in which competitive hotels are indexed, $i = 1, 2, \dots$

In period 1, hotel i offers a base good at a price p_b^i and the price of add-ons is denoted by p_a^i . Myopes comprise a fraction α of consumers, and a fraction λ of myopes are *informed myopes* who take account of add-ons if unshrouded. Unshrouding is costless to the firm and takes the

²⁸ One possibility is that myopic consumers simply do not pay enough attention. Empirical evidence is suggestive of the importance of consumers who exhibit limited attention; see, for instance, DellaVigna and Pollet (2009), Cohen and Frazzini (2008), Hirshleifer et al. (2009), Gilbert et al. (2012).

form of educating consumers about the add-ons. Sophisticates who comprise a fraction $1 - \alpha$ take account of add-ons when unshrouded.

Informed myopes and sophisticates form Bayesian posteriors, Ep_a^i , of add-ons. When add-ons are unshrouded, informed myopes and sophisticates can take account of add-ons, either by (i) choosing a firm that offers cheaper add-ons, or (ii) substitute away from the add-ons at a substitution cost $e : 0 < e < \bar{p}$; $\bar{p}$ can be thought of as the maximum price for add-ons that is allowed by government regulation. Informed myopes, like sophisticates, exert substitution effort only if $e < Ep_a$.

We only consider a symmetric price equilibrium. Denote by v_i , the surplus that a consumer gets by booking at hotel i , net of the opportunity cost of booking at the next best hotel. Variables corresponding to the next best alternative are starred. Thus, for informed myopes and sophisticates, the surplus is²⁹

$$v_i = (-p_b^i - \min\{e, Ep_a^i\}) - (-p_b^* - \min\{e, Ep_a^*\}). \quad (19.77)$$

By contrast, uninformed myopes do not take account of add-ons, so they never exert any substitution effort. Thus, their surplus is given by

$$v_i = (-p_b^i) - (-p_b^*) = p_b^* - p_b^i. \quad (19.78)$$

Let $x(v_i) \in [0, 1]$ denote the probability that a consumer books a room in hotel i ; it captures the demand by the consumer. We assume that $x'(v_i) > 0$, so if a hotel offers better value it is more likely to be chosen.

In period 2, consumers who have booked hotel i can purchase add-ons at a price p_a^i . If consumers have engaged in substitution efforts, then they do not purchase add-ons. We now state the main result.

Proposition 19.5 *Consider a symmetric equilibrium in which the belief of sophisticates and the informed myopes is that if hotel i shrouds add-ons, then $p_a^i = \bar{p}$.*

- (i) *Shrouded prices equilibrium: If $\alpha > \frac{e}{\bar{p}}$, then the price of the base good and the price of add-ons are given, respectively, by*

$$p_b^i = -\alpha\bar{p} + \frac{x(0)}{x'(0)} \text{ and } p_a^i = \bar{p}; i = 1, 2, \dots \quad (19.79)$$

In this equilibrium, only the myopes purchase add-ons.

- (ii) *Unshrouded prices equilibrium: If $\alpha < \frac{e}{\bar{p}}$, then the price of the base good and price of add-ons are given, respectively, by*

$$p_b^i = -e + \frac{x(0)}{x'(0)} \text{ and } p_a^i = e; i = 1, 2, \dots \quad (19.80)$$

All consumers purchase the add-ons in this equilibrium.

Proof: We sketch a part of the proof for (i) and leave the rest as an exercise. We are interested in computing a sequential equilibrium of the game. Thus, we need to check for (a) sequential

²⁹ We can, without loss of generality, include a benefit B of staying in a hotel, but it cancels out in the subsequent calculation (unless there is a hotel-specific benefit of booking a room).

rationality, and (b) consistency of beliefs and actions. If a firm chooses to unshroud (i.e., educate the consumers about the add-ons, so that the informed myopes behave like sophisticates), then the beliefs of the sophisticates and informed myopes are the announced value of p_a^i . Otherwise, if the firm chooses to shroud, then given the stated beliefs, sophisticates believe that $p_a^i = \bar{p}$, i.e., once a consumer is locked-in with a hotel it will be charged the maximum permissible price for the add-ons.

(i) Suppose that $\alpha > \frac{e}{\bar{p}}$ and all firms shroud. We need to check that no profitable deviations from the announced strategies are possible. Given the prices of other firms, p_b^* and p_a^* , if firm i shrouds, then the myopes do not show any awareness of add-ons, and given the stated beliefs, consumers believe that $p_a^i = \bar{p}$ for all firms. Then, from (19.77), (19.78), we get $v_i = p_b^* - p_b^i$ for sophisticates and myopes. Thus, the demand for firm i is given by $x(p_b^* - p_b^i)$ and the fraction α of myopes buy the add-ons. The profits of firm i , therefore, are given by

$$\pi = (p_b^i + \alpha p_a^i)x(p_b^* - p_b^i). \quad (19.81)$$

Firm i chooses p_b^i and p_a^i to maximize π , given in (19.81). The solution to p_a^i is directly obtained as the maximum possible value $\bar{p}$. Setting $p_a^i = \bar{p}$ in (19.81) we get $\pi = (p_b^i + \alpha \bar{p})x(p_b^* - p_b^i)$. Hence, the first order condition with respect to p_b^i is

$$-(p_b^i + \alpha \bar{p})x'(p_b^* - p_b^i) + x(p_b^* - p_b^i) = 0. \quad (19.82)$$

In a symmetric equilibrium, $p_b^i = p_b^*$. Substituting this in (19.82) we get $p_b^i = -\alpha \bar{p} + \frac{x(0)}{x'(0)}$. This is the solution given in (19.79). However, to prove that this is the optimal solution to the problem, a bit more work is needed. We need to consider the case where the firm chooses to unshroud (i.e., educate consumers of the add-ons). Unshrouding educates the sophisticates and the informed myopes who might or might not respond by engaging in a substitution effort, e , depending on whether the price of add-ons is higher or lower. Thus, in this case we need to consider two cases: $p_a^i > e$ and $p_a^i < e$. In particular, the fraction of myopes in this case is given by $\alpha_1 = (1 - \lambda)\alpha$ because the informed myopes (fraction λ of the total myopes) now behave like sophisticates. We leave it to the reader to consider these two cases and show that in each case, the profits of the firm are lower relative to the case of shrouding.

(ii) Proof left to the reader as an exercise. ■

Proposition 19.5 illustrates two different equilibria. In Proposition 19.5(i), firms choose to shroud and charge the maximum permissible price for add-ons, $p_a^i = \bar{p}$ and subsidize the base price relative to the unshrouded equilibrium in Proposition 19.5(ii). Empirical evidence for this prediction is reviewed in Gabaix and Laibson (2006) and Ellison (2005). The intuition for charging the maximum permissible price for add-ons is that shrouding induces the worst possible beliefs about the price of add-ons. Since $e < \bar{p}$, sophisticates engage in substitution effort. Thus, only the myopes buy add-ons, and once they book into the hotel, the hotel might as well charge them the maximum allowed for add-ons.

This is the sense in which the equilibrium in Proposition 19.5(i) is not efficient because only the myopes purchase the add-ons. Furthermore, the sophisticates take advantage of the subsidized base price to improve their surplus under shrouding relative to no shrouding.

Proposition 19.6 *The consumer surplus of the sophisticates is higher under shrouding relative to unshrouding when there is perfect competition among firms.*

Proof: Since perfectly competitive firms face an infinitely elastic demand curve, so $\frac{x(0)}{x'(0)} = 0$. From (19.79), (19.80) the optimal base price is $p_b^i = -\alpha\bar{p}$ under shrouding and $p_b^i = -e$ under unshrouding. Thus, the surplus of sophisticates under shrouding is $-p - e = \alpha\bar{p} - e$; since $e < \alpha\bar{p}$, the surplus is strictly positive. However, under unshrouding, the surplus is $-p - e = e - e = 0$. ■

The fact that sophisticated consumers prefer to engage with firms who shroud and charge high add-on prices but subsidize by offering low base prices is known as the *curse of debiasing*. An example is the low price of printers and the relatively high prices of patented printer cartridges. Sophisticated consumers, in this case, can buy unbranded low-priced printer cartridges on the Internet. In the neoclassical case with no myopes, advertising by firms will eliminate shrouding (Shapiro, 1995). However, the presence of myopes ensures that a captive set of consumers will pay a high price for add-ons in the presence of shrouding. We can see from Proposition 19.5 that if there are no myopes ($\alpha = 0$), then the condition $\alpha > \frac{e}{\bar{p}}$ required for the shrouded equilibrium in Proposition 19.5(i) is never satisfied.

By contrast, in Proposition 19.5(ii), the firm chooses to unshroud and chooses a price for add-ons that just makes consumers indifferent between substitution or not. We assume that the tie breaking rule is that in this case they purchase the add-ons (otherwise the price can be dropped by a penny to ensure that this occurs). This case arises when $\alpha < \frac{e}{\bar{p}}$, so the fraction of myopes is too low to make it worthwhile to extract sufficient consumer surplus by high prices of add-ons; this makes shrouding ineffective.

19.17 Limited attention

In this section, we consider the empirical evidence on *limited attention*. We also outline a general theoretical framework to model this phenomenon.

19.17.1 Limited attention and the poor

Why do individuals take payday loans with associated annual interest rates as high as 400%? Bertrand and Morse (2011) locate the reason in the limited attention of borrowers. They run several treatments in which potential borrowers are provided with a range of information on the real costs of payday loans. The information includes the weekly accumulation of fees (e.g., \$270 in 3 months on a \$300 loan), comparison of borrowing costs from other financial institutions, and refinancing using their credit card borrowing. These interventions succeeded in reducing the extent of borrowing; in the four treatments, the percentage reduction in borrowing was 5.9%, 12%, 16%, and 23%, respectively. These results suggest a potential role for legislation in making the provision of such information mandatory. However, the relatively small reductions in borrowing also indicates that limited attention alone does not account for the take-up of payday loans.

Stango and Zinman (2015) find that drawing attention to fees for overdrawn accounts, reduces overdraft fees for up to 2 years since the first reminders are sent out. The effects are more pronounced for individuals who have lower education and financial literacy. Limited attention might be expected to play a greater role when the number of options is greater. In an empirical study of the take-up of loans, Bertrand et al. (2010) find that reducing the number of options from four to one, increased loan take-ups; the effect is equivalent to a 2.3% reduction in the interest rate.

In Mexico City, only 39% of low-income individuals could spot the lowest cost credit product in bank brochures. However, 68% could spot the lowest cost product in a user-friendly summary sheet (Giné et al., 2014); limited attention and shrouded attributes are possible explanations. These examples also highlight the need for financial market regulation for framing information in a more transparent manner in the presence of limited information.

Limited attention is likely to reflect the cognitive costs of engaging in greater attention. These cognitive costs, which might take the form of depletion of scarce willpower required to exert greater attention, may create undesirable outcomes, especially for the poor. The poor may need to expend a disproportionate amount of scarce attention and willpower costs on important daily issues. These issues include arranging food for the table, dealing with debt, or arranging for clean drinking water; the associated willpower costs may be termed as a form of *cognitive tax* (Shah et al., 2012; Mullainathan and Shafir, 2013).

The result of this depletion in willpower is that the poor are more likely to engage their System 1 in making other potentially important decisions; the control exerted by System 2 on System 1 falls as willpower depletes. Thus, the poor may engage in insufficient deliberation on many important issues (Banerjee and Mullainathan, 2008); this is sometimes referred to as the problem of *low bandwidth*. Hence, poor individuals may make bad decisions, which perpetuate poverty that is the cause of the initial problem.

Indeed, the actions of the poor may not fully reflect their intentions; this is known as the *intention-action divide*. In the case of the poor, this problem is mostly caused by (1) intense willpower depleting focus on pressing current problems, and (2) present-biased preferences. Mani et al. (2013) find that farmers score much lower on cognitive ability tests before a harvest when they have high levels of financial stress. Measured IQ is 10 points lower before the harvest, relative to the period after the harvest.

Karlan et al. (2012) find that text messages sent out to individuals in Peru, Bolivia, and the Philippines, reminding them of their savings goals, actually improve savings by 16%. Presumably, these messages work because of issues of limited attention and procrastination (see Part 3 of the book). Furthermore, the most effective messages were the ones that emphasized a particular goal, for instance, purchase of a house or an appliance. Limited attention might also explain why many sophisticated financial literacy programs are not successful. Drexler et al. (2014) found that reducing the degree of sophistication of the financial programs, by emphasizing simple heuristics, improves financial literacy of firms and micro-entrepreneurs in the Dominican Republic.

The solutions to the problems of cognitive tax and low bandwidth can sometimes be surprisingly simple and offer a low-cost means of bringing about large welfare changes for the poor; Datta and Mullainathan (2014) term this as *low hanging policy fruit*. This suggests that poverty is not simply a problem of insufficient resources, so this offers a radically new perspective on the problem (WDR, 2015). Development policy could then play a key role in the timing of important decisions. For instance, sugarcane farmers have lower consumption in the months prior to the harvest and they may even pawn jewelry. However, if part of their harvest money could be put into an account with pre-specified withdrawal rates, then they may better smooth their consumption levels. Such default options have been shown to have important effects in the West for pensions and savings.

In the presence of limited attention and limited cognitive abilities, a benevolent or cunning third party might wish to frame the context within which individuals make decisions. This is known as the problem of *choice architecture*. Framing of context can take many forms; for instance, which information to make salient, which comparisons among options to highlight, which anchors to choose and so on. One example of choice architecture is Thaler and Sunstein's

(2008) idea of a *nudge*. For instance, on account of limited attention, if the healthiest snacks are put at eye level, they are more likely to be chosen; these issues are considered in more detail in Part 8.

19.17.2 Limited attention and taxes

Does consumer demand respond more to easily observed taxes relative to less salient taxes? For instance, sales taxes may be (i) posted on price stickers in shops, or (ii) not posted on the price sticker but added at the sales register at the point of sale. If demand or its responsiveness turns out to be different in the two cases (see Figure 19.11), then two competing explanations arise.

First, consumers may be unaware of the sales tax. Second, hidden fiscal attributes are less salient. Chetty et al. (2009) surveyed consumers at a grocery store where the sales tax was displayed for some but not other consumer goods. They found that the median consumer was able to successfully identify the sales tax on seven out of eight products on the survey. This finding, and the fact that consumers could easily inquire about the level of sales tax from the shop attendant if they wished, suggests that taxes that are not displayed on price stickers reduce tax salience.

Consider a simple model of consumer demand. Suppose that there are two consumer goods, x and y . Good y is an untaxed good whose price is normalized to 1. The pretax price of good x is p and the post-tax price is $q = (1 + \tau)p$, where τ is an ad-valorem sales tax. Consumers observe the pre-tax price on the posted price stickers in shops, but the sales tax, τp , is not included in the posted price.

Denote the solution to the consumer's optimal demand for x by $x(p, \tau)$. In traditional public finance theory, the consumer takes account of the post tax price, q , even if the sales tax is not posted on the price sticker. Thus, we have $x(p, \tau) = x((1 + \tau)p, 0)$, so there is no need to separately observe the sales tax if the consumer already takes account of the post tax price. In this formulation, an increase in the post tax price, q , whether it is caused by an increase in τ or an increase in p , has identical effects on the demand for x .

Let $\epsilon_{x,p}$ and $\epsilon_{x,1+\tau}$ denote, respectively, the elasticity of demand for x with respect to p and $1 + \tau$. Under classical public finance theory we should, by definition, observe that

$$\epsilon_{x,p} \equiv -\frac{\partial \log x}{\partial \log p} = \epsilon_{x,1+\tau} \equiv -\frac{\partial \log x}{\partial \log(1 + \tau)}. \quad (19.83)$$

In contrast to (19.83), Figure 19.11 shows that $\epsilon_{x,p} > \epsilon_{x,1+\tau}$. One way of representing limited attention to taxes that are not displayed on price stickers is to write $x(p, \tau)$ in a log-linear form.

$$\log x(p, \tau) = \alpha + \beta \log p + \theta \beta \log(1 + \tau). \quad (19.84)$$

In (19.84), the coefficients β and $\theta\beta$, respectively, give the two elasticities defined in (19.83). Thus,

$$\theta = \frac{\frac{\partial \log x}{\partial \log(1+\tau)}}{\frac{\partial \log x}{\partial \log p}} = \frac{\epsilon_{x,1+\tau}}{\epsilon_{x,p}}. \quad (19.85)$$

If $\theta = 1$, then we have the classical case in (19.83). However if $\theta < 1$, then we get the observed result in Figure 19.11. If one could uncover independent variation in p and τ , then $\epsilon_{x,p}$ and $\epsilon_{x,1+\tau}$ can be estimated separately to compute θ .

Chetty et al. (2009) conducted an experiment in a supermarket for three weeks in 2006. Initially, all products in the store excluded the sales tax on the price sticker; the sales tax was added later at the point of purchase. The price tags on a selected list of 750 products were altered to separately display the pre-tax price and the sales tax of 7.375% in order to increase the salience of the sales tax. The products included three groups: cosmetics, deodorants, and hair care accessories. The posting of the tax on the price stickers reduced demand by 8% relative to two control groups: other products in the store and similar products in two nearby stores.

The reduction in demand was statistically significant. One can reject the null hypothesis that posting of the sales tax on price stickers has no effect by using a t-test and non-parametric permutation tests. The effect of tagging products with the sales tax caused a fall in demand that was identical to a 7.373% increase in price. Hence, consumers pay limited attention to taxes that are not displayed on price stickers. θ is estimated to be 0.35. A separate survey conducted on supermarket shoppers showed that the median shopper was aware of the level of sales tax. Hence, the observed effect on demand was not the result of dispelling of ignorance, but an increase in the salience of the sales tax.

In a second study, the authors consider separately the effect of different kinds of indirect taxes on alcohol consumption. They considered the effect of excise duties that are included in the alcohol price and sales taxes that are only included at the point of sale, hence, these are shrouded; see Figure 19.11. The main empirical result is that, based on data over the period 1970–2003, excise taxes (more visible) reduced alcohol consumption far more than sales taxes (shrouded).

Goldin and Homonoff (2013) find that attention to sales taxes levied at the sales register differs among low-income consumers (who pay greater attention) relative to high-income consumers. This can be exploited by the fiscal authorities to reduce the degree of perceived tax regressivity by taxing goods consumed by the rich, relatively more.

Results of a similar nature have been obtained elsewhere. For instance, the earned income tax credit (EITC) is the largest cash transfer program in the US that transfers income to low-income families. However, evidence suggests that the marginal incentives implied by the program, e.g., should one work an extra hour, are not salient, hence, they would be ignored by individuals with limited attention. The EITC program appears to have had a relatively small impact on the intensive margin, i.e., hours worked, relative to the extensive margin, i.e., the decision to participate or not in the labor force (Hotz and Scholz, 2003).

Chetty and Saez (2013) arrange for tax experts to provide information on the intensive margin for potential EITC beneficiaries. Information provision takes several forms: a verbal description; a graph that reveals the behavior of the EITC as earnings change; and a table that lists key information on the EITC. Furthermore, potential beneficiaries were given advice that was particularly tailored to their personal circumstances. For instance, they were told if it would pay for them to work more or not. Yet, despite these attempts, the average behavior of the potential beneficiaries along the intensive margin did not change. One potential reason why providing information does not work in this case is that this problem is relatively more cognitively challenging relative to simply adding a sales tax on the price tag. Another reason is that the intensive margin may be determined by factors entirely outside this model.

Buchanan (1967) and Brennan and Buchanan (1980) note the potential link between tax salience and the size of the government. The basic idea is that governments can expand by increasing taxes that are not salient, ruling out adverse electoral consequences for political parties.

Finkelstein (2009) confirms these findings by examining the effect of electronic toll collections (ETC) in car plazas that automatically deduct the relevant toll; there is no manual exchange of cash between two parties in this case. Comparing customers who pay by manual tolls and by ETC, she finds much less awareness of the level of the toll among the ETC customers.

She also finds that there is an increase in the toll rates (the analog of an increase in the size of the government) following the introduction of an ETC. Indeed, toll rates are 20–40% higher in the presence of ETC as compared to manual toll collections. Furthermore, the timing of ETC increases is less sensitive to the electoral calendar relative to manual tolls, suggesting that politicians are less concerned about its electoral implications, most likely because people pay limited attention to these changes.

19.17.3 *Biased numerical attention*

It has been noted that people might unduly focus on the leftmost digit of a multi-digit number, relative to the digits on the right (Korvost and Damian, 2008). This appears to be a form of imperfect attention to some digits over others; we may term it as the *left-digit bias*. Basu (1997, 2006) explained the prevalence of prices ending in the digit 9; for instance, a price of $\$x.99$ (where x is a whole number), rather than an integer price of $\$x + 1$; e.g., a price of $\$39.99$ rather than $\$40$. He shows that in a rational expectations equilibrium, all firms choose prices ending in 99 cents and consumers, rationally, do not benefit by observing the full price (say, $\$3.00$, when the quoted price is $\$2.99$). This explanation is quite distinct from the one that ascribes such biases to limited attention arising from cognitive limitations.

Lacetera et al. (2012) explore the presence of the left-digit bias in the wholesale used car market. Potential buyers may exhibit the left-digit bias when presented with information on the odometer readings of cars. In particular, they find that there are discontinuous drops in the retail prices of cars at 10,000 mile thresholds. For instance, cars with odometer readings 79,900–79,999 miles sell for (i) $\$210$ more than those with odometer readings in the range, 80,000–80,100, yet only (ii) $\$10$ less than those with odometer readings 79,800–79,899. There are also discontinuities at the 1000 mile thresholds but these are smaller and we shall ignore them here. Furthermore, there are spikes in the volume of cars presented for sale just before they hit these odometer thresholds, suggesting awareness of this bias among car sellers. These results are also shown to hold for the retail used car market (Busse et al., 2013).

Buying a car is a large stakes decision and odometer readings are easily observed. Hence, this evidence is an important demonstration of the presence of heuristics in real markets. As with all heuristics, individuals commit systematic biases relative to the benchmark of classical statistics. The left-digit bias may be a useful heuristic in the fast-and-frugal sense, and this needs to be explored in future work. For instance, buying a car just below the threshold is costly because it will soon cross the threshold but the benefit might arise if this information is generally correlated with other useful characteristics of the car.

Consider a formal model of left-digit bias that illustrates the limited attention aspect of the bias. Suppose that the reading on the odometer is a k digit number, denoted by $m = d_1 d_2 \dots d_k$, where d_1 is the leftmost digit and d_k , the rightmost digit; each digit belongs to set $\{0, 1, \dots, 9\}$. Suppose that m , the actual odometer reading, is a positive whole number less than or equal to 100,000. A potential buyer perceives the odometer reading as $\hat{m}$, where³⁰

³⁰ A slightly different formula applies when $m > 100,000$ because the changes in this case are not to the leftmost digit until m exceeds 200,000. The reader may wish to think about the necessary modifications in the formula.

$$\hat{m} = d_1 10^{k-1} + (1 - \theta) \sum_{j=2}^k d_j 10^{k-j}, \quad (19.86)$$

where $\theta \in [0, 1]$ is the *inattention parameter*; thus the weight put on the leftmost digit, 1, is greater than the weight that is put on any of the other digits, $1 - \theta < 1$. The classical case corresponds to full attention, $\theta = 0$. An individual is maximally inattentive if $\theta = 1$, in which case, only the leftmost digit is salient. For instance, for $m = 4326$, (19.86) implies that $\hat{m} = 4 \times 10^3 + (1 - \theta) (3 \times 10^2 + 2 \times 10^1 + 6 \times 10^0)$. Whenever $0 < \theta \leq 1$ we have $\hat{m} < m$, i.e., the perceived mileage is lower than the actual mileage.

The actual value of the car is $V = V(m)$. However, the perceived value of the car is $\hat{V} = V(\hat{m})$, which takes a linear form

$$\hat{V} = V(\hat{m}) = v - \alpha \hat{m}, \quad (19.87)$$

where $v \geq 0$ is the component of valuation that is independent of the mileage, and $\alpha \geq 0$ is a measure of how quickly the car is perceived to depreciate. As consumers become maximally inattentive, $\theta = 1$, $V(\hat{m})$ becomes a step function.

Consider the change in perceived value for any two odometer readings m_1 and m_2 such that $m_1 - m_2 = 1$. Two cases arise.

- (i) m_1, m_2 do not straddle two different thresholds, e.g., $m_1 = 43,246$ and $m_2 = 43,245$.
- (ii) m_1, m_2 straddle different thresholds, e.g., $m_1 = 50,000$ and $m_2 = 49,999$.

For these two cases, (i) and (ii), a simple calculation, using (19.86), shows that

$$\hat{m}_1 - \hat{m}_2 = \begin{cases} 1 - \theta & \text{Case (i)} \\ 1 + \theta \times 9999 & \text{Case (ii)} \end{cases}. \quad (19.88)$$

Using (19.87) and (19.88), we get $V(\hat{m}_1) - V(\hat{m}_2) = -\alpha (\hat{m}_1 - \hat{m}_2)$, thus,

$$V(\hat{m}_1) - V(\hat{m}_2) = \begin{cases} -\alpha(1 - \theta) & \text{Case (i)} \\ -\alpha(1 + \theta \times 9999) & \text{Case (ii)} \end{cases}. \quad (19.89)$$

Clearly, there is a discontinuous drop in the perceived value of the car as we move from Case (i) to Case (ii). Consider Case (ii). From the second row in (19.89), we can write $V(\hat{m}_1) - V(\hat{m}_2) = -\alpha [(m_1 - m_2) + \theta \times (10000 - (m_1 - m_2))]$. Thus,

$$\lim_{m_2 \rightarrow m_1} [V(\hat{m}_1) - V(\hat{m}_2)] = -\alpha \theta \times 10000,$$

gives the discontinuous drop in car value at a threshold.

These predictions are nicely summarized in the self-explanatory Figure 19.12 for $50,000 \leq m \leq 100,000$. Since we have assumed a constant rate of depreciation, embodied in a constant value of α , the discontinuous downward jumps at each threshold are of the same size. However, comparing cars of different makes or models that have different rates of depreciation, more rapidly depreciating cars (higher α) will have larger discontinuous jumps at the thresholds.

The data for testing these predictions comes from wholesale markets in used cars. A simple theoretical model can be sketched that the reader might wish to formalize further. Suppose that




Figure 19.12 Predictions of a model of limited attention on used car values.

Source: Lacetera et al. (2012) with permission from the American Economic Association.

there are N consumers with unit demands for cars and whose perceived value for a car is given in (19.87). Cars are sold in a competitive retail market by dealers who, in turn, purchase their cars from auctions. The auctions are either second price auctions or ascending-bid English auctions. There are $M \leq N$ cars in the auctions; each car displays an odometer mileage, and its reserve price is zero. Since the market for dealers is competitive, they make no profits. Hence, in the auction market, they bid as if they had the valuations of the buyers in (19.87). Given these assumptions, the auction produces an efficient outcome, so for a car of any mileage, m , the auction price equals $V(\hat{m})$.

Figure 19.13 shows the raw prices of cars based on data for 22 million cars sold during the sample period, 2002–2008, at one of the largest wholesale car price auctions in the US. In the figure, the mileage is rounded to the nearest 500 (so, e.g., 42,315 is rounded to 42,500). The vertical lines in the figure represent each of the 10,000 mile thresholds from 0 to 125,000 miles. Average prices decline with mileage. Within two consecutive threshold values, the decrease in price is continuous. However, at each threshold value, there is a discontinuous downward jump in price, as predicted by models of limited attention.

Repeating the same analysis by netting out the effects of various fixed effects, such as the model, make, year, body style of cars, and plotting the price residuals gives a result very similar to Figure 19.13; the price plot is even smoother in this case but the price discontinuities survive. There are spikes in the volume of cars brought to the auctions with odometer readings just short of the 10,000 mile thresholds; this is most noticeable at thresholds starting at 60,000 miles. This suggests that car sellers are aware of the problem of limited attention on the part of car buyers. Thus, the left-digit bias is a feature of the demand side of the market rather than the supply side.

Using data from German used car markets, Englmaier et al. (2014) replicate the results above about the fall in the value of used cars when the odometer hits critical threshold values. However, they also find support for an additional and plausible factor in explaining discontinuous drops in used car values, the vintage or year of registration of the car. The left-digit bias cannot be used




Figure 19.13 Raw prices of cars based on data for 22 million cars sold during the sample period, 2002–2008.
Source: Lacetera et al. (2012) with permission from the American Economic Association.

to account for the vintage effect, so they propose a slightly different theoretical model to account for both effects

19.17.4 A theoretical framework

In neoclassical economics, an individual may be expected to take account of a very large number of variables in solving a problem. However, cognitive limitations may allow limited attention to some factors and greater attention to more salient factors; variables given limited attention might be substituted by default values. The choice of which factors to treat as salient may be an endogenous one. Once the individual decides on the set of salient factors, he proceeds as in neoclassical models in economics. Models with these features, where the decision is made in a manner that is consistent with neoclassical models of rationality, are sometimes known as *rational inattention models*. Our main focus in this section is on the *sparse max operator* model of Gabaix (2014) that we consider to be psychologically better grounded, and more tractable. We offer some brief comments on other models at the end.

Consider the classical problem of maximizing utility $\max_{x \in X} u(x, \theta)$, where $X \subset R^k$ is the k dimensional real domain of choices under consideration, and θ is a vector of parameters drawn from some set of parameters $\Theta \subset R^n$. We wish to adapt this problem to reflect limited attention by defining a *sparse max* operation,

$$s \max_{x \in X} u(x, \theta). \quad (19.90)$$

Denote $\theta = 0$ as the *default parameter* and the corresponding optimal action $x^d = \max_{x \in X} u(x, 0)$ as the *default action*. We assume that u is concave in x and twice continuously differentiable around $(x^d, 0)$. Suppose that the parameter θ is a linear combination of n other parameters $\theta_1, \theta_2, \dots, \theta_n$.

$$\theta = \sum_{i=1}^n \mu_i \theta_i, \text{ and } \forall i, \mu_i \in R, \theta_i \in \Theta.$$

Limited attention is captured as follows. The perceived value, $\hat{\theta}_i$, of the parameter θ_i is given by

$$\hat{\theta}_i = a_i \theta_i, i = 1, \dots, n, \quad (19.91)$$

where $a_i \in [0, 1]$ is the attention given to θ_i and $a = (a_1, a_2, \dots, a_n)^T \in [0, 1]^n$ is known as the *attention vector*; superscript T denotes the transpose. $a_i = 1$ corresponds to the case of full attention in neoclassical economics. At the other extreme, when $a_i = 0$, the parameter θ_i is given no attention. Once the optimal attention vector $a^* = (a_1^*, a_2^*, \dots, a_n^*)^T$ is endogenously determined, we can calculate the perceived or sparse parameter

$$\hat{\theta}(a^*) = \sum_{i=1}^n \mu_i \hat{\theta}_i = \sum_{i=1}^n \mu_i a_i^* \theta_i.$$

The individual then solves the following problem

$$x^* \in \arg \max_{x \in X} u(x, \hat{\theta}(a^*)). \quad (19.92)$$

The first order condition is $u_x(x^*(\hat{\theta}(a^*)), \hat{\theta}(a^*)) = 0$, where $x^*(\hat{\theta}(a^*))$ is the optimal solution. Substituting the optimal solution, $x^*(\hat{\theta}(a^*))$, in the direct utility function, u , we get the indirect utility function

$$v(a) = u(x^*(\hat{\theta}(a^*)), \theta). \quad (19.93)$$

Using the implicit function theorem on the first order condition and evaluating derivatives at $(x, \theta) = (x^d, 0)$, we get

$$x_{\theta_i} \equiv \frac{\partial x}{\partial \theta_i} \big|_{(x^d, 0)} = - \frac{u_{x\theta_i}}{u_{xx}} \big|_{(x^d, 0)}. \quad (19.94)$$

Thus, x_{θ_i} is the change in the optimal action when θ_i increases by one unit at the default option.

Consider, henceforth, the special case of quadratic utility and $x \in \mathbb{R}$.

$$u(x, \hat{\theta}) = \frac{-1}{2} \left(x - \sum_{i=1}^n \mu_i a_i^* \theta_i \right)^2. \quad (19.95)$$

In this case, the optimal solution is

$$x^* = \sum_{i=1}^n \mu_i a_i^* \theta_i. \quad (19.96)$$

We now model the optimal determination of attention. Suppose that attention is cognitively costly and the cost function of attention to θ_i is given by

$$C(a_i) = ca_i^\alpha, c > 0, \alpha \geq 0. \quad (19.97)$$

When $\alpha = 0$, there is a fixed cost of attention equal to c . In classical game theory, attention is costless, so $c = 0$. Suppose that the vector θ is drawn from some distribution $g(\theta)$ such that the means of the marginal distributions are zero ($E\theta_i = 0$), the standard deviations are $\sigma_i = E[\theta_i^2]^{1/2}$, and the covariances are given by $\sigma_{ij} = E[\theta_i\theta_j]$. The total cost of attention is given by

$$C(a) = \sum_i C(a_i).$$

The individual wishes to choose the attention vector, a , in order to maximize the expected value of indirect utility in (19.93), net of the cost of effort given in (19.97). Thus, the individual's problem is

$$\max_{a \in [0,1]^n} Ev(a) - C(a).$$

To simplify the problem we take a Taylor expansion of $v(a)$ around $a \equiv i = (1, 1, \dots, 1)^T$, where i is a vector of ones.

$$v(a) \approx v(i) + Dv(i)(a - i) + \frac{1}{2}(a - i)^T D^2v(i)(a - i), \quad (19.98)$$

where $Dv(i)$ is the matrix of first order partial derivatives of v and $D^2v(i)$ is the matrix of second order partial derivatives. $(a - i)$ is a row vector with the j th element $a_j - 1$. Using the definition of indirect utility in (19.93), $v(a) = u(x^*(\hat{\theta}(a)), \theta)$, and (19.91), differentiating once with respect to a_i we get $v_{a_i} = u_x x_{\hat{\theta}_i} \theta_i$. Since $E\theta_i = 0$ we have $EDv(i)(a - i) = 0$ (this is the zero vector). Differentiating v_{a_i} with respect to a_j we get $v_{a_i a_j} = x_{\hat{\theta}_i} u_{xx} x_{\hat{\theta}_j} \theta_j$ (recall that utility is quadratic, (19.95), and the optimal solution is (19.96), so the cross partials are zero). Taking expected values we get $Ev_{a_i a_j} = E[\theta_i \theta_j] x_{\hat{\theta}_i} u_{xx} x_{\hat{\theta}_j} = \sigma_{ij} x_{\hat{\theta}_i} u_{xx} x_{\hat{\theta}_j}$. Using these results, and taking the expected value in (19.98) we get

$$E[v(a) - v(i)] \approx \frac{-1}{2}(a - i)^T \Lambda (a - i),$$

where Λ is the *cost of inattention matrix* such that

$$\Lambda_{ij} = -\sigma_{ij} x_{\hat{\theta}_i} u_{xx} x_{\hat{\theta}_j}. \quad (19.99)$$

Thus, an individual who uses the sparse max operator will choose the vector of attention a to

$$\text{maximize}_{a \in [0,1]^n} \frac{-1}{2}(a - i)^T \Lambda (a - i) - C(a). \quad (19.100)$$

We summarize this in a formal definition.

Definition 19.6 (*Unconstrained sparse max operator*): The individual undertakes the following two-step optimization exercise.

- (1) Choose optimal attention, a^* , to solve (19.100). This gives the sparse representation of θ_i , $\hat{\theta}_i = a_i^* \theta_i$, $i = 1, \dots, n$.
- (2) The individual then solves the optimization problem in (19.92) giving rise to the vector of optimal choices, x^* .

Consider a simple one-dimensional example, $k = 1$, $n = 1$, quadratic utility, and $c = 1$. In this case, the analog of (19.100) is

$$\max_{a \in [0,1]} \frac{-1}{2} (a - 1)^2 \sigma^2 - a^\alpha. \quad (19.101)$$

The optimal values of a for three different values of $\alpha = 0, 1, 2$, denoted, respectively, by a_0^* , a_1^* , and a_2^* , are given by

$$a_0^* = \begin{cases} 0 & \text{if } \sigma^2 < 2 \\ 1 & \text{if } \sigma^2 \geq 2 \end{cases}.$$

$$a_1^* = \begin{cases} 0 & \text{if } \sigma^2 < 1 \\ 1 - \frac{1}{\sigma^2} & \text{if } \sigma^2 \geq 1 \end{cases}.$$

$$a_2^* = \frac{\sigma^2}{2 + \sigma^2}.$$

A plot of the optimal values against σ^2 (in bold) is shown in Figure 19.14; a_0^* (left panel), a_1^* (middle panel), and a_2^* (right panel). It is obvious that a_0^* , a_1^* induce sparsity. In both cases, there is a range of values for which the optimal attention is 0; a_1^* is continuous but a_0^* is not continuous at $\sigma^2 = 2$. In these cases, the individual is rationally inattentive in the sense that he takes account of the costs of attention as well as the stochastic structure of the problem. However, a_2^* never induces sparsity. We leave the proof of the next proposition for the reader.

Proposition 19.7 (Gabaix, 2014): Suppose that $\sigma_{ij} = 0$ for all $i \neq j$. Then the optimal attention, a_i^* , is increasing in σ_i^2 , x_{θ_i} , and decreasing in c .

Optimal attention is increasing in σ_i^2 and x_{θ_i} because these factors increase the loss in utility from imperfect attention (see (19.99)). An increase in c increases the cost of paying attention (see (19.100)), hence, it reduces optimal attention.

Gabaix (2014) offers plausible psychological interpretations of his model. For instance, Step 1 in Definition 19.6 may be considered to be a System 1 operation that decides how much attention to




Figure 19.14 Determination of the optimal attention parameters as a function of the variance of the underlying parameter. The horizontal axis in each of the three panels is σ^2 . The vertical axes in the three panels are: a_0^* (left panel), a_1^* (middle panel), and a_2^* (right panel).

give to a problem. Once this problem is solved, then Step 2 appears akin to a System 2 problem. The case of no attention, $a_i^* = 0$, may be considered to arise when System 1 operates under intense time pressure, and so assigns the default option $(x^d, 0)$. One may also consider this as an anchoring and adjustment model in which one begins at the default option $(x^d, 0)$ and then slowly adjusts towards the optimum (x^*, a^*) although this often falls short of the full attention optimum $(x, 1)$.

Once this is accepted, the model can be extended to constrained optimization. Suppose that now the problem is to optimize, as in (19.90), subject to the constraint $f(x, \theta) \geq 0$. Define the Lagrangian

$$L = u(x, \theta) + \lambda^d f(x, \theta), \quad (19.102)$$

where λ^d is the Lagrangian multiplier.

Definition 19.7 (*Constrained sparse max operator; Gabaix, 2014*): In the constrained problem, the sparse max operation, $s \max_{x \in X} u(x, \theta)$, subject to $f(x, \theta) \geq 0$, is defined as follows. The individual undertakes the following two-step optimization exercise.

- (1) Choose optimal attention, a^* , to solve (19.100), but instead of (19.99), using $\Lambda_{ij} = -\sigma_{ij} x_{\hat{\theta}_i} L_{xx} x_{\hat{\theta}_j}$ and $x_{\hat{\theta}_j} = -\frac{L_{x\theta_j}}{L_{xx}} |_{(x^d, 0)}$ where x^d solves $\arg \max_x L(x, 0)$. This gives the sparse representation of θ , $\hat{\theta}_i = a_i^* \theta_i$, $i = 1, \dots, n$.
- (2) Define the function $x(\lambda) = \arg \max_x [u(x, \hat{\theta}) + \lambda f(x, \hat{\theta})]$. Now solve the problem

$$\lambda^* \in \arg \max_{\lambda} u(x(\lambda), \hat{\theta})$$

subject to

$$f(x(\lambda), \hat{\theta}) \geq 0.$$

The solution to this optimization problem gives rise to the vector of optimal choices, $x^*(\lambda^*)$.

Gabaix (2014) uses this framework to re-derive basic results in consumer theory in microeconomics and in competitive equilibrium. Some of these results include the following. Individuals optimally exhibit money illusion. This can arise because consumers are inattentive to an increase in price but attentive to an increase in their incomes. The Slutsky matrix may not be symmetric under limited attention, and a behavioral analog of the Edgeworth box is derived. The weak axiom of revealed preference may also not hold.

Overall, this is an important contribution that is better psychologically founded than many of the other theoretical contributions. But, in the spirit of the approach in the rest of this book, it must await stringent testing before it gains greater acceptance. For that it must address questions such as the following. Is imposing a stochastic structure on the unknown parameters of the utility function the most attractive approach to modeling limited attention? What are the experimental/field analogs of σ_i^2 , x_{θ_i} , and c ? In several of the examples that we considered above, such as a fall in demand when sales taxes (known to people) are made salient, or changes in the pattern of Internet shopping when the less-salient delivery charges (known to people) are varied, the observed environments appear to be non-stochastic. How should we model limited attention in these cases? One may expect much progress to be made in future years on these questions.

In a popular class of models of rational inattention, information is modeled as a reduction in uncertainty. Uncertainty, in turn, is modeled as a reduction in *entropy*. Finally, an upper limit on the flow of information is imposed (Sims, 1998, 2003, 2006). Consider the simple exposition of the basics of this approach in Wiederholt (2010). Suppose that a random variable Y is distributed normally with mean μ_Y and variance σ_Y^2 . The *unconditional entropy* of Y is given by

$$S(Y) = \frac{1}{2} \log_2(2\pi e \sigma_Y^2). \quad (19.103)$$

The *conditional entropy* of Y , conditional on another normally distributed variable Z , is given by

$$S(Y | Z) = \frac{1}{2} \log_2(2\pi e \sigma_{Y|Z}^2), \quad (19.104)$$

where $\sigma_{Y|Z}^2$ is the conditional variance. A measure of the information that the random variable Z contains about Y is given by

$$\Delta S = S(Y) - S(Y | Z).$$

Observing Z reduces uncertainty about the variable Y . In this literature, limited attention is modeled as an upper bound, $b > 0$, on the information flow

$$\Delta S \leq b. \quad (19.105)$$

Using (19.103), (19.104) in (19.105) we get $\log_2 \left(\frac{2\pi e \sigma_Y^2}{2\pi e \sigma_{Y|Z}^2} \right) \leq 2b$, or

$$\frac{\sigma_Y^2}{\sigma_{Y|Z}^2} \leq 2^{2b}. \quad (19.106)$$

We may think of Z as a noisy signal of Y so that $Z = Y + \varepsilon$ where $\varepsilon \sim N(0, \sigma_\varepsilon^2)$. In this case, (19.106) can be written as

$$\frac{\sigma_Y^2}{\sigma_\varepsilon^2} \leq 2^{2b} - 1. \quad (19.107)$$

Thus, in this literature, limited attention imposes bounds on the signal to noise ratio in a purely statistical manner. These ideas have been used successfully in physical communication channels such as characterizing the rate of information flow in modems and Internet connections. It is an open question if these statistical techniques also describe similar information flows for humans. In rational inattention models, one minimizes two kinds of costs: departures of decisions from an optimum and the cost of paying attention, subject to a constraint such as (19.107). One criticism of the earlier models in this tradition is that they cannot model source-dependent inattention but recent work has attempted to address this criticism (Maćkowiak and Wiederholt, 2009; Woodford, 2012). Gabaix (2014, p. 1696) argues that these models still lack the tractability and generality of the sparse max operator.

19.18 Do experts exhibit biases?

Does market experience eliminate behavioral anomalies? The evidence gives a strongly negative answer, although experience may lessen some anomalies. In this section, we look at expert-bias in the context of judgment heuristics.

Existing empirical work suggests that experts are not immune to biases. Tetlock (2002) looks at the prediction of experts of political events. His results are sufficiently striking to justify a detailed quote.

Across all seven domains, experts were only slightly more accurate than one would expect from chance Most experts thought they knew more than they did. Across all predictions that were elicited, experts who assigned confidence estimates of 80% or higher were correct only 45% of the time . . . Expertise thus may not translate into predictive accuracy but it does translate into the ability to generate explanations for predictions that experts themselves find so compelling that the result is massive over-confidence.

Most evidence shows that experts are typically more overconfident relative to lay people who lack experience (Heath and Tversky, 1991; Frascara, 1999; Glaser et al., 2007a,b) and that overconfidence increases with experience (Kirchler and Maciejovsky, 2002). The predicted distributions of stock market returns by senior finance professionals, in a large dataset, are too narrow, and realized returns are within their 80% confidence intervals, only 36% of the time (Ben-David et al., 2013).

In this chapter, we have reported many empirical findings that demonstrate several kinds of expert-bias that we quickly summarize now. For surveys, see, for instance, Gilovich et al. (2002), Kahneman (2003), Kahneman (2011), and Tetlock (2006). Tetlock (2010) is a useful response to alternative critical views. Mathematical psychologists exhibit the law of small numbers (Tversky and Kahneman, 1973). The affect heuristic determines experts' perceived riskiness of various hazardous substances (Slovic et al., 1999) and the expert-assigned likelihood of offending of recently discharged mental patients (Slovic et al., 2000). Clinical psychologists underweight the base rate relative to a Bayesian decision maker (Meehl and Rosen, 1955). Expertise does not shield decision makers from the hindsight bias, as shown in a meta-study (Guilbault et al., 2004). Traders in a large investment bank, in both its international branches, were found to be hindsight-biased (Biais and Weber, 2009).

People are often willing to pay a premium for expert services that are mistakenly perceived to be competent (Powdthavee and Riyanto, 2015). Evidence supports the important role of anchoring on a list price by estate agents (Northcraft and Neale, 1987) and in legal judgment (Chapman and Bornstein, 1996; Englich and Mussweiler, 2001; Englich et al., 2006). The false consensus effect is exhibited by judges (Solan et al., 2008). Finance professionals also exhibit a false consensus effect by imputing on others, their own risk preferences (Roth and Voskort, 2014).

Experts, such as physicians and World Bank staff, exhibit framing effects, often of a similar magnitude to those observed with student populations. When presented with identical options, expert responses can flip depending on whether the outcomes are framed as gains or losses relative to a reference point (McNeil et al., 1982; Kahneman and Tversky, 1984; WDR, 2015). There is evidence of limited attention in wholesale car markets by wholesale car dealers (Lacetera et al., 2012).

In the experiments of Redelmeier and Shafir (1995), physicians were given a case study in which a patient suffers chronic hip pain. The physician, in the case study, is on the verge of referring the

case to an orthopedic consultant. At this point, the physician in the case study learns new information about drugs that the patient has not tried. In one treatment, physician-subjects learn that the patient has not yet tried Ibuprofen, and in the second treatment they learn that the patient had not yet tried Ibuprofen and Piroxicam. In the first treatment there were two choices: Ibuprofen+referral and plain referral. In the second treatment, in addition to these two choices, an extra choice was added: Piroxicam+referral. In the second treatment, despite more choices (three instead of two), 72% of the physician-subjects chose plain referral as compared to 53% in the first treatment.

Being subject to judgment heuristics and framing effects, experts can make systematic errors just like anyone else. Thus, market experience does not eliminate non-standard behaviors. But these are not the only reasons that experts exhibit potential biases. Moskowitz and Wertheim (2011) argue that the home ground advantage in sports arises not just from insider-knowledge of home players but also from expert-biases. The experts, in this case, are the match referees. Even if the referees are completely impartial and honest, they are subject to a natural human tendency to avoid the displeasure of those in the proximity. Thus, their discretionary authority typically tends to be exercised in favor of the home team, as for instance, in adding discretionary injury time minutes at the end of a football game. The idea of “home ground advantage” in sports is also used to explain the behavior of experts in the 2008 financial crises by Barth et al. (2013). Many financial experts held senior private sector positions in the financial sector before holding similar senior positions in the US Fed. The home team, in this case, becomes the private financial sector and the away team, the dispersed individual investors in the financial sector, whose interests were insufficiently protected.

In an admirable piece of work, the WDR (2015) documents several kinds of biases among its professional staff. These include confirmation bias, susceptibility to sunk costs, and the influence of framing. The WDR (2015, p. 18) is candid in its assessment of expert-bias: “This finding suggests that development professionals may assume that poor individuals may be less autonomous, less responsible, less hopeful, and less knowledgeable than they in fact are.” The WDR also suggests potential solutions to the problem of expert-bias. These include *dogfooding* (experts signing up and playing their own programs for real) and *red teaming* (having an adversarial outside team test that tests the proposals).

CHAPTER 20

Mental Accounting

20.1 Introduction

The idea of *mental accounting* is an essential part of behavioral economics with wide-ranging implications. Thaler (1999, p. 186), who has done more than anybody else to bring mental accounting to the center-stage in behavioral economics writes: “*I wish to use the term ‘mental accounting’ to describe the entire process of coding, categorizing, and evaluating events.*” Elsewhere (p. 183), he provides a definition of financial mental accounting: “*mental accounting is the set of cognitive operations used by individuals and households to organize, evaluate, and keep track of financial activities.*”

Many readers of this book are unlikely to have formally encountered the mental accounting literature. For this reason, the introduction is split into two parts. The first part gives some examples of mental accounting, while the second part gives an introduction to the chapter.

20.1.1 Some examples of mental accounting

Mental accounting is best motivated by the following insightful series of three examples that demonstrate how individuals create a system of mental accounts.

Example 20.11 (Thaler, 1999, p.183): *A former colleague of mine, a professor of finance, prides himself on being a thoroughly rational man. Long ago he adopted a clever strategy to deal with life’s misfortunes. At the beginning of each year he establishes a target donation to the local United Way charity. Then, if anything untoward happens to him during the year, for example an undeserved speeding ticket, he simply deducts this loss from the United Way account. He thinks of it as an insurance policy against small annoyances.*

Example 20.12 (Thaler, 1999, p.183): *A few years ago I gave a talk to a group of executives in Switzerland. After the conference my wife and I spent a week visiting the area. At that time the Swiss franc was at an all-time high relative to the US dollar, so the usual high prices in Switzerland were astronomical. My wife and I comforted ourselves that I had received a fee for the talk that would easily cover the outrageous prices for hotels and meals. Had I received the same fee a week earlier for a talk in New York though, the vacation would have been much less enjoyable.*

Example 20.13 (Thaler, 1999, p.184): *A friend of mine was once shopping for a quilted bedspread. She went to a department store and was pleased to find a model she liked on sale.*

The spreads came in three sizes: double, queen and king. The usual prices for these quilts were \$200, \$250 and \$300 respectively, but during the sale they were all priced at only \$150. My friend bought the king-size quilt and was quite pleased with her purchase, though the quilt did hang a bit over the sides of her double bed.

Classical economics assumes that individuals derive utility only from final levels of wealth.¹ In Example 20.12, the final wealth is the same whether the fee is received a week ago or after the talk (discounting is hardly an issue over such a small time interval), hence, if liquidity constraints do not bite, neoclassical economic theory cannot directly explain the stylized facts in the example.

In Example 20.13, the individual needed a queen sized quilt but purchased the oversized quilt because they are priced the same. Assuming that the individual derives greater utility from a better fitting quilt, why did she buy the oversized quilt?

Suppose that individuals have a system of *mental accounts*. Let us think of mental accounts, at the moment, as simply the assignment of transactions to different exogenously created categories, and money is not fully fungible across these categories. Classical economics assumes that one has a global financial account that simply aggregates all the individual accounts using the assumption of perfect fungibility of money. This makes the individual accounts redundant. Consider two salient properties of these mental accounts, supported by the evidence, that can explain Examples 20.11, 20.12.

1. Money is not fungible across mental accounts.
2. The individual does not exclusively focus on maximizing overall financial wealth. For instance, the objective might also be to limit the size of losses in individual accounts.

In Example 20.11, the Professor has an *account for unexpected losses* that is initially, and notionally, debited with some level of contribution to the charity. As unexpected losses occur, the contributions to charity are reduced to offset them so that the unexpected losses account does not show up a loss. In Example 20.12, the Thalers create a separate account for their visit, the Swiss mental account, that includes only incomes and expenditures from the Swiss trip. If the fee had been received a week earlier, it would not have been credited to this mental account. Within the Swiss mental account, the receipt of the fee offsets the loss arising from an appreciation of the franc. The explanation of Example 20.13 requires the concepts of *acquisition* and *transaction* utility that we develop later. We consider two more examples to illustrate the range of mental accounting phenomena.

Example 20.14 *In Kenya, many households do not invest sufficiently in health care prevention such as buying insecticide-treated mosquito nets. While credit constraints of households are part of the explanation, the bednets are not prohibitively expensive either. Dupas and Robinson (2013) apply mental accounting techniques to design a simple intervention. In the baseline intervention, individuals are provided with a lockable box with padlock, keys, and a passbook that they can use to keep track of the cash deposited in the lockable box. In three other treatments, they vary the degree of self-control by not providing the keys, or exerting a degree of social pressure. This intervention facilitates the creation of a mental account for health emergencies and is extremely effective in the take-up of the relevant health prevention. After 12 months, the take-up rates were between 66% and 97% in the four treatments.*

¹ Reference dependence can be readily incorporated into the classical subjective expected utility framework. However, without a theory of how human behavior differs in the domains of gains and losses, this adds little value; see Part 1 of the book for a discussion. In contrast, prospect theory offers a fundamentally different view of human nature in the two domains.

Example 20.15 *One might use mental accounting to categorize different sources of income. In the context of one-shot dictator games,² Cherry (2001) finds that the offer proposed by dictators depends on the source of money. When dictators are proposing a split over earned money, they exhibit self-regarding behavior 76% of the time. However, if they are using experimentally allocated money, they exhibit self-regarding behavior only 26% of the time. In this case, earned income resides in a mental account that has a different degree of flexibility towards supporting prosocial behavior.*

20.1.2 An introduction to the chapter

Section 20.2 briefly considers *framing effects* as a simple form of mental accounting. Under framing, informationally equivalent ways of presenting identical information that highlights different information cues, elicits different responses by the subjects. Since this phenomenon is pervasive and examples abound in this book, our discussion is brief and based on two classic experiments (McNeil et al., 1982; Kahneman and Tversky, 1984).

What kind of theoretical framework is suitable for an analysis of the wide diversity of mental accounting phenomena? One of Thaler's (1999) important insights is that Kahneman and Tversky's (1979) prospect theory can explain several mental accounting phenomena. Section 20.3 presents a general framework for when to segregate or integrate outcomes and applies this insight to the following simple example of a mental accounting problem. Consumers are more likely to travel to get a discount of \$5 on an item that costs \$15 as compared to a more expensive item that costs \$125. In each case, the saving is identical, which makes it difficult for neoclassical economics to account for the result.

Section 20.4 considers the *behavioral life-cycle model* (Shefrin and Thaler, 1988, 2004; Thaler and Shefrin, 1981; Thaler, 1980, 1985). The classical life-cycle model assumes perfect fungibility of money across all financial accounts. In contrast, evidence suggests that individuals code financial transactions into separate mental accounts and money is imperfectly fungible across all accounts. So, for instance, one might put aside money in a separate account for childrens' education. This would be unnecessary if money were fully fungible across accounts. Furthermore, the evidence suggests that the marginal propensity to consume out of these accounts is different. Thus, one might be less willing to raid the childrens' education amount relative to say, a mental account for drinking and leisure expenses.

If the marginal propensity to consume across distinct mental accounts differs, then the effect of salary composition, for instance, between wage and bonus, could have an effect on one's savings. Hence, the savings implications of defined pension contribution plans could be quite different from those suggested by the life-cycle model. The behavioral life-cycle model can also account for the reluctance of households to consume out of home equity, for the dependence of mandatory pension programs on demographic composition, and for the differential savings effects of small and large windfalls and of bonuses.

Why does the ticking of the taxi meter reduce one's pleasure from the taxi ride? Why do we normally prefer to prepay for a vacation? At the opposite extreme, why do we often prefer to buy consumer durables on installments? The life-cycle model in economics does not directly answer these stylized facts. For instance, delaying payments is preferable in the presence of time discounting, but in experiments with zero interest rate, subjects still prefer to prepay for vacations.

² A one-shot dictator game is a two-person game. In this game, a proposer (or dictator), proposes a division of a given sum of money between himself and a passive receiver who has an empty action set.

In Section 20.5, we use the *red* and *black* of mental accounting to address this class of questions (Prelec and Loewenstein, 1998). People prefer to have their mental accounts in the black (debt free) rather than in the red (under debt); this is a form of *debt aversion*.

We combine this insight with other insights from mental accounting. For instance, consumption that has already been paid for can be enjoyed as if it were free in the current period (*prospective accounting*). Furthermore, there is a *coupling* of benefits and costs. The thought of an impending payment reduces the current flow of utility from a good (ticking taxi meters). These insights allow us to explain the prepayment and postpayment for different kinds of goods, growing use of debit cards, use of tokens at casinos, dependence of mortgage payments on property values, popularity of fixed fee payment schemes, Club Med vacations, and gift giving among couples.

Individuals tend to break down complex problems into simpler ones that only require a subset of the information. These issues are examined in Section 20.6, under the rubric of *choice bracketing* (Read et al., 1999). An army general who focuses only on the battle at hand (*narrow bracketing*) might lose sight of winning the war (*broad bracketing*). Recall the problem of finding a taxi on a rainy day in New York that we considered in Part 1. One possibility is that taxi drivers set a daily target for earnings (narrow bracketing) rather than an income target over a longer period (broad bracketing). Narrow bracketing, in this case, leads to a backward bending labor supply curve (Camerer et al., 1997). Choice bracketing can also explain strategies used in “alcoholics anonymous,” i.e., set a daily target of relative abstinence (narrow bracketing) relative to a long-term target (broad bracketing). Other applications of choice bracketing include the desire for risk aggregation (Samuelson, 1963), and the equity premium puzzle (Benartzi and Thaler, 1995).

In classical economic theory, actual choices arise from well-defined initial preferences and well-defined subjective probability distributions over uncertain events. Contrary evidence has long been recognized in the work on *constructed preferences* (Slovic, 1995) and anchoring on irrelevant anchors (Tversky and Kahneman, 1974). Section 20.7 focuses on the experimental results from novel situations in which individuals are unlikely to have well-defined initial preferences (Ariely et al., 2003).

These results show that initial anchors have an inordinate effect on initial preferences. In this sense, initial preferences are *arbitrary*. However, once these preferences are embedded, subsequent preferences are *coherent*, for instance, in the sense that individuals display a well-defined demand curve; hence, the term *coherent arbitrariness*. These results can also explain the initial hype before mega sporting events, reluctance of firms to cut wages and inter-industry wage differentials, among others.

Section 20.8 considers the implications of sunk costs. In classical theory, sunk costs are irrelevant to current decisions. In actual practice, sunk costs often have an important bearing on current decisions (Kahneman and Tversky, 1984; Thaler, 1999; Shafir and Thaler, 2006). The reason is that individuals create a mental account for a particular category of expenditure. Sunk costs push this mental account into the red, which is aversive. Individuals then take actions to avoid closing accounts that are in the red. There is also some evidence that individuals react differently to sunk costs when they take the form of dollar amounts as compared to costs in terms of time (Soman, 2001). Empirical work also supports the idea that individuals have mental accounts for time (Rajagopal and Rha, 2009).

In Section 20.9, we consider an assortment of other mental accounting phenomena. In Section 20.9.1, we consider the distinction between *transaction utility* and *acquisition utility*. Transaction utility is salient when we make a purchase. The distinction between the two concepts can explain why we might buy less appropriate items at a sale, simply because they are cheaper, as in

Example 20.13. Section 20.9.2 answers the question of when to close an account (Thaler, 1999). The ideas developed here can explain the tendency to hold on to loss making stocks for too long and sell winners too early (Odean, 1998).

Section 20.9.3 shows that under mental accounting, *in-kind transfers* can dominate equivalent *cash transfers*. We give an illustrative example from the food stamps program in the US. The rationale for the existence of this program is in contrast to a standard result in neoclassical economics that demonstrates the dominance of cash transfers over in-kind transfers.

Section 20.9.4 extends loss aversion to a dynamic context using the observation that prior gains in casinos can induce individuals to take greater subsequent risks; this is known as the *house money effect* (Thaler and Johnson, 1990). However, no clear pattern on subsequent risk taking emerges in the case of prior losses. We also briefly comment on the relation between narrow bracketing and stock prices (Barberis and Huang, 2001).

In Section 20.9.5, we consider the *diversification heuristic* (Simonson, 1990; Read and Loewenstein, 1995). Diversification is generally good practice and consistent with rationality. However, in the presence of complex choices, individuals appear to engage in a naive form of diversification. For instance, when presented with an option to invest their wealth in n different options, individuals might diversify equally among the options. This is known as the $\frac{1}{n}$ *diversification heuristic* and has important implications for actual choices. For instance, in the domain of pension portfolios, it may lead people to choose portfolios that are dominated by stocks or by company assets (Benartzi and Thaler, 2001).

Section 20.9.6 considers the *mental accounting of goals*. Individuals often set goals and targets, e.g., completing term papers, quitting an addiction, losing a pre-specified amount of weight by a certain date, and so on. A literature that we consider in more detail in Part I, emphasizes the positive effects of goals (Locke and Latham, 1990). However, goals that are too ambitious are likely to be missed and can be counterproductive (Soman and Cheema, 2004); this is predicted by our model in Part I and supported by the evidence. Finally, we examine the *behavioral economics of retirement savings* in Section 20.9.7, and consider the role of several behavioral factors such as default options (opt-in or opt-out), initial contribution rates, and automatic escalation (Thaler and Benartzi, 2007; Benartzi and Thaler, 2013).

20.2 Framing and mental accounting

Mental accounting is used to describe the entire process of coding, categorizing, and evaluating events (Thaler, 1999). *Framing*, i.e., the presentation of a problem in informationally equivalent ways that highlight different information cues, is a form of mental accounting. Framing is not hard to demonstrate.

Example 20.16 (Kahneman and Tversky, 1984): *Imagine that the US is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows:*

Positive Framing: If program A is adopted, 200 people will be saved. If program B is adopted, there is a one-third probability that 600 people will be saved and a two-thirds probability that no people will be saved. Which of the two programs would you favor?

Negative Framing: If program C is adopted, 400 people will die. If program D is adopted, there is a one-third probability that nobody will die and a two-thirds probability that 600 people will die. Which of the two programs would you favor?

Under positive framing, the programs are framed in terms of *lives saved*. The result was that 72% of the respondents chose A, the safe option, and 28% chose B, the risky option. Under negative framing, only 22% chose the safe option, C, and 78% chose the risky option, D. The corresponding options (A and C; B and D) in each of the two frames are identical. Thus, in neoclassical economics, if A is preferred to B, then C must be preferred to D (or vice versa). The results show that individuals are risk averse for gains and risk loving for losses, which is an essential feature of prospect theory (see Part 1); the results also show that framing matters.

These results have been replicated for experienced professionals. For instance, when these experiments were conducted with World Bank staff, the outcome was similar. In this case, 75% chose the safe option under positive framing and 66% chose the risky option under negative framing (WDR, 2015, pp. 181–2).

Consider the following problem in which experienced physicians took part. Subjects in experiments received statistical information on the outcomes of two different treatments for lung cancer. Two different frames, the survival frame (positive framing) and the mortality frame (negative framing), were used. Both frames had identical information content that, in neoclassical economics, is not predicted to affect the outcomes.

Example 20.17 (McNeil et al., 1982): *In the “survival frame,” the choices were as follows:*
Surgery: Of 100 people having surgery 90 live through the postoperative period, 68 are alive at the end of the first year and 34 are alive at the end of five years.
Radiation therapy: Of 100 people having radiation therapy all live through the treatment, 77 are alive at the end of one year and 22 are alive at the end of five years.
In the “mortality frame,” the choices were as follows:
Surgery: Of 100 people having surgery 10 die during surgery or the post-operative period, 32 die by the end of the first year and 66 die by the end of five years.
Radiation therapy: Of 100 people having radiation therapy, none die during treatment, 23 die by the end of one year and 78 die by the end of five years.

The two informationally identical frames have a significant effect on outcomes. In the survival frame, 18% of the respondents favored radiation therapy. But in the mortality frame, this rose to 44%. The framing effect was exhibited equally by experienced physicians, statistically sophisticated business students, and clinic patients. Clearly, expertise does not mitigate mental accounting.

20.3 Prospect theory and mental accounting

Prospect theory is a convenient vehicle for modeling several aspects of mental accounting (Thaler, 1985). The central concepts in prospect theory such as reference point, loss aversion, and different attitudes to risk in the domain of gains and losses, can aid in the explanation of several mental accounting phenomena. Consider the following well-known experiment.

Example 20.18 (Tversky and Kahneman, 1981, p. 457): *Imagine that you are about to purchase a jacket for (\$125)/[\$15] and a calculator for (\$15)/[\$125]. The calculator salesman informs you that the calculator you wish to buy is on sale for (\$10)/[\$120] at the other branch of the store, located 20 minutes drive away. Would you make the trip to the other store? Two versions of the problem are given—in one version the figures in parentheses are given (\$125 for the jacket, \$15 for the calculator, and \$10 for the calculator in the sale) while in the other*

version, figures in square brackets are given (\$15 for the jacket, \$125 for the calculator, and \$120 for the calculator in the sale). At a price of \$15(\$125), 68%(29%) of subjects are willing to travel.

Let us now develop a framework to explain the finding in Example 20.18, as well as other similar findings. Denote wealth by y , and the reference point by r . Define $x = y - r$ as the increment in wealth relative to the reference point. Then, the power form of the utility function, v , in prospect theory (see Part I of the book for details) is given by

$$u(x) = \begin{cases} x^\beta & \text{if } x \geq 0 \\ -\lambda(-x)^\beta & \text{if } x < 0 \end{cases}. \quad (20.1)$$

Let us use the experimentally observed values of the parameters of the utility function: $\beta = 0.88$ and $\lambda = 2.25$ (Tversky and Kahneman, 1992). Using the properties of the prospect theory preferences, all of which are encapsulated in (20.1), Thaler (1985) gives the following plausible mental accounting conjectures.

1. *Multiple gains are segregated:* The utility function for gains under prospect theory is strictly concave. Hence, if there are several positive gains, they should not be combined (because the marginal utility of a larger amount is smaller). Instead *gains should be segregated*, perhaps temporally. The following analogy is helpful: If one receives a box of sweets one typically does not finish it in one go but tries to savor the enjoyment over time. Also, we typically do not like to wrap all Christmas presents in one box. As a numerical example, if $x = 100$, then one prefers to split it up into two equal gains of 50 each; using (20.1) we get

$$(100)^{0.88} < (50)^{0.88} + (50)^{0.88} \Rightarrow 57.544 < 62.535,$$

which is true. This result holds more generally for a strictly concave function v in gains such that $v(0) = 0$. Let $a > 0$ and $b > 0$ be two strictly positive outcomes, expressed relative to the reference point. Then strict concavity of v in gains implies that $v(b+a) - v(a) < v(b) - v(0)$, which can be rewritten as $v(b+a) < v(a) + v(b)$. Thus multiple gains are segregated.

2. *Multiple losses are integrated:* The utility function under prospect theory is strictly convex for losses. Hence, using the converse argument to the previous case, *losses should be integrated*. For instance, if $x = -100$, then the decision maker will prefer to integrate two equal losses of 50; using (20.1), we get

$$-2.25(100)^{0.88} > -2.25(50)^{0.88} - 2.25(50)^{0.88} \Rightarrow -129.47 > -140.7,$$

which is true. More generally, consider two strictly positive outcomes, $a > 0$ and $b > 0$, expressed relative to a reference point. Suppose that v is strictly convex in losses and $v(0) = 0$, then $v(0) - v(-a) > v(-b) - v(-(b+a))$ or $v(-b) + v(-a) < v(-(b+a))$. Thus, it is optimal to integrate losses. Consequently, sellers will attempt to lump together items for sale; for instance, adding options to a car at the time of purchase. It is as if the additional options have an inelastic demand, once they are combined with a larger purchase.

3. *Mixed net gains are integrated:* Individuals using prospect theory exhibit *loss aversion*. So the absolute value of the disutility arising from a small loss is equivalent to the utility arising from a larger gain. Insofar as individuals do not like their mental accounts to go into the

“red” (see Section 20.5), they should *integrate small losses with larger gains*. Thus, using (20.1), a loss of \$5 can be offset against a gain of 12.565 because

$$-2.25 (5)^{0.88} + (12.565)^{0.88} = 0.$$

This is also sometimes known as the *principle of cancellation*; one can cancel the smaller loss against a larger gain. If the two outcomes (12.565, -5) are integrated for a net outcome of 7.565, then the utility is $(7.565)^{0.88} > 0$. Consider any outcome pair $(b, -a)$, where each outcome is measured relative to a reference point. Let $0 < a < b$, so $b - a > 0$ (net gain), then we have $v(b) + v(-a) < v(b - a)$. The proof is straightforward. By the definition of loss aversion, we have $v(a) < -v(-a)$, thus,

$$v(b) + v(-a) < v(b) - v(a). \quad (20.2)$$

But v is strictly concave in gains, so

$$v(b) - v(a) < v(b - a) - v(0). \quad (20.3)$$

$v(0) = 0$ due to reference dependence, thus, from (20.2), (20.3) we get

$$v(b) + v(-a) < v(b - a),$$

as claimed.

4. *Mixed net losses and the silver lining principle*: Consider any outcome pair $(b, -a)$, where each outcome is expressed relative to a reference point such that $0 < b < a$, so $b - a < 0$ (net loss). In this case, $v(b) + v(-a) \leq v(b - a)$, so, in general, we cannot say anything about integration/segregation of outcomes. If b is small relative to a , for instance, the pair (40, -6000), then we have large net losses $b - a$. In this case, we are likely to have $v(b) > v(b - a) - v(-a)$, provided that the curvature of the utility function is not too flat or the net loss is very large. Suppose that the parameter of the utility function in (20.1), $\beta = 0.75$, then for the pair (40, -6000), the condition $v(40) > v(40 - 6000) - v(-6000)$ requires that

$$(40)^{0.75} > -2.25 (5960)^{0.75} - (-2.25 (6000)^{0.75}),$$

or $15.9 > 7.7$, which is true. In this case, one should *segregate large losses from small gains*; this is known as the *silver lining principle*. At the other extreme, when b is not too different from a , as in the pair (40, -50), then integration ($v(b) + v(-a) < v(b - a)$) may be optimal. For the pair (40, -50), integration is optimal if $v(40) + v(-50) < v(40 - 50)$. Applying the prospect theory utility function with $\beta = 0.75$, we get

$$(40)^{0.75} - 2.25 (50)^{0.75} < -2.25 (10)^{0.75}$$

or $-26.4 < -12.7$, which is true.

Let us revisit Example 20.18 in the light of the silver lining principle. Consider the outcome pair $(b, -a) = (5, -125)$; the item is expensive and costs \$125, and one can save \$5 by travelling to the other store. So $0 < b < a$, but b is small. From the silver lining effect, in this case, one

segregates the large loss from the small gain, so one does not travel to get the extra discount. In terms of the prospect theory utility function with $\beta = 0.75$,

$$(5)^{0.75} > -2.25(120)^{0.75} - (-2.25(125)^{0.75}),$$

or $3.3437 > 2.5362$, which is true. In the other case of the cheaper product that costs \$15, we have $(b, -a) = (5, -15)$, the converse arguments show that integration is optimal, and one undertakes the trip to the other store to save \$5. The usual calculation in this case is

$$(5)^{0.75} - 2.25(15)^{0.75} < -2.25(10)^{0.75},$$

or $-13.806 < -12.653$, which is true.

20.4 The behavioral life-cycle model

A central assumption in various versions of the life-cycle model in economics is that money is fungible across various accounts. This assumption allows one to write the intertemporal budget constraint in the following well-known form: The present discounted value of consumption is bounded above by the present discounted value of income from all sources. In a series of papers, Shefrin and Thaler (1988, 2004), Thaler and Shefrin (1981), and Thaler (1980, 1985) have questioned this assumption.

Another important feature of their critique, in relation to the life-cycle model, is to model the self-control problems that one's own future selves might have and how, in anticipation of these problems, one's previous selves take remedial measures. The previous selves are encapsulated in the form of a *planner*, who cares for the utility of all the subsequent selves. All subsequent selves, one for each time period, termed as the *doers*, are myopic in that they care only about maximizing their own utility.³ Shefrin and Thaler describe models with these features as *Behavioral Life-Cycle Models* (BLC).

We have found it more useful to separate the issues arising from a planner–doer framework and those that arise from mental accounting. The planner–doer framework and issues of self-control are treated in Parts 3 and 6 of the book. In this chapter, we focus on life-cycle considerations that arise when one takes into account issues of mental accounting.

20.4.1 *Mental accounting and the life-cycle model*

In the dynamic context, as in the static context, empirical evidence suggests that individuals create a system of accounts such that money is not fully fungible across the accounts. Evidence suggests fairly detailed and disaggregated mental accounting. However, for tractability, we consider disaggregation into the following three main mental accounts (Shefrin and Thaler, 2004). A *current income account* (*I*), a *current asset account* (*A*), and a *future wealth account* (*F*). A distinguishing feature of these accounts is that some are more tempting to invade than others. In particular, accounts that are considered “wealth” are less tempting than those that are considered as “income”.

³ The idea that future selves have potential self-control problems that a current self might correct (the two-self models) can also be found in Elster (1979), Margolis (1982), and Schelling (1984).

In classical economic theory, there is no issue of “invading” financial accounts. However, much empirical evidence indicates that individuals have present-biased preferences (see Part 3 of the book); such individuals are particularly susceptible to invade financial accounts.

Suppose that an individual has a stream of income over T periods, given by $y_1, y_2, \dots, y_T$. The individual retires in period T , so $y_T = 0$. There are perfect capital markets and the interest rate is zero. Suppose that the individual saves a fraction, λ , of income each period for retirement; perhaps the individual takes part in a pension scheme that deducts a fraction of income in each period. In time period $0 \leq t < T$, the account position is as follows.

The *current income mental account* is given by

$$(1 - \lambda)y_t \equiv I_t. \quad (20.4)$$

The *current asset mental account* is given by the cumulative sum

$$\sum_{\tau=1}^{t-1} [(1 - \lambda)y_{\tau} - c_{\tau}] \equiv A_t, \quad (20.5)$$

where $(1 - \lambda)y_{\tau} - c_{\tau}$ is the discretionary, or non-pension, savings made in period $\tau \leq t - 1$. The *future wealth mental account* is given by

$$\sum_{\tau=t+1}^T (1 - \lambda)y_{\tau} + \sum_{\tau=1}^T \lambda y_{\tau} \equiv F_t. \quad (20.6)$$

The first term on the LHS in (20.6) gives the disposable income (after pension deductions have been made) summed across all future periods and the second term gives the pension wealth over the lifetime.

It is easiest for the individual to spend from the current income account, more difficult to invade the current asset account, and most difficult to invade the future wealth account. The current asset account is sometimes disaggregated into sub-accounts. For instance, a childrens' education account could be designed to finance a child's education; such an account would be very difficult for an individual to invade.⁴

There is also evidence that small windfalls are coded by individuals as part of the current income account while larger windfalls are coded as part of the current asset account (Milkman and Beshears, 2009). Home equity is treated by some individuals as part of the current asset account, while others might treat it as part of the future wealth account. Asset income is usually kept in the current asset account but there is some evidence that dividend income might be kept in the current income account (Shefrin and Statman, 1984). Shefrin and Thaler (2004, p. 401) write: “While the mental accounting system described here may seem bizarre to economists, it is remarkably similar to the accounting systems used by most private universities.” Graham and Isaac (2002) present survey evidence, which supports this mental accounting practice.

The difference in the degree of difficulty of invading various accounts suggests that the individual has an *account specific marginal propensity to consume*. Shefrin and Thaler (2004) report an experiment with MBA students in Santa Clara University that provides support for this view.

⁴ A long-standing co-author of mine has a practice of keeping a book account. His large, and constantly growing, book collection is a testament to the effectiveness of such accounts.

Courant et al. (1986) find that the marginal propensity to consume out of the current asset account is about 0.25, but the marginal propensity to consume out of the future wealth account is significantly smaller.

Individuals are also likely to code labor income as current income and capital income as current asset income. Holbrook and Stafford (1971) find that the marginal propensities to consume out of labor and capital income are, respectively, 0.9 and 0.7. When the marginal propensity to consume is account-specific and constant, we get that consumption, c_{tj} , at time t , financed from account $j = I, A, F$ is given by

$$c_{tI} = \alpha_I I_t, \quad (20.7)$$

$$c_{tA} = \alpha_A A_t, \quad (20.8)$$

$$c_{tF} = \alpha_F F_t, \quad (20.9)$$

where $\alpha_I, \alpha_A, \alpha_F$ are, respectively, the constant marginal propensities to consume from the income, current asset, and the future wealth mental accounts. Shefrin and Thaler (2004) interpret the empirical evidence as suggesting the following magnitudes

$$\alpha_I \approx 1; 0 < \alpha_A < 1; \alpha_F \approx 0. \quad (20.10)$$

Levin (1998) used the Longitudinal Retirement Survey for individuals in their late 50s or older and distinguished between ten consumption categories and finds support for the pattern of marginal propensity to consume given in (20.10).

Using (20.7) to (20.10), one expects that

$$c_{tI} \approx I_t; c_{tA} = \alpha_A A_t; c_{tF} \approx 0. \quad (20.11)$$

One role of mental accounting is to act as a mechanism of self-control (Thaler, 1999). But how does the individual decide which mental accounts to create? Is it always clear which mental account a new item of expenditure should be assigned to? It is quite possible that there is some ambiguity with respect to both questions. Is it possible that an individual might engage in creative book-keeping or try to find loopholes to sidestep the very controls that were initially designed to help in the self-control problems? These sorts of questions are posed by Cheema and Soman (2006). They cite the following example from Mark Twain. “Having limited himself to one cigar a day, he started shopping for bigger and bigger cigars until he had each made to such proportions that he ‘could have used it as a crutch.’” This is an example of a loophole in the original mental account (1 cigar a day) that did not specify the size of the cigar.

Elsewhere, Soman and Cheema (2001) ask which mental accounts should unexpected windfalls be classified into? Because such windfalls are unexpected, there is potential ambiguity about which mental account they could be classified into. It is shown that this gives an opportunity for individuals to use such gains to clear accounts that are in the red. For instance, such gains could be used to write off sunk costs.

20.4.2 *Some predictions of the mental accounts model*

We now examine some of the predictions arising from the application of mental accounts to the life-cycle model.

(1) *Pensions and savings*: Choi et al. (2009) document evidence consistent with mental accounting from the 401(k) pension plans in the US, based on a change in rules in 2003. In contrast to the assumption of perfect fungibility across accounts, they found that individuals do not adequately consider the spillover effects across various accounts in their portfolios. Thus, they act as if they considered these accounts to be separate, as in the mental accounting literature.

Suppose that an employee pension scheme makes it mandatory to contribute an extra dollar of income. Thus, the total contribution to the pension scheme is now the sum of the initial contribution, λy_t , and the extra dollar, for a total of $1 + \lambda y_t$ dollars in contribution. What can we say about the effect on the total saving made by the individual? The prediction of the standard life-cycle model is that aggregate savings are unaffected. This is because money is fungible across accounts and the individual offsets one unit of private savings to make up for the extra unit of mandatory saving towards the pension account. These arguments are routinely used to illustrate neutrality effects in fully funded pension plans in the classical life-cycle model.

Now suppose that in the presence of mental accounting, the marginal propensities to consume differ across the three main accounts (I , A , and F). The method of financing extra pension contributions will now affect the amount of private savings that are offset. If the extra dollar of income comes out of the current income account, then, from (20.7), (20.11), consumption falls by one dollar so private savings are unaffected. If the extra dollar of pension contributions comes out of the current asset account, then it is shared between a fall in consumption and a fall in savings; the respective shares depend on the size of α_A ; evidence supports this prediction (Diamond and Hausmann, 1984; King and Dicks-Mireaux, 1982). Insofar as the young have a higher marginal propensity to consume out of the current asset account, α_A , relative to the old, the implications of mandatory pension programs are also contingent on the demographics of a society.

In neoclassical theory, the mix between dividends and capital gains is irrelevant for consumption. However, empirical evidence suggests that dividend and capital gains incomes are not fungible and realized dividends have a larger, positive, impact on consumption (Baker and Wurgler, 2007). Thus savings of individuals are lower if they receive a greater fraction of their income as dividends, relative to capital gains. Furthermore, dividends are withdrawn from the household portfolio at a relatively larger rate than capital gains, which is consistent with the predictions of the mental accounting model.

(2) *Housing wealth*: What are the implications of rapidly rising house prices on the consumption and savings behavior of households? The mental accounts approach would first attempt to uncover evidence on how households code an increase in equity income arising from an increase in house prices. Evidence suggests that households code housing wealth somewhere in between the current asset account and the future wealth account, so the marginal propensity to consume out of an increase in housing equity is relatively small; evidence indicates it is not statistically different from zero (Skinner, 1989). This result contrasts with the life-cycle model in which individuals should interpret an increase in housing equity as an increase in the present discounted value of income. Thus, they should have the same marginal propensity to consume out of it as they have out of current income. Evidence, however, suggests that even the elderly are reluctant to consume out of housing equity. The possibility that such behavior on the part of the elderly simply reflects the bequest motive can be ruled out (Hurd, 1987).

(3) *Bonuses and windfalls*: In the life-cycle model, bonuses are one component of income, so whether a part of one's income is received as bonuses or otherwise has no effect on consumption and savings behavior. However, in the Thaler–Shefrin framework, bonuses are coded as a part of the current asset account. Therefore, the marginal propensity to consume out of bonuses is lower than that of regular income. The implication is that bonuses increase the savings of individuals.

Thaler and Shefrin (2004) report supportive evidence from an empirical study based on Japanese data.

What about windfall income? The prediction of the mental accounting model depends on the size of the windfall income (Thaler, 1990). When windfalls are small relative to income, they are coded as part of the current income account where the marginal propensity to consume is 1. However, large windfalls are coded as a part of the current asset account where the marginal propensity is relatively lower. Landsberger (1966) considered the receipt of money received by Jewish families as restitution from Germany after World War II. Consistent with the prediction of the mental accounting model, those who received the largest windfalls as a fraction of their incomes, had the smallest marginal propensity to consume.

Further evidence supportive of mental accounting in the domain of small windfalls is presented in Heilman et al. (2002), Milkman and Beshears (2009), and Abeler and Marklein (2013). Milkman and Beshears compared the consumption behavior of online grocery store consumers when they were given small vouchers to spend at the store (about \$10) and when they were not given such vouchers. Controlling for customer fixed effects, they found that a \$10 voucher increased spending by \$1.59. Furthermore the target of the increase were items that consumers would typically not buy from the store. By contrast, since these small windfalls comprise a negligible fraction of total lifetime income, they should leave current consumption unaffected in the classical model; this is rejected by the evidence.

The source of the windfall matters too. Income from the sale of a security is typically coded as a part of the current income account, hence, the marginal propensity to consume is 1 (Summers and Carroll, 1987). However, unrealized capital gains are treated as a part of the current asset account, hence, relatively less likely to be reflected in current consumption (Hatsopoulos et al., 1989). Whether windfalls are anticipated or unanticipated is reflected in consumption-effects that exceed those that one would expect from the neoclassical life-cycle model (Arkes et al., 1994).

20.5 The red and black of mental accounting

Prelec and Loewenstein (1998) distinguish between two binary states of mental accounts. A mental account is said to be in the *red* if it is in debt and in the *black* if it is debt free. Evidence indicates that individuals derive satisfaction if their accounts make a transition from red to black. In other words, individuals behave as if they have *debt aversion*.

The standard model of discounted utility in discrete time proceeds as follows. Suppose that the consumer makes a purchase and T_b , T_c , respectively, are the sets of discrete time periods for which the benefits accrue and payments are made. The purchase generates a dated stream of utilities or benefits, $\{u_t\}_{t \in T_b}$, and a dated stream of payments or costs, $\{p_t\}_{t \in T_c}$. So, for instance, if a consumer buys a washing machine that gives utility over periods 1, 2, 3, and requires payments over periods 2, 3 (as in mortgage purchases), then $T_b = \{1, 2, 3\}$, $T_c = \{2, 3\}$. Suppose that the individual applies the discount factors δ_b , δ_c , respectively, to the streams of benefits and costs. Then, in the discounted utility model, the consumer should make the purchase if

$$\sum_{t \in T_b} \delta_b^t u_t - \mu \sum_{t \in T_c} \delta_c^t p_t > 0, \quad (20.12)$$

where μ is the marginal utility of money that converts a payment term into its utility equivalent. The *discounted utility decision criterion* given in (20.12) implies the following.

1. DU1: On account of discounting, the consumer prefers to pay later rather than earlier.
2. DU2: For any two arbitrary but different streams of benefits and costs *with identical present discounted values*, the purchase decision is identical. In other words, the characteristics of the two streams such as whether payments precede or succeed the stream of benefits are not relevant to the purchase decision.

Prelec and Loewenstein (1998) give the following two examples, based on experimental evidence, that violate DU1, DU2.

Example 20.19 (*Vacation*): Suppose that you are going to be away on vacation, six months from now. The cost of the vacation is \$1200. There are two methods of payment. In method 1, you pay six monthly payments of \$200 prior to the vacation. In method 2, you pay six payments of \$200 after returning from the vacation. In experiments, 60% of the subjects preferred method 1 (preponement of payment) despite an implicit interest penalty of about \$50 arising from paying earlier.

Example 20.20 (*Consumer durables*): Six months from now you will be purchasing a clothes washer/dryer that costs \$1200. There are two methods of paying. In method 1 you can pay in six equal monthly installments of \$200 starting now while in method 2, you pay in six equal monthly installments of \$200 with the first installment beginning in six months time. Now 84% of subjects prefer method 2 (postponement of payment).

In other examples, Prelec and Loewenstein find that the preference for prepaying does not apply only to luxury goods (e.g., vacations) but also to items covering miscellaneous living expenses. In another experiment, they find that a majority of subjects decline the offer of advance payment for work. Each of these examples show a form of debt aversion, which takes the form of aversion to consumption before paying for it, or getting advance payments for future work. One possible explanation for some of these results is that individuals prefer an increasing profile of consumption. In light of these anomalies, Prelec and Loewenstein propose an alternative theory that is based on the following main considerations.

- A1: (*Prospective accounting*) Consumption that has already been paid for can be enjoyed as if it were free in the current period. If payments are made prior to the enjoyment of benefits from a good, then, at the time of making the payments, the pain arising from the payments is mitigated by thoughts of future benefits.
- A2: (*Coupling*) While enjoying the benefits from a good, consumers could be thinking of the payments. Conversely, while making the payments, consumers could be thinking of the benefits. Prelec and Loewenstein term this as *coupling* between benefits and costs. Alternative methods of payment have different degrees of coupling. So, for instance, credit cards weaken coupling while cash payments strengthen it.
- A3: (*Prorating*) Essentially this implies that at any time, the remaining future benefits or costs are equally divided (prorated) over the remaining periods when the consumer constructs hedonic benefits and costs.

The theory can be formalized as follows.

First, we identify how the consumption experience (the benefit from consumption) of the consumer is modified by thoughts of payments. Define *hedonic benefits* or *net experienced utility* in period t as

$$u_t - \mu \hat{p}_t, \quad (20.13)$$

where u_t is the *experienced utility* at time t ; $\hat{p}_t$, to be formally defined below, is the imputed (not actual) cost in period t ; and μ is the marginal utility of money. The imputed psychological cost, $\mu\hat{p}_t$, reduces some of the benefits from consumption.

Next, we take into account how the payment experience (the cost of consumption) is influenced by thoughts of benefits received from the good. Define hedonic costs or *net experienced disutility* in period t as

$$\hat{u}_t - \mu p_t, \quad (20.14)$$

where $\hat{u}_t$ is the imputed (not actual) psychological benefits when one makes a payment of p_t . The imputed benefit offsets some of the cost of the good and if $\hat{u}_t - \mu p_t > 0$, it more than fully offsets the cost.

Second, we introduce a loss-aversion by a parameter $\eta \in [0, 1]$; as in prospect theory, losses are more salient than gains. Loss aversion applies to both experienced utility and experienced disutility, i.e., when $u_t - \mu\hat{p}_t < 0$ or $\hat{u}_t - \mu p_t < 0$. The “loss-aversion weights” are defined as follows.

$$\begin{cases} 1 - \eta & \text{if } u_t - \mu\hat{p}_t \geq 0 \text{ or } \hat{u}_t - \mu p_t \geq 0 \\ 1 + \eta & \text{if } u_t - \mu\hat{p}_t < 0 \text{ or } \hat{u}_t - \mu p_t < 0 \end{cases}$$

Define $T_b^+ \subset T_b$ for which $u_t - \mu\hat{p}_t \geq 0$ and $T_b^- \subset T_b$ for which $u_t - \mu\hat{p}_t < 0$. Similarly, define $T_c^+ \subset T_c$ for which $\hat{u}_t - \mu p_t \geq 0$ and $T_c^- \subset T_c$ for which $\hat{u}_t - \mu p_t < 0$.

Third, we use the following *purchase criterion* that captures the *red and black of mental accounts*. The individual buys the good if and only if

$$\begin{aligned} & \left(\sum_{t \in T_b^+} \delta_b^t (1 - \eta) [u_t - \mu\hat{p}_t] + \sum_{t \in T_b^-} \delta_b^t (1 + \eta) [u_t - \mu\hat{p}_t] \right) \\ & + \left(\sum_{t \in T_c^-} \delta_c^t (1 + \eta) [\hat{u}_t - \mu p_t] + \sum_{t \in T_c^+} \delta_c^t (1 - \eta) [\hat{u}_t - \mu p_t] \right) > 0. \end{aligned} \quad (20.15)$$

The first bracket gives the present discounted value of hedonic benefits while the second bracket is the present discounted value of hedonic costs. In the absence of loss aversion, i.e., $\eta = 0$ and the absence of hedonic affects, i.e., $\hat{u}_t = \hat{p}_t = 0$, (20.15) collapses to the standard discounted utility case of (20.12).

To operationalize the purchase criterion given in (20.15), we must define how the consumer imputes the values $\hat{u}_t, \hat{p}_t$. In principle, $\hat{u}_t$ and $\hat{p}_t$ can depend non-linearly on the entire stream of benefits and costs respectively. However, assuming a linear relation between imputed costs/benefits and the entire stream of costs/benefits, we get

$$\hat{p}_k = \sum_{t \in T_c} \omega_{tk}^c p_t; \quad k \in T_c, \quad (20.16)$$

where ω_{tk}^c is the fraction of period t price, p_t , that contributes to the imputed cost at time k , $\hat{p}_k$. If past prices do not influence current imputed costs, then $\omega_{tk}^c = 0$ for $t < k$.

Similarly, define

$$\hat{u}_k = \sum_{t \in T_b} \omega_{tk}^b u_t; \quad k \in T_c, \quad (20.17)$$

where ω_{tk}^b is the fraction of date t benefit that contributes towards the imputed date k benefits. Notice that variations in $\hat{p}_k$, $\hat{u}_k$ over time are caused by variations in the coefficients $\omega_{tk}^c, \omega_{tk}^b$ across time.

Definition 20.8 *The entire set of coefficients $\{\omega_{tk}^c, \omega_{tk}^b\}_{t,k}$ is called a consumer's mental accounting system.*

There are several alternative ways of constructing an individual's mental accounting system. Prelec and Loewenstein use *prospective accounting*, *coupling*, and *prorating over multiple events* (see A1–A3, above) to operationalize (20.16), (20.17) as follows.

$$\hat{p}_k = \alpha \left(\frac{u_k}{\sum_{t>k} u_t} \right) \sum_{t>k} p_t, \quad (20.18)$$

$$\hat{u}_k = \beta \left(\frac{p_k}{\sum_{t>k} p_t} \right) \sum_{t>k} u_t, \quad (20.19)$$

where $\alpha \geq 0, \beta \geq 0$ are the *coupling parameters*. α denotes the marginal impact of the importance of future payments in reducing net experienced utility or hedonic benefits at time t . β denotes the marginal impact of the importance of future benefits in buffering the pain of current payments or hedonic costs. So, for instance, if $\alpha = \beta = 0$ and $\eta = 0$, then (20.15) collapses to the standard discounted utility model in (20.12). In forming imputations $\hat{u}_k, \hat{p}_k$ in period k , the consumer ignores costs and benefits that arose in earlier periods (*prospective accounting*). The benefits and costs are *prorated* (which is reflected in the terms $u_k / \sum_{t>k} u_t, p_k / \sum_{t>k} p_t$).

Substituting (20.18), (20.19) in (20.13), (20.14) we get the hedonic benefits (or net experienced utility) and hedonic costs (or net experienced disutility) at time k as follows

$$\text{Hedonic benefits: } u_k - \mu \alpha \left(\frac{u_k}{\sum_{t>k} u_t} \right) \sum_{t>k} p_t. \quad (20.20)$$

$$\text{Hedonic costs: } \beta \left(\frac{p_k}{\sum_{t>k} p_t} \right) \sum_{t>k} u_t - \mu p_k. \quad (20.21)$$

Suppose that an individual prepays for an item, as in the vacation example. Then, given prospective accounting, hedonic benefits at date k , using (20.20), are simply given by actual benefits at date k , u_k (since $\sum_{t>k} p_t = 0$). However, if the individual buys an item on installments to be paid after the purchase, as in the case of consumer durables, then the thoughts of future payments mitigate the current benefits, as shown in (20.20).

Analogously, (20.21) captures the effect on hedonic costs when there are some future benefits to be derived from the purchase. In the special case when all payments are to be made after all benefits from the item have ceased to exist, then hedonic costs at date k are simply given by




Figure 20.1 Hedonic costs (light vertical lines) and hedonic benefits (bold shapes) for the two cases of prepayment and postpayment for vacations.

$-\mu p_k$ (since $\sum_{t>k} u_t = 0$). Hence, the “state” of the cost and the benefit accounts, *red or black*, influences hedonic costs and benefits; and, by implication, the purchase decision in (20.15).

Prepaying increases the current hedonic benefits by shrinking the term $\sum_{t>k} p_t$ (see (20.20)). However, the purchase decision in (20.15) clarifies the disadvantage; delayed payments bite less on account of the discount factor. This trade-off will sometimes make prepaying an optimal choice as in the vacation example. In other cases, it is optimal to pay later, as in the consumer durables example.

We illustrate this with the case of our vacation example. The consumer durables case is left as an exercise for the reader. The vacation example is represented in Figure 20.1. The vertical axis measures hedonic costs and benefits, as defined in (20.20), (20.21). The horizontal axis measures time. For the two alternative payment options, the bold shapes (the rectangle in panel A and the irregular shape in panel B) represent the temporal hedonic benefits. The lighter vertical lines represent the temporal hedonic costs at a particular date.

The vacation lasts over the time interval T_b . Suppose first that the vacation is prepaid in equal monthly installments of $\$p$ (panel A). Here, the benefit is received during the vacation period when all costs have been prepaid, so using (20.20), the hedonic benefits (or net experienced utility) per period equal the actual benefits during each time period in the vacation, u_t , $t \in T_b$. The experienced utility is shown by the area of the rectangle in panel A, after the vacation has been paid for by monthly installments. On the other hand, using (20.21), the hedonic costs (or net experienced disutility), shown by the vertical bars, falls over time because the outstanding payments $\sum_{t>k} p_t$ get reduced (while $\sum_{t>k} u_t, p_k$ are constant for all k). Thus, just before the vacation begins, the cost of the vacation is no longer aversive; the anticipation of future utility from the vacation coupled with prospective accounting swamps all costs.

Now consider panel B, where all payments are made after enjoying the vacation. As the vacation progresses, the remaining benefits of the vacation, $\sum_{t>k} u_t$, fall. Hence, using (20.20), the hedonic benefits fall as the vacation progresses. Towards the end of the vacation, the remaining benefits are minimal while the complete set of payments looms large so that the hedonic benefits

become negative. The vacation benefits have all been enjoyed in the past, so given the *prospective accounting* assumption, the hedonic costs stay the same in all periods following the vacation (or see (20.21)). It is quite possible that despite making adjustments for discounting, pre-paying for a vacation can be the optimal choice depending on the relative sizes of α , β and the streams of benefits and costs.

We now briefly outline some other predictions of the model.

1. **Mortgage payments and property values:** Generally speaking, hedonic benefits and hedonic costs can be thought of as two separate accounts. The accounts can be in the red or in the black. Individuals have a preference to keep the accounts in black (debt aversion). Hence, when the benefits from an item are declining, individuals would like to accelerate payments, otherwise the benefits account will go into red. This is obvious from (20.20); as $\sum_{t>k} u_t$ falls, one way to keep hedonic benefits high is to make a large current payment, p_k , in order to reduce the outstanding payments, $\sum_{t>k} p_t$. Therefore, one prediction of the model is that mortgage payments will increase if property values are anticipated to fall.
2. **Growing use of debit cards:** If individuals are debt-averse and prepayment increases the hedonic utility from consumption of the good, we should observe sellers taking advantage of this phenomenon. This potentially explains the growing use of debit cards (when they do not offer significant advantages over a credit card) because by paying upfront (as in the vacation example), individuals can ensure that their mental account does not go into red. Another possible example is the use of tokens in casinos. Whilst the act of buying a token gives hedonic pain, on account of parting with the money, the pain is cushioned by thoughts of future benefits as in (20.21). Furthermore, once the payment is in the past, the consumer gets a higher hedonic utility while actually playing in the casino (prospective accounting). This induces the individual to pay upfront for the casino tokens. A similar reasoning might apply to individuals who go on holiday to a foreign country and report that things are much cheaper abroad even when they are not. Once they have paid for the holiday and changed into the foreign currency, prospective accounting ensures that they can enjoy the holiday as if it were free.
3. **Popularity of fixed fee tariffs:** Perhaps the desire for prepayment also explains the popularity of fixed fee payment schemes such as unlimited Internet access, once an initial fee is paid. Telephone companies also engage in similar pricing schemes. Prelec and Loewenstein (1998) report that in one study, about 65% of consumers selected a flat rate telephone payment scheme when an alternative variable rate scheme would have saved them money. In contrast, only 10% of those who choose the variable rate payment option could have saved money by switching to the flat rate scheme. The mental accounts explanation would be that talking on the phone is less pleasurable under a variable rate scheme because the hedonic benefits of a currently enjoyable conversation are tempered by thoughts of future payment (see (20.20)). For a similar reason, a taxi ride is perhaps made less pleasant by the sight of the taxi meter ticking away. All these are examples of *payment decoupling* due to prospective accounting.
4. **Club Med vacations and gifts:** An increase in the coupling parameter α is bad news for the purchase decision because it reduces the hedonic benefits; see (20.20). Hence, it would seem to be in the interests of sellers to decouple benefits from costs. So, for instance, the Club Med vacation covers several activities (travel, meals, lodging and recreation) under a single fee so that the consumer feels that individual activities are uncoupled from the fee; see Thaler (1980). Another interesting application is to gift giving. Why do couples with joint accounts give gifts to one another? Perhaps the gift uncouples the payment from the

benefits for the receiver. The gift item might actually turn out to provide significant hedonic utility to the receiver who in the absence of the uncoupling might have agonized over its purchase.

20.6 Choice bracketing

In neoclassical economics, when faced with the problem of making a choice from a particular domain, individuals consider all available information and all available choices. *Choice bracketing* occurs when individuals group the available choices into sets, and view these sets in isolation from one another, or do not take into account the set of all possible choices. If not all choices are considered, or all available information is not used, then an individual is said to be engaged in *narrow choice bracketing*. By contrast, in neoclassical economics, individuals might be said to be engaged in *broad choice bracketing*.

In a dynamic setting, narrow choice bracketing occurs when individuals take actions in any time period by only considering the costs and benefits in that period, rather than optimally choosing the entire intertemporal profile of actions. A wide range of behaviors seems consistent with some form of choice bracketing. Our treatment in this section is based on Read et al. (1999). Consider the following examples under certainty.

Example 20.21 (*War strategy*): Consider an army general who is engaged in fighting a war that consists of a sequence of battles. In theory, one expects that the general is concerned about winning the war. However, if the general engages in narrow choice bracketing, he might give undue importance to winning particular battles, with the possibly unpleasant result of losing the war. So, for instance, in a battle that involves securing a particularly strategic position, such as high ground, a consideration of the battle alone (narrow choice bracketing) might dictate retreating because of the heavy losses involved. However, that could have a serious consequence for subsequent battles and for winning the war.

Example 20.22 (*Smoking*): Consider a smoker who brackets choices narrowly. Thus, the decision to smoke each cigarette is taken in terms of its current costs and benefits while ignoring the effect on future utilities. Because the future costs are underweighted in this decision, narrow choice bracketing is likely to promote smoking. However, under broad bracketing, the individual must consider that at the rate of a packet a day, such a strategy could lead her to smoke 7300 cigarettes a year. So, under broad bracketing, the individual might decide not to smoke. The choice of narrow or broad bracketing might also involve issues of emotions, temptation, and addiction, over which the individual could have limited self-control.

Example 20.23 (*New York cab drivers*): Camerer et. al. (1997) found that New York cab drivers have a particular earnings target for the day. This is a form of narrow bracketing because the cab drivers choose a very short horizon. On a busy day, therefore, they might quit earlier, generating a backward bending supply curve of labor. If they were to choose broad bracketing, say, in the form of maximizing earnings over a year, then they could presumably earn more; this appears to be the case for more experienced drivers. On the other hand, the gain from narrow bracketing might be an improvement in personal motivation (see below).

Although choice bracketing has been largely ignored in economics, models in economic theory often implicitly assume and invoke it. So, for instance, a consumer's risk aversion is often defined over her aggregate wealth. However, several kinds of insurance contracts deal with only very small

changes in wealth (such as an appliance breaking down). Individuals would have to be absurdly risk averse in some cases to buy such insurance. Hence, the presumption must be that individuals look at the cost and benefits of such actions in isolation, which is a form of narrow bracketing.

Under uncertainty, choice bracketing leads to other interesting effects as the following two examples show.

Example 20.24 (*Equity premium puzzle*): Suppose that individuals can invest either in bonds, or in equity. Consider two individuals who use prospect theory to evaluate risky choices: Mr. Anxious and Mr. Relaxed. The former checks portfolio returns far more regularly than the latter. Equity returns fluctuate more than bonds, and in any period, are as likely to go up as down. On account of loss aversion, observing equity returns will be stressful for both types of individuals, but particularly for Mr. Anxious who checks the portfolio more frequently. Hence, Mr. Anxious will hold stocks only if the return on equity is relatively higher. Within a temporal choice bracketing context, one description of Mr. Anxious is that he engages in narrow choice framing. Rather than being affected by the longer-run return on the equity, he is too sensitive to frequent fluctuations. If there are a large number of Mr. Anxious types in the population, then for equity to be held, arbitrage will ensure that there is a premium on equities, relative to the normal case of broad bracketing (Benartzi and Thaler, 1995).

Example 20.25 (*Risk aggregation*): Sometimes several gambles might be accepted if they are bracketed broadly, but not when they are bracketed narrowly. In a famous example, the Nobel laureate Paul Samuelson asked a colleague if he would accept the following gamble: win \$200 or lose \$100 with equal probability (Samuelson, 1963). The colleague replied that he would reject this bet but would accept 100 such bets.

In fact, if the colleague happens to be a prospect theory maximizer, he would prefer even two such bets to a single bet, as we now show. Suppose that the prizes (\$200 and \$100, respectively) reflect outcomes relative to the reference point. Consider the utility function in (20.1) with the empirically supported parameter values $\beta = 1$ and $\lambda = 2.25$ (see Part 1 of the book for details). Assuming linear probability weighting, the value function under prospect theory for a single bet is

$$V_1 = \frac{1}{2}(200) - \frac{1}{2}2.25(100) = -12.5 < 0. \quad (20.22)$$

When there are two bets, the following sequence of outcomes is equally possible: (200, 200), (200, -100), (-100, 200), (-100, -100). Suppose that if the individual accepts the two-period bet, he derives utility from the sum of undiscounted payoffs over the two periods. This is the sense in which he brackets broadly; narrow bracketing would take the form of treating the two separate bets as one-period bets; since the bets are identical, using (20.22) we know that $V_1 + V_1 < 0$. Given our assumptions, under broad bracketing, the four possible equally likely outcomes are 400, 100, 100, -200; the individual is in the domain of gains in the first three cases and in the domain of losses in the last case. The value function in this case is given by

$$V_2 = \frac{1}{4}(400) + \frac{2}{4}(100) - \frac{1}{4}2.25(200) = 37.5 > 0.$$

Thus, if the individual brackets broadly and aggregates the bets over the two periods, then he accepts ($V_2 > 0$). However, if he were to use narrow bracketing, and treat the two bets separately, then he rejects ($V_1 + V_1 < 0$).

Rabin and Weizsäcker (2009) apply narrow bracketing to explain the experimental findings in a problem in Tversky and Kahneman (1981). In this problem, individuals are given two sets of choices with two options within each set. Individuals are found to choose in a manner that is consistent with prospect theory: Choose the safe option in the domain of gains and the risky option in the domain of losses. However, when the two chosen options are combined, they are first order stochastically dominated by the combination of the two unchosen options. This behavior can be explained by narrow bracketing, which makes it acceptable for subjects to have missed out the stochastic dominance in the combined choices (broad bracketing).

The reason for why people engage in narrow bracketing is not well known. Read et al. (1999) conjecture the following possibilities. One possibility is that cognitive limitations preclude a consideration of the entire domain of choice. However, there is some experimental evidence that the heuristic of narrow bracketing arises even in the absence of cognitive limitations. Yet another possibility is that choice bracketing could be acquired by social norms and conventions. For example, we often divide the week into two unequal lengths: the working week and the weekend, and we label food as meals or snacks.

A particularly important justification for narrow bracketing is that it might reduce or overcome problems associated with self-control. Alcoholics under treatment are usually asked to take one day at a time (narrow bracketing) because the possibility of lifetime abstinence (broad bracketing) is too daunting. Narrow bracketing, in this case, improves motivation. Successful batsmen in the game of Test cricket, when asked about their batting strategy often say that they take one session of play at a time (narrow bracketing) rather than planning to bat out for the entire day. Or a tennis player may focus on winning one set/game at a time to keep motivation at a high level. Motivation might also explain the behavior of New York cab drivers mentioned above. By setting an attainable target over a short period of time, they are less likely to be daunted by the prospect of achieving a more ambitious target, say, a yearly income target.

Broad bracketing, under the typical assumptions in neoclassical economics, may be expected to achieve the global maximum. However, in the absence of these assumptions, narrow bracketing has certain advantages. First, broad bracketing might create relatively greater anxiety because targets and goals that it prescribes are likely to be more ambitious, so possibly more daunting as well. Second, and related to the first, narrow bracketing might be a better motivational device. Third, under broad bracketing, people need to take account of a larger set of factors that is more cognitively demanding, a factor that is important when relatively quick decisions need to be made, or when higher cognitive loads render the underlying problem too difficult/impossible to solve. Thus, in this case broad bracketing might lead to procrastination. Insofar as individuals make errors in computation and prediction, the consideration of a larger set of factors might also magnify errors. There are also other cognitive costs of engaging in broad bracketing, such as greater record keeping.

20.7 Coherent arbitrariness

Neoclassical economic theory rests on the foundation that choice reflects well-defined preferences. So, for instance, when consumers have preferences defined over consumption bundles, then, under certain conditions, the consumer has a downward sloping demand curve. Now suppose that we start with a downward sloping demand curve. Can we infer that consumers have well-defined initial preferences?

Ariely et al. (2003) consider this issue and advance the concept of *coherent arbitrariness*. They find that initial preferences of individuals might well be arbitrary. However, once arbitrary preferences get entrenched, then, economic agents behave in a coherent manner with respect to these preferences. So, for instance, a consumer's initial valuation for a product might differ noticeably from the fundamental value (*arbitrariness*). However, if the quality of the product is subsequently improved, consumers are willing to pay more, just as economic theory predicts (*coherence*).

The idea that preferences are not well defined initially, but become more clear as the decision process unfolds has also been suggested in other contexts. See, in particular, the idea of *constructed preferences* in Slovic (1995). In other contexts, it has been shown that preferences are quite malleable. For instance, in experiments on the elicitation of risk preferences conducted before and after the 2007 financial crash, the post-crash preferences are more risk averse (Gong et al., 2013).

So how do consumers *anchor* their initial preferences? The arbitrariness of the initial anchor can be illustrated by the following well-known experiment (see Section 19.6 in Chapter 19). Tversky and Kahneman (1974) found that the answers of subjects in experiments to the question of the number of African countries in the UN depended on the prior random outcome of a roulette wheel (arbitrary anchor). Similar results have been replicated in other contexts.⁵

The implication of these results is that consumers may not approach novel situations with well-defined, preexisting, preferences, nor with well-defined subjective probability distributions over outcomes. Hence, initial impressions and initial accidents of history are important anchors. So, for instance, the marketing decisions associated with the launch of a new product are often critical determinants of how much consumers are willing to pay for the product. Since marketing personnel are likely to be aware of these issues, there is a potential role for government regulation to ensure that accurate and transparent information is transmitted to consumers in the first instance.

To test the arbitrariness of anchors for initial preferences, Ariely et al. (2003) needed consumers, in this case MIT students, to value a product that they did not have any market experience of. They came up with the willingness to accept a payment (WTA) in return for the experience of listening to an unpleasant sound. Three groups of students were first asked to listen to a sample of unpleasant noise for 30 seconds. The first group was offered a payment of 10 cents (the low anchor group), the second group was offered a payment of 50 cents (the high anchor group), and the third group did not receive any anchor information. Thus, the initial payment was used as the (arbitrary) anchor. Then, they were asked if they would be willing to hear the same noise in return for a payment.

The WTA of the three groups of students was then elicited for increasing (10 sec, 30 sec, and 60 sec) and decreasing (60 sec, 30 sec, and 10 sec) durations of sounds; each sequence was repeated over three rounds. The results, for each sequence and for each of the three rounds, are shown in Figure 20.2. There are two main results. First, there is an almost linear relation between WTA and the duration of the sound (coherence). Second, the WTA was significantly higher in the high anchor group relative to the low anchor group and the no anchor group; indeed the high anchor group earned less money in the experiment, relative to the other two groups. There was no convergence of WTA among the high and low anchor groups. Repeated experience with the experiment did not override the anchoring effect.

One confounding factor in this experiment may arise if subjects infer that the anchors represent some risk to their hearing. Hence, the experiment was repeated with another group of

⁵ Green et. al. (1998) and Kahnemann and Knetsch (1993) find that a higher (lower) initial anchor elicits a higher (lower) willingness to pay for a public good. For other examples, see Section 19.6 in Chapter 19.




Figure 20.2 The WTA of the three groups of students for increasing and decreasing durations of sounds.

Source: Dan Ariely, George Loewenstein, and Drazen Prelec, "Coherent arbitrariness: stable demand curves without stable preference." *The Quarterly Journal of Economics* (2003) 118(1): 73–106, by permission of Oxford University Press.

students and the initial anchor used was the first three digits of the subject's own social security number. Also, in this experiment, the stakes were raised by a factor of 10. It turned out that subjects with smaller social security numbers required much lower payments, hence, it is anchoring that is at work. As in the previous experiment, subjects exhibited coherence by demanding higher payments for exposure to noise of a longer duration.

To test the results in a concrete economic environment, in yet another experiment, the subjects are initially exposed to an arbitrary anchor. Then their WTA was elicited in an auction. The idea is for subjects to listen to other bids and adjust their WTA in light of this market information. The high and the low anchors in this case were \$1 and 10 cents, respectively. The anchoring results survive. The mean WTA for the high (low) anchor group was 59 cents (8 cents). The WTA for the two anchor groups does not converge, although the variance of bids within each group reduces as the experiment is repeated.

Fudenberg et al. (2012) find, unusually, very weak anchoring effects for goods, and an absence of anchoring effects for lotteries. However, anchoring is one of the most robust findings in the heuristics literature (see Chapter 19). In general, the results survive with multiple anchors and non-monetary rewards. The role of anchors in determining subsequent choices can potentially explain several puzzles. For instance, why asset prices might have an arbitrary component that is unrelated to fundamentals? Anchoring might also explain the findings of Bewley (1998) that reducing payment to workers reduces morale but workers are less concerned about their wages relative to wages elsewhere; in this case, the anchor is the wage in one's own firm. This also potentially explains why inter-industry wage differentials can persist; see for example Thaler (1989).

20.8 Sunk costs and mental accounting

Neoclassical economics is based on a *marginalist principle*: Individuals are guided by the marginal costs and the marginal benefits of their actions. In particular, the fixed or sunk costs that might

have been incurred previously are irrelevant. However, experimental evidence indicates that sunk costs do have an important influence on the actual action taken.

Kahneman and Tversky (1984) report an experiment in which subjects in the experiments are told that they purchased a theater ticket for \$10 in advance (sunk cost) but subsequently lose it before the event. These subjects were less willing to buy another ticket (46%) as compared to a second set of subjects who were told that they had lost an equivalent sum of money (88%). Perhaps, consumers create a mental account on “theater outings” and losing the ticket is seen as an outflow from this account, while losing an equivalent amount of money leaves this account unaffected.

Suppose that you purchased a ticket for \$100 to see a basketball game, but there is a blizzard on the day of the game. Will you go to see the game? The cost of the ticket is a sunk cost and should not influence the decision to travel in a blizzard. However, individuals are more willing to travel in the blizzard if they already hold the ticket (Thaler, 1980). The explanation is that individuals have a mental account for the basketball game. It is debited for the \$100 when the ticket is purchased. If the game is missed, then the account must be closed in the red, which is aversive to the ticket holder, but not to a non-ticket holder. There is a cost of travelling to the game. However, Heath (1995) suggests that since these costs are not directly monetary, they are not included in the relevant mental account.

Sunk costs have real effects on decisions. However, the effect of sunk costs can wear off over time; this is known as *payment depreciation*. Thaler (1999) gives the following example. Suppose that you buy shoes that seem comfortable at the time of buying, but after wearing them for a day they feel uncomfortable. You try wearing them for the next few days but it gets worse. At this point, the shoes cannot be returned because of the wear and tear on them. If the shoes cannot be returned, then their purchase is a sunk cost. Thaler then makes the following plausible conjectures, which illustrate important mental accounting principles.

1. If the shoes are more expensive, you will try to wear them more often, despite the discomfort.
2. You will eventually stop wearing the shoes, but the more expensive are the shoes, the longer they will sit in the closet before being eventually thrown away.
3. Irrespective of how expensive the shoes, they are eventually thrown away.

The individual's behavior in 1 and 2 above is an example of the influence of *sunk costs*. The behavior in 3 is an example of *payment depreciation*, i.e., eventually people ignore sunk costs. See Arkes and Blumer (1985) for evidence on how the effect of sunk costs can persist, but the effect vanishes eventually. Gourville and Soman (1998) give evidence on payment depreciation in the context of payments for health club memberships. Thus, once the (depreciated) sunk costs fall below some lower limit, say, $\underline{c}$, the mental account is closed.

One could make use of the concepts of *acquisition utility* and *transaction utility* (see Section 20.9.1) to understand this example better. Suppose that the individual uses transaction utility to make the purchase decision, but then evaluates the subsequent experience in terms of acquisition utility. Since the shoe turns out to be smaller than expected, the ex-post valuation of the shoe, v , is lower. To consider an extreme case, suppose that ex-post, $v = 0$, so that the acquisition utility is $-p$ per period. The mental account corresponding to this purchase is, thus, in the red, equal to the dollar amount p . If losses depreciate continuously at the rate δ each period, then the mental account is closed when $pe^{-\delta t} = \underline{c}$, where t is the time period at which the equality is satisfied. It is clear that the size of t is increasing in the absolute magnitude of the loss, p . Hence, more expensive (but useless) shoes are kept in the closet longer before being thrown away.

Shafir and Thaler (2006) asked the following question to subscribers of a wine newsletter.⁶

Example 20.26 “Suppose you bought a case of a good 1982 Bordeaux in the futures market for \$20 a bottle. The wine now sells at auction for about \$75 a bottle. You have decided to drink a bottle. Which of the following best captures your feeling of the cost to you of drinking this bottle?” The respondents were given five answers to choose from: \$0, \$20, \$20 plus interest, \$75, and –\$55. The percentages of respondents choosing each answer was 30, 18, 7, 20 and 25.

This is a question on opportunity cost that lies at the very heart of any introductory course in economics. Based on an opportunity cost calculation, the correct answer is \$75, and the sunk cost of \$20 is irrelevant. However, more than half the respondents report that drinking the bottle either costs nothing or actually saves them money. There is a wide range of answers among these seasoned wine drinkers. When the experiment was posed in terms of giving the bottle away (rather than drinking it), then 30% gave the correct answer, \$75. The authors conclude that wine drinkers think of their initial purchase as an investment. People who choose the –\$55 option, for instance, think that the value of the investment turned out to be \$75, but they paid \$20, so they saved \$55. Thus, more than half the responders think that drinking the bottle either costs them nothing, or actually constitutes a saving.

Shafir and Thaler (2006) follow up this experiment with the following.

Example 20.27 “Suppose you buy a case of Bordeaux futures at \$400 a case. The wine will retail at about \$500 a case when it is shipped. You do not intend to start drinking this wine for a decade. At the time that you acquire this wine which statement more accurately captures your feelings?” (a) I feel like I just spent \$400, much as I would feel if I spent \$400 on a weekend getaway. (b) I feel like I made a \$400 investment which I will gradually consume after a period of years. (c) I feel like I just saved \$100, the difference between what the futures cost and what the wine will sell for when delivered.

Participants were asked to give a score from 1 (strongly agree) to 5 (strongly disagree) to each of the three options. The mean scores for the three choices were 3.31, 1.94, 2.88 (the lowest score indicates the strongest agreement). Hence, responders gave the maximum approval to option (b), the investment option, rather than the saving option (c).

Combining the results of the two experiments, the authors conclude the following. The wine connoisseurs think of the initial purchase as an investment. At a later date, the wine is drunk as if it were free. This is similar to the feature of *prospective accounting* in Prelec and Loewenstein (1998) that we consider in Section 20.5. This explains the title of the Shafir–Thaler paper: *Invest Now, Drink Later, Spend Never*. Furthermore, a large proportion of individuals (55%) start thinking in terms of the opportunity cost (\$75) if the question is posed in terms of the wine bottle having been broken accidentally. Indeed, breaking the bottle makes the replacement cost of the bottle more salient in a way that giving away the bottle does not.

Does the form of sunk costs matter? In particular, does it matter if sunk costs are expressed in terms of dollar value or in terms of time? The experimental evidence in Soman (2001) suggests that *sunk-time costs* have a smaller effect on actual actions relative to *sunk-dollar costs*. He gives three main distinguishing factors that are the source of the difference. First, time is perceived as

⁶ The wine newsletter in question, *Liquid Assets*, used to be published by the economist Orley Ashenfelter. The subscribers are extremely knowledgeable wine consumers who have substantial home cellars. In many cases, these are economists or business executives so they should, in principle, be able to fully appreciate the choices in the experiment.

a flow variable rather than a stock variable, which makes it difficult for an individual to make an inventory of the time expended on various activities. Second, in contrast to dollar costs, the disutility from expending time might not be directly proportional to the actual time spent. This disutility depends, for instance, on how enjoyable the activities undertaken during this time were, and how productively the time was spent. Finally, most people, particularly students, may have better training in accounting for money as compared to time. Navarro and Fantino (2009), however, find that sunk-time costs are also important, but that they can be moderated. There is also some evidence to suggest that individuals' willingness to pay is higher for risky options that involve time rather than money (Okada and Hoch, 2004).

Rajagopal and Rha (2009) ask if individuals have mental accounts for time, in a similar manner to the mental accounts that they have for money. They postulate two mental accounts for time: work and non-work. In experiments, they find that individuals face different degrees of anguish if they are delayed for work or non-work activities, suggesting that time is not fungible across these two mental accounts.

20.9 Some other mental accounting phenomena

In this section we consider an assortment of mental accounting phenomena.

20.9.1 Transaction utility and acquisition utility

Suppose that an individual has a valuation for an object, $v > 0$, and pays a price p to acquire it. Then, the net valuation of the object, or the *acquisition utility*, U_A , is

$$U_A(v, p) = v - p. \quad (20.23)$$

Acquisition utility corresponds to consumer surplus. A related concept is *transaction utility*, U_T , defined as follows (Thaler, 2008).⁷

$$U_T(\bar{p}, p) = \bar{p} - p, \quad (20.24)$$

where $\bar{p}$ is the reference price or the regular price that the consumer expects to pay. Thaler (2008) feels that the most important determinant of $\bar{p}$ is fairness, which depends on the costs of the seller. This is supported by the experimental evidence (Thaler, 2008, Table 1, p. 19). Transaction utility would dictate buying the good only if the actual price is below the reference price, i.e., $\bar{p} \geq p$. The total utility from the purchase is $U(v, \bar{p}, p) = U_A + U_T$.

Example 20.28 (Thaler, 1985, p. 189): *You are lying on the beach on a hot day. All you have to drink is ice water. For the last hour you have been thinking about how much you would enjoy a nice cold bottle of your favorite brand of beer. A companion gets up to go make a phone call and offers to bring back a beer from the only nearby place where beer is sold (a fancy resort hotel) [a small, run-down grocery store]. He says that the beer might be expensive and so asks how much you are willing to pay for the beer. He says that he will buy the beer if it costs as much or less than the price you state. But if it costs more than the price you state*

⁷ Thaler (2008) is a reprint of Thaler (1985).

he will not buy it. You trust your friend, and there is no possibility of bargaining with the (bartender) [store owner]. What price do you tell him?

The median willingness to pay for the beer from the resort and the grocery store was, respectively, \$2.65 and \$1.50 in 1984 prices. To rationalize this, suppose that the decision criterion of the consumer is to purchase the object if $U_T \geq 0$. Hence, if the consumer has a higher reference price, $\bar{p}$, for the resort relative to the grocery shop, then he would also be willing to pay a higher actual price, p . These results are quite striking because in each case, the object of consumption is an identical bottle of beer in an identical environment and there are no strategic considerations.

This would also appear to explain why we sometimes buy items at a sale that we do not find very useful subsequently. The reason is that transaction utility is salient at the time of a purchase. This rule, namely, over-purchase when the actual price is lower than the reference price can also be used to explain the puzzle in Example 20.13, where the individual buys a king size quilt when a queen size quilt would have been more appropriate.

20.9.2 *When to close an account?*

Individuals find it more difficult to close an account that is in the red. Consider a portfolio that has two assets. One of the assets has been going up in value, while the other has been going down in value. The gains and losses can either be *accrued* (so long as the stock is unsold), or they can be *realized* (if the stock is sold). Realized gains and losses seem more painful relative to accrued gains and losses (Thaler, 1999). If individuals have a mental account for returns on the stock and they find it aversive to close it in the red, then they are reluctant to sell losers. So if they have to raise some money they are more likely to sell winners. In contrast, the neoclassical prediction would be to sell losers because tax laws typically allow realized losses to be offset against other gains, allowing one to reduce the overall tax owed (Shefrin and Statman, 1985). Empirical evidence shows that people do indeed sell winners and hold on to loss making stock for too long—this is known as the disposition effect (Odean, 1998).

20.9.3 *Self-control and gift giving*

Let us take as given that individuals construct mental accounts that deal with distinct items of consumption. It is observed that people often set *low* budgets for some accounts. For instance, an individual who is worried about excessive wine consumption might cap his “wine mental account” at some upper limit, say, $\$a$. Suppose that a wine comes up on the market at a price $\$b > \$a$ for which the individual’s willingness to pay is $\$c > \$b$. In the absence of mental accounts, it is optimal to buy the wine. However, with a capped mental account, it is optimal not to buy. It is quite possible that a gift of the wine to this individual will be better appreciated relative to receiving an equivalent sum of money. This contrasts with the well-known result that, in the presence of convex indifference curves, equivalent sums of money must always be preferred to payments in kind.

Individuals might wish to cap expenditures on bads or sin goods, such as cigarettes and alcohol, because they realize that they may have self-control problems in the future. One form of expenditure capping for bads is an increased demand for smaller packets of the sin goods. The demand for such items, packaged in small amounts, can potentially create a premium above the price at which they are normally sold; this was found to be true for cigarette packets (Werthenbroch, 1998).

AN EXAMPLE: THE FOOD STAMP PROGRAM

Consider the Food Stamp Program in the US, which is the cornerstone of the *Federal food assistance programs*. The Food Stamp Program enables low-income families to buy nutritious food with coupons and Electronic Benefits Transfer (EBT) cards. Food stamp recipients spend their coupons to buy eligible food from authorized retail food stores. Neoclassical economic analysis predicts that government assistance in the form of equivalent cash benefits leads to higher utility for consumers, relative to in-kind assistance such as food coupons. We demonstrate the reverse to be true when the individual engages in mental accounting.

Suppose that a consumer derives utility from two goods only: food (Y) and all other goods lumped together as a composite good (X). The individual has a fixed, and low, amount of income, $I > 0$. The price per unit of X is normalized to unity, while the price of Y per unit is $p > 0$. The consumer's problem in the absence of any government intervention is to choose X, Y , to maximize a utility function $U(X, Y)$ that is increasing, concave, and twice continuously differentiable, subject to two constraints. The first constraint is the usual budget constraint that constrains expenditure to be no greater than income.

$$X + pY = I. \quad (20.25)$$

The second constraint is a mental accounting constraint. The individual has a system of mental accounts such that he evaluates expenditures on food arising from his own income in a separate account. The individual caps the expenditure on the food account *from his own income* at some number $\$c > 0$.

$$pY \leq c \Leftrightarrow Y \leq \frac{c}{p} \equiv y. \quad (20.26)$$

If the individual is given some *gift* of food, say, $\hat{y}$, then he is willing to consume up to an amount $\hat{y} + y$, but never buys more than y units from his own income. We assume, for simplicity of analysis, that the mental accounting constraint (20.26) is invariant with respect to changes in income or government policy. The individual's indifference curves in (food, non-food) space are very flat, hence, the individual buys too little food. We also make the mental accounting assumption that the individual treats all gifts of income as his own income.

In Figure 20.3, the vertical line at y indicates the maximum food that the individual is willing to buy with own income. The budget constraint in (20.25) is shown by the line AC . The intersection of the two constraints gives the convex choice set AB . Since the price of X is normalized to unity, the height of point A also gives I . The optimum is at B (an indifference curve is shown passing through it).

Direct Income Transfer Scheme: Suppose that the government provides assistance to the poor directly through an *income transfer scheme*. The individual's income, inclusive of the additional income, is now $\bar{I} : \bar{I} > I$. The new budget constraint under the *income transfer scheme* is given by

$$X + pY = \bar{I}, \quad (20.27)$$

and is shown by the line DH in Figure 20.3. Given our assumptions, the constraint (20.26) is unchanged. The intersection of the two constraints gives the convex choice set DF . The consumer's new optimum is at F .




Figure 20.3 The relative optimality of food stamps and direct cash benefits.

Food Stamp Program: Suppose now that the government undertakes a *food stamps program* and provides food stamps equal to F_s , expressed in units of Y , to the household, as shown in Figure 20.3. To ensure comparability of the two alternative government programs, we assume that the nominal transfer of income is the same in each case, i.e.,

$$pF_s = \bar{I} - I. \quad (20.28)$$

The budget constraint under the food stamp program is given by

$$X + p(Y - F_s) = I. \quad (20.29)$$

Given our assumptions, with the introduction of the food stamp program, the individual is willing to buy a maximum amount $F_s + y$ of food. Substituting (20.28) in (20.29) we get that $X + pY = \bar{I}$, which is identical to the budget constraint in the case of the direct income transfer scheme. Thus, under the food stamp program, the choice set of the individual is given by the convex choice set EGJ .

The optimum in the case of the food stamp program at point G , is on a higher indifference curve relative to the optimum for the income support program (at point F). Hence, under mental accounting, in-kind transfers may dominate cash transfer. Although our example is contrived to directly show the efficacy of the food stamp program under mental accounting, it highlights a more fundamental idea. If policymakers can understand the mental accounting system used by the poor, then policy intervention may lead to a Pareto improvement; the World Bank Development Report, 2015, highlights several such channels.

20.9.4 The house money effect

Thaler and Johnson (1990) extend the insights from static loss aversion in prospect theory to a dynamic context. Subjects in their experiments were given the following three problems. The percentage of subjects making each choice is given in square brackets.

Problem 1. You have just won \$30. Now choose between:

(a) A 50% chance to gain \$9 and a 50% chance to lose \$9. [70]

(b) No further gain or loss. [30]

Problem 2. You have just lost \$30. Now choose between:

(a) A 50% chance to gain \$9 and a 50% chance to lose \$9. [40]

(b) No further gain or loss. [60]

Problem 3. You have just lost \$30. Now choose between:

(a) A 33% chance to gain \$30 and a 67% chance to gain nothing. [60]

(b) A sure \$10. [40]

In Problem 1, subjects have made a prior gain. This induces more than two thirds of them to subsequently choose the risky option, a phenomenon known as the *house money effect*; gamblers refer to money they have won in a casino as *house money*. Furthermore, they are likely to put their winnings in a separate mental account from their own money (Thaler, 1999). Hence, a prior gain can induce risk seeking. Problems 2, 3 illustrate the ambivalent nature of a prior loss. Problem 2 shows the converse result from Problem 1; a prior loss can induce more conservative behavior. However, Problem 3 shows that a prior loss can also stimulate risk seeking. The difference between Problems 2 and 3 is that option (a) in Problem 3 gives a chance to break even, while this is not the case in Problem 2.

The house money effect, namely, that current attitudes to risk are conditioned by previous gains and losses has several applications. Barberis and Huang (2001) apply this idea to the following puzzle in finance. At the level of an individual stock, returns over time have a high mean, are highly volatile, and lagged returns have some predictive ability to forecast future returns. However, the cross-sectional evidence is that stocks whose price relative to fundamentals is low have higher than average returns.

Barberis and Huang postulate that the utility of individuals is influenced not only by the final levels of consumption, but also by the gains and losses in the value of individual stocks. Furthermore, the evaluation of current gains and losses depends on previous gains and losses. Thus, for instance, if a stock, say stock A, has been performing well in the recent past, investors become less concerned about future losses on the stock. In other words, previous gains offset (in a psychic sense) future losses. An implication is that investors view the stock as less risky. Conversely, if a stock, say stock B, has been making losses in the recent past or not doing too well, then it is perceived as more risky. This, in turn, induces perceived changes such that possible future losses from the stock become relatively more salient. For the individual to be willing to hold stock B, its returns must increase in equilibrium. Hence, stocks fluctuate in value, below-par performance being followed by above-par performance.

Furthermore, the more the return fluctuates, the greater the psychic costs on the loss averse individual, so the greater must be the average expected return on the asset for it to be held in the portfolio. When investors in Barberis and Huang (2001) derive utility from changes in the value of individual stocks, they engage in a form of *narrow bracketing*. An alternative, which is considered by the authors, is for investors to focus on the changes in the aggregate value of their portfolios as, for instance, would be recommended under *broad bracketing*. However, empirical evidence in this case is more supportive of the narrow bracketing approach. Information on the gains and losses of individual stocks is more readily available relative to aggregate information on various portfolios.

Rockenbach (2004) reports on experimental evidence whereby investors keep separate mental accounts for risky and riskless assets. Investors exhibit features of loss aversion when they react to changes in the value of assets in the risky portfolio. However, the presence of separate mental

accounts hinders the exploitation of arbitrage opportunities across the different accounts, especially if money is not fungible across the accounts. This multi-layered portfolio approach supports the behavioral finance theory of Shefrin and Statman (2000).

20.9.5 *The diversification heuristic*

The *diversification heuristic* exemplifies the human tendency to choose too much variety when presented with a larger set of alternatives, *each of which can be simultaneously chosen*. The situation may be quite different when only one alternative can be chosen, in which case a greater number of choices may result in a choice impasse. The term diversification heuristic was used by Read and Loewenstein (1995) and Benartzi and Thaler (2001) credit Simonson (1990) with the first demonstration of this heuristic. Some diversification is always optimal and lack of diversification is most likely to be suboptimal. However, in the presence of a sufficient degree of complexity in the choices, the observed extent and nature of diversification is often quite naive and unlikely to be the outcome of a neoclassical optimization exercise.

In an extreme form of the diversification heuristic, when faced with many different options, subjects may choose each option equally. Consider the following quote from Benartzi and Thaler (2001, p. 79): “Writing in about the fourth century, a Rabbi Issac bar Aha gave the following asset allocation advice: ‘A man should always place his money, a third into land, a third into merchandise, and keep a third at hand.’ There is anecdotal evidence that this rule is still in use.”

The growth in defined pension contribution plans raises the issue of the financial sophistication of individual decision makers to optimally allocate funds in their portfolios. Benartzi and Thaler (2001) ask employees to make a choice among alternative retirement funds. They find that individuals prefer to diversify by using simple heuristics. One of the retirement funds offered one bond option and one stock option. On average, half the portfolio was allocated to the stock. But if another stock was added to the possible portfolio, then the average allocation towards each stock was a third. This motivates the authors to propose the idea of a $\frac{1}{n}$ heuristic: When individuals have n funds to choose from, they divide their portfolio equally among the alternatives. While some diversification is generally good, here the extent of diversification is naive.

Picking from funds that contain predominantly stocks, an individual who uses the $\frac{1}{n}$ heuristic will load his portfolio towards stocks. Likewise, picking from a fund that contains predominantly bonds, the portfolio will be loaded towards bonds. Benartzi and Thaler (2001) confirm this by comparing two different actual pension plans. In one plan, the TWA plan, the choices comprise of five core stocks and one core bond. The average portfolio chosen in this case has 75% stocks. Another pension plan, the University of California plan, comprises of four core bonds and one core stock. Here, the average portfolio has 34% stocks, well below the national average of 57% stocks in retirement portfolios.

Other supporting evidence from the US shows that individuals invest excessively in the stocks of their own companies (Poterba, 2003; Mitchell and Utkus, 2004). In a related problem, individuals invest too much in home stocks that they are more familiar with, forgoing the diversification benefits from investing in foreign stock (French and Poterba, 1991; Cooper and Kaplanis, 1994; Lewis, 1999; Solnik and Zuo, 2011). However, there is some evidence that investors are improving the extent of their diversification (French, 2008).

Read and Loewenstein (1995) focus on the behavior of trick-or-treaters on Halloween night as they visited different houses.⁸ In the first treatment, the sequential choice condition, the children

⁸ This is one of four different experiments that supports the diversification heuristic in their paper.

were offered a choice between two candies, Three Musketeers and Milky Way, in two consecutive houses. In the second treatment they were told at the first house they visited to “choose whichever two candy bars you like” from a large pile of both types of candies. In the first treatment, 52% of the children picked up two pieces of their favorite candy at the two houses and 48% picked different candies. In the second treatment, they showed a strong desire to diversify and all children picked each candy of a different type.

20.9.6 *The mental accounting of goals*

Individuals often set a variety of goals for themselves. Goals can be very explicit (all or nothing goals). For instance, losing x kilos of weight over t time periods. Or goals might be less well-specified (graded goals). For instance, a simple target of reducing weight over time. Other examples of goals include writing term papers, completing referee reports, cleaning one's room, quitting an addiction such as smoking, and resuming physical activity after surgery.

Evidence shows that setting goals improves performance (Locke and Latham, 1990; Orbell and Sheeran, 2000). Two important factors that are conducive to an improvement in performance are whether the goal is well-specified (as opposed to vaguely specified goals such as “do your best”) and “the level of difficulty” of the goal. The evidence on the positive effects of goals is considered in Part 1 of the book.

However, goals that are too difficult may be less helpful for the individual; for a theoretical analysis of goals, see Part 1 of the book. Soman and Cheema (2004) ask if such goals are counterproductive. In particular, they consider cases where individuals miss their goal, and the effect of a failure to achieve the goal on subsequent behavior. They look at two well-specified goals: a saving's goal and a task completion goal that imposes a deadline. Soman and Cheema point out the following main drawbacks of goals.

First, on account of mental accounting, individuals might code outcomes as gains or losses. Missing a goal is coded as a loss, so it reduces utility much more sharply. This is particularly true of the all-or-nothing goals where, ex-post, one cannot rationalize missing a goal. However, for graded goals, one can rationalize one's behavior better, and possibly reduce the set of outcomes in the domain of losses. Second, missing an all-or-nothing goal generates strong negative emotions. The individual then engages in *emotional repair*. This takes the form of diverting attention to those activities that generate a feeling of well-being and positive emotion. Emotional repair could take the form of a pleasurable, but eventually harmful activity or a redirection of goals that diverts from the original goal. Third, following a period in which one suppresses an activity in order to achieve a goal (e.g., smoking), one's craving for it can reach a level that is higher relative to the level without suppression. This is termed as the *rebound effect*.

In conjunction, these effects might make goals counterproductive in some cases. The experimental results in Soman and Cheema support their assertion.

20.9.7 *The behavioral economics of retirement savings*

Thaler and Benartzi (2007) outline the behavioral economics approach to retirement savings; we return to these issues with a complementary discussion in Part 8. We summarize their main findings as follows.

1. Enrollment defaults: In actual practice, the tax rules in most countries provide significant inducements to contribute to pension plans. These rules include tax deductible contribu-

tions and tax deferred accumulations of pension savings. However, the take-up of pension plans is often suboptimal, possibly due to procrastination and inertia in joining these plans. There is significantly more participation in retirement savings plans when the default option automatically enrolls, but individuals can choose to opt-out if they desire (Choi et al., 2002, 2004); this is the opt-out approach. However, in the converse case of the opt-in approach, participation is much lower.

Why do defaults influence choices? One potential explanation is that initial preferences, particularly for new or unknown options, need to be constructed (Slovic, 1995; discussion in Section 20.7). However, constructing preferences is cognitively costly, while accepting a default option is costless, hence, more likely to be accepted. This channel is probably even more important in the face of procrastination arising from present-biased preferences.

2. Complexity of the pension plan: There is a negative correlation between the number of investment options offered in the savings plan and the participation rates. Furthermore, a simplification of pension plans in terms of a pre-selected contribution rate and asset allocation reduces procrastination and improves the take-up of pension plans (Choi et al., 2009; Beshears et al., 2013a).
3. Initial contribution rates: Initial contribution rates in many savings plans are found to be too low. The typical default savings rate is 2–3% and the investment choice is normally a conservative one (e.g., money market account). Savers often use mental heuristics to decide on contribution rates as multiples of 5 or save at the minimum rate required to get employers to match at the upper allowable limit.
4. Portfolio allocation: Many pension plan participants use the $\frac{1}{n}$ decision heuristic, i.e., given a choice between n investments, choose to allocate the portfolio equally among the options; see Section 20.9.5, above. Experts in finance need not be exempt from the diversification heuristic. Thaler and Benartzi (2007) report the following quote from Harry Markowitz. “I should have computed the historic covariances of the asset classes and drawn an efficiency frontier. Instead, . . . I split my contributions fifty-fifty between bonds and equities.” Depending on the fraction of risky/riskless assets in the portfolio, the $\frac{1}{n}$ decision heuristic can expose the plan participant to a non-optimal level of risk. If the diversification heuristic is important, then it would seem that adding options to the 401(k) plans will have little effect when the number of options is already large. Two additional features of portfolio allocation are of interest.

First, individuals invest a disproportionate fraction in the stock of their own company (Mitchell and Utkus, 2004). Thaler and Benartzi (2007) report that 33% of individuals who own company stock do not realize that it is more risky, relative to a diversified portfolio. It is possible that company stock is a unique asset that is neither viewed as equity nor as bonds (Benartzi and Thaler, 2001). Portfolios in retirement plans that do not have company stock are normally split evenly between bonds and stocks. But in the presence of company stock, pension portfolios invest 43% in company assets and the remaining amount is split evenly between stocks and bonds. Hence, those with company stock in their pension portfolios end up investing a disproportionate amount of their pension savings, about 71%, in equities. Second, in evaluating the attractiveness of assets, individuals overly extrapolate past performance of a stock for the foreseeable future (Benartzi, 2001). Third, the availability of faster information on the Internet and other newer communication media means that the variability of stock prices will be even more salient. The latter will possibly create even greater myopic loss aversion.

5. Intervention by plan providers: There is little evidence to indicate that education improves saving rates among workers. However, elements of plan design might play a beneficial role. An example is the “Save more tomorrow” (SMT) plan (Benartzi and Thaler, 2004). Suppose workers were asked to increase their plan contributions now, conditional on the current wage. Given loss aversion, this might be viewed as a loss in current consumption, so it is less likely to be successful in increasing savings. Hence, a more sensible strategy is to have workers agree to deduct a fraction of a future wage increase as additional saving contributions, up to a certain limit (sometimes known as *automatic escalation*). If the reference income level of workers is the initial wage level, then additional contributions out of potential future income do not induce loss aversion. Evidence indicates that the SMT plan has led to an increase in savings (Benartzi and Thaler, 2013).

CHAPTER 21

Bounded Rationality in Financial Markets

21.1 Introduction

Financial markets have traditionally been held up as an ideal example of efficiently functioning markets. These markets are often used in introductory courses in economics to demonstrate why it is not restrictive to focus exclusively on fully rational economic agents. By way of extra realism, it is typically added that economics does not rule out irrationality; however, it is argued, this can be safely ignored. Two arguments are typically given to justify this conclusion on a-priori grounds. First, individual errors will cancel out in the aggregate. Second, irrational market participants (or *noise traders*, a term due to Black, 1986), if any, will be driven out of the market by rational market participants (sometimes known as *arbitrageurs*).

In stark contrast to the natural sciences, fundamentally important assumptions are often justified in economics on grounds of plausibility, rather than an appeal to rigorous evidence. In the case of financial markets, the theory of efficient financial markets has not been very successful in accounting for the evidence. Furthermore, a great deal of progress has been made in the theoretical models in behavioral finance to better explain the economic implications of noise traders. The term *investor sentiment* that describes the activities of noise traders is now routinely used, even in the popular press.

The *efficient markets hypothesis* (EMH) is the cornerstone of classical finance. The EMH stresses that, at any instant in time, the price of a security equals its fundamental value, i.e., the present discounted values of all future income flows (e.g., dividends) associated with the security. In Section 21.2.1, we outline the concept of *short-selling* or *shorting* of an asset, a mechanism that lies at the heart of EMH and effective arbitrage. While this concept will be quite familiar to those who have taken a basic course in finance, for others it is important to grasp its importance.

Section 21.2.2 gives a statement of the EMH, and considers six different implications of EMH that have been subjected to empirical testing. These include, for instance, the *weak form* of EMH (stale information should not affect current prices); *semi-strong form* of EMH (publicly available information should not affect current prices); *strong form* of EMH (insider trading should be ineffective); prices of substitute securities should not differ; and prices should not react to no-news.

In Section 21.2.3, we outline the theoretical objections to the EMH; formal models of these objections are given in subsequent sections of this chapter. On account of various judgment heuristics, individual investors may exhibit *investor sentiment* that can lead to systematic departures of prices from fundamental values. In the face of persistent investor sentiment, arbitrageurs

might not be able to hold on to their rational, but loss making, positions for too long, hence they might have to exit the market. Furthermore, professional arbitrage does not solve the problem either, for a variety of reasons that we discuss in more detail, below.

The recent literature has focused on two misperceptions in decisions under risk. First, overconfident individuals assign higher probabilities to their stocks doing well. Second, they overweight small probabilities (see Part 1). These misperceptions have led to interesting insights in the emerging literature on the *psychology of tail events* (Barberis, 2013). For instance, the market will price skewness in the returns to assets, and the two misperceptions also potentially explain the equity premium puzzle.

In Section 21.2.4, we consider the empirical evidence for EMH. Stock prices are far too volatile relative to the predictions of EMH (Shiller, 1981). The weak form of EMH is contradicted in many empirical studies (De Bondt and Thaler, 1985). Stale information can predict future prices; past winners over a long horizon give low subsequent returns and past losers give higher subsequent returns. Stocks with high *market to book value ratios* (MBVR) in the past do worse in the future (Lakonishok et al., 1994). There is also a well-documented *momentum effect*; stocks that have done well in the recent past continue to do well while those that have recently done badly continue to do so (Jegadeesh and Titman, 1993). The semi-strong form of the EMH is contradicted by evidence of a *post-earnings-news drift* in which prices continue to move in the direction of initial public news (Bernard and Thomas, 1989; Michaely et al., 1995).

Prices of strictly substitute securities are found to differ systematically. There is particularly clean evidence on this effect from *Siamese-twin companies* that have merged in the past, but continue to operate and trade in different markets, e.g., Royal Dutch Petroleum, and Shell Transport and Trading, PLC. A security in such a company is a claim to the total income of the merged companies in the ratio in which they originally merged their assets. However, the observed market ratios of the security prices of these companies differ systematically from the ratio predicted by theory (Froot and Dabora, 1999). Security prices also react to no-news and many stock market crashes occur without any apparent change in fundamental values (Cutler et al., 1991). The price of freshly inducted stocks in the S&P 500 index jumps upwards without any new information about the fundamental values (Shleifer, 1986; Wurgler and Zuravskaya, 2002).

Defenders of the EMH often cite the difficulty of beating the stock market as proof that financial markets are efficient. However, this is simply a confusion between necessary and sufficient conditions (Barberis and Thaler, 2003). The EMH implies that there are no arbitrage opportunities; however, the absence of arbitrage opportunities does not logically imply that EMH holds. Several interesting findings have emerged about the performance of individual investors in stock markets (Barber and Odean, 2013). On average, individual investors do indeed earn below average market returns on longer-term investments, and have poor stock selection ability, while their success at shorter horizons is better. Furthermore, there is evidence that many institutional investors, and the most skilled individual investors, can earn abnormal returns that are correlated across time. There is substantial heterogeneity in the performance of individual investors that is correlated with age, overconfidence, and gender.

In Section 21.3, we describe a seminal model of noise trader risk in which there are two assets, a riskless asset and a risky asset (De Long et al., 1990a). The only source of uncertainty is the noisy beliefs of the noise traders about the future price of the risky asset, which otherwise is subjected to no underlying fundamental uncertainty. The presence of noise trader risk creates uncertainty about the future price of the risky asset. Hence, it affects the current demand of the risky asset. Risk averse arbitrageurs are deterred by the presence of future uncertainty, so they reduce their demand for the risky asset. However, due to investor sentiment, the noise traders

might be currently bullish and demand higher amounts of the risky asset. If the future price turns out to be high due to the future bullishness of noise traders, then the profits of noise traders who had overinvested in the risky asset are higher. Under appropriate conditions, that we outline, noise traders may earn higher average profits than arbitrageurs.

Modern financial markets are characterized by a separation between brains and capital. Small and dispersed investors typically lack detailed knowledge of the working of financial markets, and often use the services of professional investors. Professional investors invest the capital of others, and charge a fee in return for their services. Since professional investors are trained and supposedly knowledgeable, can we be sure that the heuristics and biases of individual investors that impede market efficiency, no longer operate? While the traditional biases might not operate, new biases are introduced, which once again preclude market efficiency. We consider this question in several places in Part 7 (see, for instance, Sections 19.18 and 21.7).

In Section 21.4 we consider a formal model of *professional arbitrage* (Shleifer and Vishny, 1997). Since small investors do not have the expertise to evaluate professional arbitrageurs, they resort to a simple heuristic. Namely, invest more in a professional arbitrageur whose past market returns are high; this is known as *performance based arbitrage*. However, in any period, a negative shock to investor sentiment reduces current prices and the returns of professional arbitrageurs. Small investors who rely on performance based arbitrage then withdraw funds from professional arbitrageurs precisely when current prices are low and so the potential returns on investing are high. This hampers the ability of professional arbitrageurs to push prices in the direction of fundamental values.

An important assumption underlying the EMH is that the flow of any relevant information is instantaneous. However, a great deal of evidence shows that information flows in financial markets may be gradual, so that stock prices adjust slowly (Hong et al., 2000). In Section 21.5, we outline a formal model that combines gradual flow of information with heterogeneity among investors (Hong and Stein, 1999, 2007). We show that there can be intense trading even when prices equal fundamental values because investors may hold incorrect models of the real world. Information that may correct their models flows gradually. We also show that *momentum* and *trading volume* are closely related.

In Section 21.6, we consider the twin phenomena of stock market *underreaction* and *overreaction*. Stock market underreaction takes place when good news about a stock is not immediately incorporated into its price so that we expect its return to be higher next period. Overreaction takes place if following a sequence of good news, lower returns are expected in the future, indicating that the market initially overreacted. We use a formal model in which investors hold on to two incorrect models of the world, a mean reverting model and a trending model, and based on stock market news, update their beliefs about the models (Barberis et al., 1998). Within such a framework, we derive appropriate restrictions on the parameters that ensure the model can match the data on underreaction and overreaction.

Section 21.7 considers the subject of *behavioral corporate finance*. Classical corporate finance studies the interaction between managers and shareholders of a firm and decisions such as the source of finance, investment decisions, and design of managerial compensation. Behavioral corporate finance enriches the domain of this enquiry by considering the possibility that managers and shareholders might not be fully rational in the sense of classical finance. There are two main difficulties in this area. First, unlike the EMH in behavioral finance, there is no clear benchmark to test. Second, corporations are naturally reluctant to part with internal data that would allow for more stringent tests of theory. Despite these hurdles, impressive progress has already been made.

Section 21.7.1 considers a range of biases exhibited by managers of firms. Many managerial decisions are of an infrequent nature (e.g., capital restructuring, mergers) that offer limited learning opportunities. These decisions are often taken under extreme uncertainty, so not only is it unclear what the optimal decision is, methods for evaluating the quality of these decisions are also vague, and subject to alternative interpretations. On account of these reasons, when outcomes are bad, managers can blame luck, but when the outcomes are good they can ascribe them to their ability levels; this is known as *attribution-bias* (Langer and Roth, 1975; Bettman and Weitz, 1983; Bertrand and Mullainathan, 2001).

Managers and CEOs are also known to exhibit a great deal of *overconfidence*, and evidence indicates that more overconfident managers, purely because they might have been lucky, rise to the top (Goel and Thakor, 2008). Due to overconfidence, managers might overvalue their firms, therefore, relying on borrowing rather than raising equity for investments (Malmendier and Tate, 2005). Overconfident managers are also more likely to merge because they are overconfident of their abilities in turning around the target firm (Roll, 1986; Malmendier and Tate, 2008; Gervais et al., 2011).

Sections 21.7.2 and 21.7.3 consider two kinds of polar possibilities. In Section 21.7.2, we consider the possibility that managers are rational, but investors are potentially irrational and draw the implications for merger activity (Shleifer and Vishny, 2003). We consider the short- and long-run impacts of acquiring a firm through three different methods: cash, stock, and outright purchase of capital. We also consider empirical evidence on three main objectives of mergers (Baker and Wurgler, 2012). These objectives are: *maximizing fundamental values*; *catering* to increase stock prices above fundamental values; and *market timing*, i.e., taking advantage of temporary stock mispricing to issue new shares or repurchase old shares. In Section 21.7.3, we consider the other polar extreme of irrational managers and rational investors. Irrationality of managers takes the form of managerial overconfidence (Gervais et al., 2011). In this context, we also clarify the role of risk aversion, manager's ability, and alternative incentive schemes.

Finally, in Section 21.7.4, we consider *behavioral corporate signaling*. Classical corporate finance explains the existence of dividends, IPO underpricing, capital restructuring, and debt as costly signals to investors of the quality of the firm (Ross, 1977; Myers and Majluf, 1984; Harris and Raviv, 1985). However, empirical evidence shows that the main reason for issuing dividends is to take advantage of temporary mispricing of securities. Indeed, survey evidence rejects the main implications of the classical signaling model in corporate finance (Brav et al., 2005). When surveyed, managers reject the idea that dividends are costly signals of the competence of firms that might reduce the value of the firm. We consider the possibility that benevolent managers may signal the underlying earnings of the firm through dividends (Baker and Wurgler, 2011). Yet these signals do not reduce the value of the firm, nor do they signal the strength of the firm.

21.2 The efficient markets hypothesis

21.2.1 *An introduction to short selling*

We shall need the standard concept of a *short position* or *shorting* to understand some of the material in this chapter. Suppose that there are two time periods, $t = 1, 2$ and that there is a single security priced at P_1 in period 1. Ignore issues of time discounting. Suppose that investor A believes that the second period price, P_2 , will be lower than P_1 . It makes sense for investor A to sell the security in period 1 and buy it back in period 2. For physical goods, one generally needs

to have possession of the goods before one can sell them. This is not the case for financial instruments such as securities. It is possible that another investor, say, investor B, actually possesses the security. Investor A can then ask a broker to lend him n units of the security for one period. The broker borrows n units of the security from investor B and lends it to investor A for one period.

Investor A sells the security in period 1 and receives an amount of money equal to nP_1 . In the next period, period 2, investor A purchases n units of the security at a price P_2 and returns the security to investor B, through the broker. The broker, who performs the financial intermediation, much as banks do, receives a brokerage fee for the service. The profit of investor A equals $n(P_1 - P_2) \geq 0$. Thus investor A can make a profit if beliefs turn out to be correct, but can also make a loss if they do not. In particular, this strategy of shorting is successful only if the actual price falls; thus, shorting is typically a response to expectations of a price fall in the future. If any dividends or interest payments accrue to the security during the period that it is in the possession of investor A, then these must be passed on to investor B, via the broker. There are also regulations on which types of securities can be shorted and when.¹ Indeed, these issues were part of the regulatory response to the 2007–2008 financial crises.

In the examples below, we use the following structure. Consider only two time periods, 1 and 2, and assume that at the end of period 2, investors simply sell all shares and spend the money on consumption. Consider an investor who has a choice between buying securities either in Ford or in General Motors. Since these companies sell similar products, we should expect that they are exposed to the same *fundamental risk* (say, involving the car industry as a whole). *Idiosyncratic risk*, on the other hand, is risk that is specific to each company such as a problem that might arise with brake disks or software in one of the companies; e.g., the problem with the so-called “defeat software” in VW cars in 2015 is an example of idiosyncratic risk. We assume that ex-ante, each company carries the same idiosyncratic risk. Suppose that the fundamental value of a security in each time period for each of the two companies is \$30. Since the risk and return on securities issued by both firms is similar, they are *close substitutes*.

Example 21.29 (*Efficiency in financial markets*): Let us work through an example to see how the arbitrage argument advanced by Friedman (1953) and Fama (1965), works to produce the efficient markets hypothesis. For whatever reason, say, on account of irrational investors, suppose that Ford shares drop to \$20 in period 1. Arbitrageurs jump on this opportunity to take a short position in GM shares. Consider an arbitrageur who borrows 6 GM shares in period 1. He sells them for $30 \times 6 = \$180$ in period 1 and buys $\frac{180}{20} = 9$ Ford shares with the money. The buying of Ford shares in this manner will continue until excess demand puts sufficient upward pressure on the price of Ford shares for them to equal the fundamental value of 30. Suppose that, as expected, on account of the activity of arbitrageurs, Ford shares revert to their fundamental value of 30 in period 2. Then the arbitrageur sells the 9 Ford shares for $30 \times 9 = \$270$ in period 2. Using this money he buys 6 GM shares for $30 \times 6 = \$180$ in period 2, which are returned back to the broker. The total profits of the investor equal $270 - 180 = \$90$. The efficient markets hypothesis holds because each security now equals its fundamental value; the presence of substitute securities, shorting activity, and the arbitrage activities of rational arbitrageurs are all critical to this result.

Example 21.30 (*Substitute securities, shorting and protection from fundamental risk*): The main advantage of taking a short position in a substitute security (GM in this case) is that it protects the investor from fundamental risk. To see this, suppose that in the second period, on

¹ These and other insitutional details can be picked up from any basic text on finance.

account of fundamental risk, the fundamental value of both securities unexpectedly drops to 20. In this case, in period 2, the arbitrageur sells the 9 Ford shares for $20 \times 9 = \$180$. Using this money he buys 6 GM shares for $20 \times 6 = \$120$ in period 2, which are returned back to the broker. The total profits of the investor then equal $180 - 120 = \$60$. So, while fundamental risk reduces the price at which the Ford shares can be sold (fall in revenue from \$270 to \$180) it also reduces the cost of purchasing the GM shares (fall in cost from \$180 to \$120) that have to be returned back to the broker in period 2. Hence, the act of shorting, reduces the exposure of the investor to fundamental risk. However, legal constraints might prevent many investors, particularly mutual fund and pension managers from taking a short position. This exposes them to fundamental risk, hence, they might undertake insufficient arbitrage to ensure that price equals the fundamental value of the security.

Example 21.31 (Exposure to idiosyncratic risk remains): Taking a short position does not, however, protect the investor from idiosyncratic risk. Continue to assume the setup of Examples 21.29, 21.30. However, now suppose that in period 2, on account of idiosyncratic risk (say, recall of Ford cars because there was a problem with faulty brakes), the price of a Ford security drops to \$20 in period 2. The arbitrageur then sells the 9 Ford shares for $20 \times 9 = \$180$ in period 2. Using this money he buys 6 GM shares for $30 \times 6 = \$180$ in period 2, which are returned back to the broker. The total profits to the investor then equal $180 - 180 = \$0$. Worse, if the Ford securities drop to a price below \$20, the investor makes a loss.

21.2.2 Statement and implications of EMH

The efficient markets hypothesis (EMH) has been a cornerstone in neoclassical finance. We give a brief outline of its definition and assumptions.

Definition 21.9 EMH implies that the value of a security at any point in time equals its fundamental value, i.e., the present discounted value of all future dividends and earnings. All publicly available information is immediately incorporated into security prices. Otherwise prices would not be the same as fundamental values at some instant in time.

The EMH rests on the following main assumptions.

- A1: Rationality:** By and large, investors are assumed to be rational. In finance, rational investors are called *arbitrageurs*. Rationality has two requirements in this case.
 - (i) In response to new information, investors update their priors in a Bayesian manner.
 - (ii) When faced with a risky situation, investors have beliefs over outcomes and use subjective expected utility to evaluate the risky situation.
- A2: Irrationality, if any, cancels out in the aggregate:** Any irrationality, if it exists at all, is random and so, in the aggregate, it is expected to cancel out.
- A3: Irrational investors eventually exit the market:** If irrational traders bet on a price that is different from the fundamentals, then arbitrageurs will step in and make money off them. Since, irrational traders cannot continue to lose money forever, they must eventually exit the market.

How does one test the EMH? The implication of EMH is that prices should respond instantly to news that alters fundamental values. Thus, instantaneously, price should equal the new fundamental value. EMH implies the following six testable implications, T1–T6.

- T1: Stale information should *not* affect current prices of securities. It must already have been incorporated into security prices on account of the activity of the arbitrageurs. Thus, past prices and returns should not affect current stock prices. This is known as the *weak form of EMH*.
- T2: Any publicly available information should *not* affect current prices of securities. Such information is expected to be instantaneously incorporated into security prices. This is known as the *semi-strong form of EMH*.
- T3: Investors should *not* be able to earn more than the average market profits by using information that might not be available to others (absence of *insider trading*). The expectation is that such information would be quickly picked up by other investors, wiping out any potentially abnormal profits. This is known as the *strong form of EMH*.
- T4: Security prices must *not* first underreact to news about fundamental values, but then overreact to it before finally reverting to the new fundamental value, possibly over several periods.
- T5: There should be *no* consistent differences in the prices of strictly substitute securities.
- T6: Security prices should *not* react to *no-news*, i.e., news that is unrelated to fundamental values.

21.2.3 *Evaluating the theoretical case for EMH*

We first catalog, informally, some theoretical objections to the EMH that are treated more formally in subsequent sections.

1. *Availability of substitute securities*: Substitute securities are not always available.² The existence of such securities is essential to the argument for EMH because they allow investors to take a short position and, hence, protect themselves from fundamental risk (Examples 21.29, 21.30). The absence of substitute securities and fears of inability to insure against fundamental risk might dissuade arbitrageurs from undertaking the right level of economic activity that is required to ensure that prices of assets equal fundamental values.
2. *Investor sentiment and judgment heuristics*: The EMH assumes that investors use Bayes' rule to update their priors. Our discussion of judgment heuristics in Chapter 19 suggests that investors often do not use Bayes' rule correctly (e.g., base rates are underweighted). People are known to assign confidence intervals to their estimates that are too narrow and 90% of those surveyed report that they have above-average levels of intelligence and emotional ability.³ At the same time, individuals can be too cautious and underestimate the likelihood of the sample. This is known as *conservatism* (see Chapter 19).

In other contexts, individuals rely on *representativeness* and exhibit the law of small numbers. People use the *availability heuristic*, relying on their memory for the available information. However, all memories are not equally retrievable or salient. For instance, more recent events and more unusual events might be more salient and more retrievable, which can result in individual actions being biased. Any departure from Bayesian rationality, typically on account of judgment heuristics, is termed as *investor sentiment*. Other names for such investors include *unsophisticated investors* (Kyle, 1985) and *noise traders* (Black, 1986).

² See, Figlewski (1979), Roll (1988), Campbell and Kyle (1993). For a review see Shleifer (2000), Barberis and Thaler (2003).

³ See Weinstein (1980) for an early study and Barberis and Thaler (2003) for references to the subsequent literature.

In example 21.29, suppose that the decrease in the original price of Ford securities in period 1 from \$30 to \$20 is caused by investor sentiment, say, noise traders. Arbitrageurs (or rational traders) then face an additional risk when they “short” GM securities. Namely, that investor sentiment may get worse in period 2. This problem may become even more serious due to the presence of *positive feedback traders*, i.e., traders who base their decision to buy an asset only if it performed well last period (De Long et al., 1990b). Following the lead of noise traders, such investors may be even more reluctant to buy Ford securities in period 2, increasing the downward pressure on Ford securities. Thus, mispricing of Ford securities in the first period, caused by investor sentiment, may persist in period 2.

In this case, Ford shares do not revert to their fundamental value of 30 in period 2. Suppose that on account of noise traders and positive feedback traders, the period 2 price of Ford securities turns out to be 15. In the two-period setup of Example 21.29, the arbitrageur sells the 9 Ford shares for $15 \times 9 = \$135$ in period 2. The investor then needs to purchase the 6 GM shares which cost $30 \times 6 = \$180$ in order to return them back to the broker. The total profits of the investor equal $135 - 180 = -\$45 < 0$. This is the risk created on account of noise traders and positive feedback traders, which might well limit arbitrage.

Indeed, if arbitrageurs can see through the behavior of noise traders then they might take actions that bet on mispricing getting worse. It might also be the case that noise traders take a long time to lose money and to exit the market. The result could be that arbitrageurs liquidate their positions before noise traders can exit the market because they cannot hold on to their positions for too long (Figlewski, 1979). This takes away the normal countervailing effect of arbitrageurs to bring price back in line with fundamental values; Section 21.3 below formalizes these issues. This simple example also demonstrates the dangers of holding on to entrenched positions, based on purely ex-ante plausibility, such as the following arguments given to undergraduate students in economics. (1) If there are any irrational agents, then rational agents will make money off them and drive them out of the market. (2) Any irrationality in the market cancels out in the aggregate. The work on judgment heuristics (Chapter 19) rejects this assumption. It shows that laypersons and experts have persistent and *systematic biases*. Hence, the probability that such errors point in the same direction, rather than canceling out, is much greater, with serious implications for financial stability.

Noise traders is not just a theoretical construct. Consider the evidence from the market price of shares in closed-end funds. The manager of a closed-end fund raises a fixed initial amount of money towards a fund. Additional money can neither be added to, nor withdrawn from the fund, although shares in the fund trade actively in the stock market. EMH predicts that the market price of a security in the fund should equal the net asset value (NAV) of the fund. However, systematic differences between the market price and NAV were found by Lee et al. (1991). To get some idea of these differences, consider the updated figures from Thaler (2015) drawn from closed-end funds: Shares in Gabelli Utility Trust sold at a premium of 18.2% over the NAV, while shares in Firsthand Technology Val sold at a discount of 37.4%.

Lee et al. (1991) also provide evidence to explain the following four *closed-end fund puzzles* that cannot be accounted for by classical models in finance. At the time of starting the fund, security prices trade at a sizable premium of about 7% due to start-up costs and underwriting costs. Within six months, the fund trades at a discount of 10%. There is

substantial variation in the premia and discounts over time. However, when there is a merger or liquidation or conversion of the fund to an open-end fund (to which money can be added and withdrawn), the market price converges to the NAV. Why do these systematic patterns of returns persist if all market participants are rational?

It turns out that the closed-end puzzles can be explained by the existence of noise traders. The basic idea is that individual investors are more likely to invest in the shares of small firms (which includes closed-end funds), relative to larger institutional investors. Thus, insofar as investor sentiment is more likely to be exhibited by small investors rather than institutional investors, the time variation of investor sentiment is likely to be reflected in the relative stock performance of small and large companies. Lee et al. (1991) were able to show that they could correlate well the premium/discount in the security prices of closed-end funds with the difference between the performance of small and large companies. For more empirical evidence on noise traders and a detailed discussion of the closed-end fund puzzles, see Shleifer (2000).

In addition to incompatibility with Bayes' rule and choosing on the basis of representativeness, or anchoring, several other forms of investor sentiment have been documented. Hindsight bias can lead to underestimation of financial volatility (Biais and Weber, 2009). In order to be induced to hold stocks, sellers must offer too high a premium on equities relative to bonds; this is known as the equity premium puzzle (Benartzi and Thaler, 1995). Investors also hold on to losing stocks for far too long and sell winning stock too early relative to the predictions of standard theory; this is known as the *disposition effect* (Odean, 1998). Both these phenomena can be explained using prospect theory (see Part 1 of the book). Benartzi and Thaler (1995) also demonstrate the importance of *framing effects* on investor choice. In particular, the decision of investors to buy stock depends on whether news about the performance is presented as *high, long-term, returns* or a *more finely presented series of fluctuating short-term stock returns*.

The following quote from Shleifer (2000, p. 10) illustrates a set of other judgment heuristics and biases that run contrary to the assumption of EMH, further undermining the theoretical case for EMH:

... it is difficult to sustain the case that people in general, and investors in particular, are fully rational. At the superficial level, many investors react to irrelevant information ... Investors follow the advice of financial gurus, fail to diversify, actively trade stocks and churn their portfolios, sell winning stocks and hold on to losing stocks thereby increasing their tax liabilities, buy and sell actively and expensively managed mutual funds, follow stock price patterns and other popular models. In short, investors hardly pursue the passive strategies expected of uninformed market participants by the efficient market theory.

Heuristics also explain why there is far too much trading and churning of assets. Thaler (1999) reports that 700 million shares are traded on the NYSE. For 2005, Hong and Stein (2007) give a figure of \$51 trillion in trading volume for the 50 members of the World Federation of Exchanges. This is too large a volume to be plausibly explained by classical theories in finance. Barber and Odean (2000) show empirically that there is far too much trading relative to a rational benchmark. Indeed, in their sample, and after taking account of trading costs, investors would have done better by trading less. Furthermore, in one possible manifestation of overconfidence, investors who traded the most, earn the lowest average returns.

The phenomenon of too much trading by market participants is different from the puzzle of limited stock market participation. Given the relatively high mean rates of returns, why are so many potential investors unwilling to invest? Barberis et al. (2006) argue that the answer lies in narrow framing of options in which individuals segregate stock market returns from other sources of income. Once this is assumed, then they invoke the tendency of individuals to reject small gambles, even those that are actuarially favorable (Tversky and Kahneman, 1992). Thus, a combination of first order risk aversion and narrow framing provides one possible rationalization of the puzzle of stock market participation.

Expected utility theory underpins EMH. In the general case under risk, the formula—prices of stocks equal the present discounted value of future dividends—relies on a first order linear stochastic difference equation that exploits the arbitrage relation between the returns on safe and risky assets (Blanchard and Fischer, 1989, ch. 5). We have seen in Part I of the book that, contrary to expected utility theory, the bulk of the empirical evidence supports *non-linear probability weighting*. In particular, individuals overweight small probability events and underweight high probability events. Furthermore, individuals may often think that an event is more probable than it actually is, e.g., the probability of an air-crash. In stock markets, overconfidence is a particularly salient factor that might cause *overassignment of probabilities*.

These two misperceptions (non-linear probability weighting and overassignment) have several implications that are studied under the emerging topic of the *psychology of tail events*. An important implication is that the market will price skewness in asset returns (Barberis and Huang, 2008). Suppose that the return on an asset is positively skewed. Then there is a small probability that the investor will earn very high returns. On account of the two misperceptions mentioned above, the demand for such assets is higher than it should ideally be, pushing up the prices and reducing the average returns. Empirical evidence is supportive of this prediction (Boyer et al., 2010). These ideas can also be used to explain the equity premium puzzle (De Giorgi and Legg, 2012): The two misperceptions lead investors to overestimate and overweight the small probability of a stock market crash, hence, they demand a premium for holding stocks relative to bonds.

In Part I of the book we documented two kinds of behaviors for very low probabilities; the first, overweighting of small and underweighting of large probabilities has been mentioned above. A second kind of behavior arises when we approach even lower probabilities; these extremely low probabilities are either underweighted or simply ignored (e.g., we do not scurry around to get insurance against an asteroid hitting us). This explains why individuals are reluctant to purchase insurance for low probability natural hazards and a host of other behaviors (see Part I for details). Taleb (2007) draws attention to this kind of behavior in his book, *The Black Swan*, where individuals assign very low probabilities to tail events. In this case, the misperceptions are the exact opposite of the two misperceptions reported above.

How do we know which of the two sets of misperceptions (overweighting or underweighting of very small probabilities) applies? A plausible suggestion is that the availability heuristic may help distinguish between the two cases (Barberis, 2013). Events that are rare and unusual do not bring to mind similar instances. Thus, individuals might have underweighted the probability of the stock market crash of 2008 because they could not recall how events similar to 2005 and 2006 led to a crash in the past.

3. *Professional arbitrageurs*: In actual practice, there are a large number of small and dispersed investors. Professional arbitrageurs typically make the bulk of their investments in financial markets on behalf of these small investors; we may term this as the *separation of brains and capital* (Shleifer and Vishny, 1997). Several features of the operation of professional arbitrageurs are relevant for the success of EMH. These mainly arise from the inability of dispersed and small investors to fully comprehend the operations of financial markets; thus, they rely on indirect signals of the ability of professional arbitrageurs. Among the indirect signals used by small investors are the following. (1) The end of year reports on the performance of professional arbitrageurs. (2) How well one professional manager does relative to another? (3) The performance of professional arbitrageurs relative to some publicly observed target, such as the S&P 500 Index.

Each of these features creates serious distortions. Feature (1) often forces professional arbitrageurs to liquidate their positions in a loss making asset too quickly, even when holding on to such a position would have been beneficial in the longer run. This makes it difficult to counter the activity of noise traders.

Feature (2) induces professional arbitrageurs to invest in similar assets so that each can avoid looking worse than each other; this is a form of *herding behavior* (Scharfstein and Stein, 1990). Feature (3) induces professional arbitrageurs to choose portfolios whose returns are closely aligned to the S&P 500 Index. There is some evidence that these distortions reduce the performance of professional arbitrageurs below that of passive investment strategies (Ippolito, 1989; Lakonishok et al., 1992).

To these factors, one could add transaction costs of arbitrage which might prevent market efficiency. Transaction costs include the cost of finding information about which shares are mispriced relative to their fundamental values, as well as the brokerage fees for taking a short position (D'Avolio, 2002).

21.2.4 *The empirical evidence on efficiency of financial markets*

The early evidence in the 1960s till about the mid 1970s repeatedly confirmed the EMH (Fama, 1965, 1970; Jensen, 1978). It focused mostly on T1 and T2 above, i.e., the weak and the semi-strong forms of the EMH. Fama (1965) found that stock prices approximately follow a random walk. Fama et al. (1969) suggested the use of *event studies* to test for the EMH and these became very popular. The basic idea was to focus on a particular event. For instance, the announcement of earnings or dividends, news of a takeover or divestiture, news about the issue of new shares or repurchases of existing shares, and other news about changes in management or of management strategy. Researchers then checked to see if there was an effect on security prices, subsequent to these events, as information became stale. The expectation was that the effect of the news would be incorporated immediately, or even prior to the announcement of news. The results were predominantly supportive of the EMH; see Fama (1970) for a review.

Since the late 1970s there have been serious challenges to the EMH. The theoretical and empirical objections have not only focused on T1 and T2 above, but also on T3–T6. Unfortunately for the EMH, and despite econometric and data caveats, all of the implications in T1 to T6 were not only rejected, but the theoretical case for them was found to be on shaky ground (Shleifer, 2000; Barberis and Thaler, 2003). We have considered some elements of the theoretical case in Section 21.2.3 above and more details are considered later, but in this section we focus on the empirical evidence.




Figure 21.1 A plot of price and dividends against time.

Source: Shiller (2003) with permission from the American Economics Association.

The empirical evidence against EMH has been extensively documented.⁴ Shiller (1981) provided an early empirical challenge to the EMH. He found *excess volatility* in stock prices that could not be explained by fundamental values. Denote by P_t , the price of a stock at time t , by I_t the information set at time t , and by $E[D | I_t]$ the conditional expected present value of subsequent dividends. Then the EMH states that $P_t = E[D | I_t]$, or $P_t = D + u_t$, where u_t is a mean zero random variable. Thus,

$$D = P_t - u_t. \quad (21.1)$$

Under EMH, all information is incorporated in prices, thus, u_t and P_t are uncorrelated. Denoting by $V(x)$, the variance of x , (21.1) implies that:

$$V(D) = V(P_t) + V(u_t). \quad (21.2)$$

Since variances are non-negative, it follows that an implication of EMH is

$$V(P_t) \leq V(D). \quad (21.3)$$

A plot of P_t and D against time is shown in Figure 21.1; these variables are labeled respectively as “Real Stock Price (S&P 500)” and “PDV, Constant Discount Rate.” D was proxied by the present value of actual dividends paid in the following year. From Figure 21.1, it is evident that $V(P_t) >$

⁴ There are several good surveys; see, for instance, Shleifer (2000), Barberis and Thaler (2003), Subrahmanyam (2008), and Barber and Odean (2013). There are two textbook treatments that use a behavioral economics perspective to cover a wide terrain in behavioral finance; see Akerlof and Shiller (2008) and Shiller (2015).




Figure 21.2 Cumulative average residuals for W and L portfolios of 35 stocks.

Source: De Bondt, W. F. M. and Thaler, R. (1985). "Does the stock market overreact?" *The Journal of Finance* 40: 793–805. © American Finance Association. With permission from John Wiley & Sons.

$V(D)$, contradicting EMH. In other words, there is excessive stock price volatility relative to the predictions of the theory. This should not come as a surprise to behavioral economists. Stock market decisions are often likely to be based on investor sentiment and behavioral features that create another source of price volatility; these factors are not taken into account under classical finance. These results were challenged by Marsh and Merton (1986), mainly in terms of the proxies used for D and P_t as well as issues of time-invariant interest rates. However Shiller's basic result was borne out by several papers that strove to address the Marsh and Merton (1986) critique (Barberis and Thaler, 2003).

The weak form of EMH is contradicted by several studies. De Bondt and Thaler (1985) look at the performance of winners (W) and losers (L), over the past three years, on the New York Stock Exchange over the time period 1926–1982. They then compute the returns on W and L stocks for the following five years, controlling for the relative risks of the two stocks. They find that W stocks and L stocks give, respectively, low and high subsequent returns. They explain their results by positing that the stock prices have overreacted in the past. The returns on W (respectively, L) stocks overshot (respectively, undershot) their fundamental values. The result is that the market gradually corrects for the mispricing. This directly contradicts weak form EMH because past performance is able to predict future returns. Figure 21.2 plots the "cumulative average residuals" for Winner and Loser portfolios of 35 stocks over five years (60 months); for any asset, the residual is an estimate of the difference between the return on the asset and the market portfolio.

Similar arguments can be used to explain why companies with high *market to book value ratios* (MBVR) in the past do worse in the future (De Bondt and Thaler, 1987; Fama and French, 1992; Lakonishok et al., 1994). Stocks of such companies are also sometimes called *glamour stocks*. A high MBVR in the past could indicate overreaction, which is gradually corrected in the future by a drop in stock prices. This contradicts weak-form EMH because stale information on MBVR helps to predict future returns.

One response in classical finance has been to interpret the MBVR as a measure of fundamental risk of a company. Hence, a high MBVR is associated with higher risk and lower subsequent returns (Fama and French, 1993, 1996a). However, not everyone is persuaded by the interpretation of MBVR as a measure of fundamental risk (Lakonishok et al., 1994; Mackinlay 1995; Daniel and Titman, 1997).⁵

Another violation of the weak form of EMH is the *momentum effect* (Jegadeesh and Titman, 1993). They classify portfolios into deciles based on performance in the previous six months. The main finding is that portfolios in the top deciles in the previous six months give about 10% higher returns relative to those in the bottom deciles for the next six months.⁶ They consider a trading strategy that buys past winners and sells past losers. This strategy earns a compounded excess annual return of 12.01% on average.

The semi-strong form of the EMH is contradicted by several pieces of evidence. EMH requires public news to be immediately reflected in stock prices. However, there is evidence of a *post-earnings-news drift*; prices move in the same direction as the initial public news (up if the news is good and down if it is bad) for a significant amount of time. This post-public-news drift is well documented for a range of public news such as earnings news (Bernard and Thomas, 1989, 1990; La Porta et al., 1997), share repurchase news (Ikenberry et al., 1995), stock issue news (Loughran and Ritter, 1995; Spiess and Affleck-Graves, 1995), and new dividend issues (Michaley et al., 1995).

Another piece of evidence comes from the *January effect*. Using data over the period 1904–1974, Rozeff and Kinney (1976) find that the average stock market returns during the month of January were 3.48%, compared with an average monthly return of 0.42% during the remaining 11 months of the year, which is eight times higher. The January effect is found in many other studies (De Bondt and Thaler, 1985). Rozeff and Kinney used an equally weighted index of New York Stock Exchange prices, hence, their results could have been caused by the overweighting of small firms. This was confirmed by Reinganum (1983). The January effect cannot be fully explained by risk considerations and issues of tax liability. Since the company size and the time of the year are both public information, this evidence contradicts the semi-strong form of EMH.

EMH requires prices of strictly substitute securities to show no consistent differences (see T5 above). This is rejected by the evidence (Lamont and Thaler, 2003). This evidence is immune to the Fama (1970) critique mentioned above. Froot and Dabora (1999) examine the share prices of large *Siamese twin* companies. These are companies that (i) have merged at some point in the past, but operate in geographically separate areas, and (ii) pool their cashflows, but each twin retains its own stock. So if companies A and B are Siamese twins who have merged on an $x : y$ basis ($x + y = 1$), then shares in company A are a claim to $x\%$ of the total cashflows of the merged company. If, as EMH suggests, prices of shares equal fundamental values, then the ratio of share prices of companies A and B should also be in the proportion $x : y$.

Froot and Dabora (1999) examine three pairs of Siamese Twins. (1) Royal Dutch Petroleum and Shell Transport and Trading, PLC. (2) Unilever NV and Unilever PLC. (3) SmithKline Beecham. Rosenthal and Young (1990) had already shown that there is a significant departure of the relative

⁵ A more satisfactory objection is that the book value might not adequately capture the fundamental value of the firm. For instance, firms that have excellent future growth prospects have high book to market values; in contrast those that are perceived to face agency problems have low book to market values. For attempts to construct a better measure of MBVR, but in the context of merger decisions of firms, see Ang and Cheng (2006) and Dong et al. (2006).

⁶ Grinblatt and Moskowitz (2004) use tax factors to explain a part of the momentum effect. However, tax factors cannot explain the momentum effect in its entirety.




Figure 21.3 Log deviations of prices from the predicted 1.5 : 1 ratio for Royal Dutch and Shell Transport.

Source: Reprinted from *Journal of Financial Economics* 53(2): 189–216. Kenneth A. Froot and Emil M. Debora. (1999). "How are stock prices affected by the location of trade?" Copyright © 1999, with permission from Elsevier.

share prices of Royal Dutch-Shell and Unilever NV and Unilever PLC from the predicted ratio under EMH. Froot and Dabora report two main findings:

(a) Traded price is rarely in the proportion $x : y$, contradicting the EMH. Consider, for instance, Royal Dutch and Shell Transport, which merged in 1970 on a 60:40 or 1.5:1 basis. So a share in Royal Dutch is a claim to 60% of the cash flow of the merged companies. The percent deviations of the ratio of prices from the predicted 1.5:1 basis are significant and shown in Figure 21.3. Similar disparities are reported for all three pairs of companies in Froot and Dabora.

(b) The price of the stock of each twin is highly correlated with the market index shares where the twin predominantly trades. For instance, when the US stock market performs relative to the UK market, then share prices of Royal Dutch, predominantly traded on the US stock market, rise relative to Shell, which is predominantly traded in the UK market. The same effect is obtained when the US dollar appreciates against the British pound.

A corollary of T5 is that controlling for the degree of differences in assets, their returns should be similar. However, the well-known equity premium puzzle (see Part 1 of the book), is a counterexample. Stocks are riskier than bonds, however, the 7% a year premium on stocks over bonds, averaged over a sufficiently long period, is too large to be explained, even accounting for the level of risk (Mehra and Prescott, 1985; Benartzi and Thaler, 1995).

Finally, we turn to the last implication of EMH that security prices should *not* react to no-news (see T6 above). The stock market crash of 1987, in which the Dow Jones fell by 22.6%, did not follow any apparently relevant news. Cutler et al. (1991) find that no major news accompanied the 50 largest one-day stock price movements since World War II. A corollary of T6 is that prices of stocks should not be affected by factors that are unrelated to fundamentals; for the evidence from Treasury bills see Duffee (1996), Krishnamurthy (2002), and Greenwood and Vayanos (2010).

Evidence relevant for T6 also comes from the inclusion of new stocks into the S&P 500 index. From time to time, some of the existing stocks in the index are withdrawn because, say, companies merge or are taken over. New stocks are then used to replace the withdrawn stocks with the objective of maintaining the representativeness of the index. The EMH predicts no change in the price of the stock that has been newly inducted into the index. The reason is that the induction, per se,

does not convey any new information about the company that was not available earlier. However, the evidence is that there is substantially increased trading in the freshly inducted stocks, contradicting T6; see, for instance, Harris and Gurel (1986), Shleifer (1986), Wurgler and Zuravskaya (2002).

Another piece of evidence comes from the literature on bubbles/sunspots in asset prices (Shiller, 2003). Speculative bubbles could cause departures of the price of an asset from fundamental values. However, as this departure worsens, the situation becomes unsustainable. Investors start expecting a reversion of prices towards the fundamental value. When the bubble bursts, there is a large change in the price of the asset without any apparent news. Similarly, stock prices often react to an announcement of an increase, or an initiation of dividends when such events convey no new news about fundamentals. The traditional explanations have sometimes relied on “dividends as a signaling device.” However, see Benartzi et al. (1997) who do not find any evidence that a dividend increase signals future changes in earnings.

For most investors, it is very hard to beat the stock market. For instance, Malkiel (1990) points out that, on average, mutual and pension fund managers are not able to beat passive investment strategies. This would, on the face of it, seem to rescue the EMH (Rubinstein, 2001). However, as Barberis and Thaler (2003) point out, on closer examination, this just turns out to be a confusion between necessary and sufficient conditions. The EMH implies that

“Prices equal fundamental values $\Rightarrow$ no arbitrage opportunities.”

However, the absence of arbitrage opportunities is only a necessary, but not sufficient condition for stock prices to equal fundamental values. Hence,

“No arbitrage opportunities $\nRightarrow$ prices equal fundamental values.”

Thus, the lack of success of actual fund managers in beating the stock market cannot be taken as evidence that prices equal fundamental values, or that the EMH holds.

The rejection of EMH has not gone unchallenged; see for instance Fama (1970). News that is favorable for returns might also be associated with changes in the risk from holding the security. So we need a model that incorporates both risk and return (the capital asset pricing model is one such model). Suppose that for some security, the level of return and exposure to risk go up following some news about the security. Furthermore, suppose we observe that prices are below fundamental values, yet sufficient arbitrage activity does not take place. The Fama (1970) critique would imply that we do not see enough arbitrage precisely because the security is more risky now and the higher return justifies the greater risk. So there is no need for further arbitrage to put downward pressure on the return.

Every empirical finding that is contrary to the EMH could then be challenged on this technical ground. However, this possibility was already checked for in the results of stock market underreaction and overreaction by De Bondt and Thaler (1985), reported above. They measured risk using the then dominant model in finance, the capital asset pricing model (CAPM), under which the only risk relevant for the pricing of an asset is the covariance of the stock returns with the market portfolio; the correlation between the two is called “beta.” They found the beta value for winners to be 1.37 and the beta value for losers to be 1.03. Thus, losers were actually less risky than winners, which goes in the reverse direction to Fama’s critique. Beginning with a series of papers in the early 1990s, written by Fama and French, the response in classical finance has been to debunk the CAPM model (Fama and French, 1996a,b). Fama and French went on to construct

three factor and five factor models with additional explanatory variables that were proposed to give a better measure of risk, but it is not clear if these models are supported by the empirical evidence (Fama and French, 2015).

It is worthwhile to review some of the evidence of the success, or lack of it, of individual investors in beating the stock market. In general, the evidence suggests that, on average, individual investors earn below market rates of pre-tax return on investments, whether one takes account of transaction costs or not. However, for very short run investments, say, a week, the performance of individual investors is better, at least based on US data. This evidence is nicely summarized in Barber and Odean (2013) that we draw on here.

Based on the trading records of 10,000 investors, Odean (1999) gave evidence documenting below-average returns of individual investors in the absence of transaction costs of trade. Barber and Odean (2000) replicate these results in the presence of transaction costs from the trading records of 78,000 investors. Their main findings are as follows. First, transaction costs are the main reason for the underperformance of individual investors relative to the market. Second, those investors who traded more are, on average, more likely to realize lower returns. For instance, the difference in annual rates of return between the 20% of investors who trade the most and the 20% who trade the least, is 7 percentage points.

These results have been widely replicated in other countries. For holding periods ranging from a day to six months, Barber, Lee, et al. (2009) examine trading data from Taiwan over 1995–1999. Individual investors buy stocks that subsequently perform poorly and sell stocks that subsequently perform strongly. The behavior of institutional investors is the opposite, suggesting poor stock picking skills on the part of individual investors. Grinblatt and Keloharju (2000) show, for Finnish data, that individual investors earn low returns. Furthermore, they have poor stock selection abilities. They are net buyers of stocks that have weak future performance, while institutional investors are net buyers of stocks that do well subsequently. Poor stock selection ability of individual investors is also found for US data (Hvidkjaer, 2008; Barber, Odean, and Zhu, 2009).

For investments over a short horizon, say, one week, the evidence from the US suggests that individual investors perform quite well (Barber, Odean, and Zhu, 2009; Kaniel et al., 2012; Kelly and Tetlock, 2013). However, these results are not supported by Taiwanese data (Andrade et al., 2008; Barber, Lee, et al., 2009). For Finland, empirical evidence suggests that investors make losses on limit orders⁷ but gains on market orders (Linnainma, 2010), but the opposite result is found for Taiwan (Barber, Lee, et al., 2009).

There is substantial heterogeneity in investors' skills. There is some evidence that there might be a positive correlation between cognitive skills and investor performance (Korniotis and Kumar, 2013; Grinblatt et al., 2012) and a negative correlation between age and investor performance (Korniotis and Kumar, 2009). A strong gender effect has been noted by Barber and Odean (2001) who report two main findings. First men trade more than women. Second, those who trade more earn lower profits. Gender effects on the extent of trading are found for Germany, although these are weakened once we take account of risk aversion (Dorn and Huberman, 2005).

Men are often found to be more overconfident than women in empirical studies. This is one possible reason why men trade more and so realize lower returns (Glaser and Weber, 2007; Grinblatt and Keloharju, 2009). Evidence from US data indicates a strong correlation between self-reported measures of competence and the extent of trading (Graham et al., 2009).

⁷ A limit order is an order placed with a brokerage firm to buy or sell a fixed number of shares at a specified price or better.

Excessive trading may also be the result of *sensation seeking*. The idea is that individuals view trading activity as a form of entertainment, or sensation seeking. Thus, there might be a degree of substitutability between various sensation seeking activities so that when individuals have access to alternative sensation seeking activities (e.g., raised stakes in state lotteries or the introduction of a new legal lottery), they engage in less stock market trading (Dorn and Sengmueller, 2009; Gao and Lin, 2014).

Coval et. al. (2005) show that market inefficiencies can be exploited by skillful individual investors to earn abnormal profits, which exceed the profits arising from other strategies based upon size, value, or momentum. Furthermore, there is persistence in profits, across time, among skillful investors. Similar results of persistence in intertemporal returns among the most skillful investors are found for Taiwanese data by Barber et al. (2014). However, the percentage of these skillful investors is very low. As one might expect, evidence shows that the special skills required to earn above-average returns in the stock market are more likely to reside with institutional investors (Grinblatt and Titman, 1989; Daniel et al., 1997).

21.3 Noise trader risk

Any departure from Bayesian rationality on account of, say, judgment heuristics that are inconsistent with classical statistics is termed as *investor sentiment*. Investors who exhibit investor sentiment are sometimes called *unsophisticated investors*, or *noise traders*. In this section, we consider the seminal model of noise trader risk due to De Long et al. (1990a).

Noise, in this context, may be defined as any information that is irrelevant to the fundamental value of a security. Rational traders, or *arbitrageurs* in neoclassical theory, are assumed to ignore such noise. In contrast, noise traders view noise as useful information. Black (1986, p. 531) describes noise traders, thus: “People who trade on noise are willing to trade even though from an objective point of view, they would be better off not trading. Perhaps they think the noise they are trading on is information. Or perhaps they just like to trade.” The presence of noise traders provides a powerful theoretical case against EMH, which, for the reasons we have documented above, serves to limit arbitrage activities.

On account of the noise traders, arbitrageurs face two kinds of risk.

1. Arbitrageurs incur a *noise trader risk* by participating in arbitrage activity. For instance, mispricing caused by the noise traders might get worse in the future, or arbitrageurs might not be able to hold on to a loss making position for the length of time that it takes for noise traders to lose enough money to exit the market. Thus, prices might deviate from fundamental values for extended periods of time, even in the absence of any fundamental risk. Classical finance does not deny the existence of noise traders. However, and purely on grounds of theoretical implausibility rather than concrete evidence, it is argued that noise traders have no more than a fleeting impact on financial markets. Furthermore it is believed that rational arbitrageurs would trade against noise traders who would then lose money and exit the market (Friedman, 1953; Fama, 1965). However, this rules out the possibility that noise traders may hold on to irrational beliefs for a sufficient length of time and arbitrageurs might have to liquidate their loss making positions before noise traders can exit.
2. Arbitrageurs may undertake insufficient arbitrage for assets with high *fundamental risk*, provided they are sufficiently risk averse. By contrast, investor sentiment may be sufficiently

strong for noise traders to engage in active trading despite the high fundamental risk and risk aversion. This takes away the normal countervailing effect of arbitrageurs in bringing prices in line with fundamental values.

We now introduce noise traders into a simple model in finance (De Long et al., 1990a). The two main results, each contradicting the supposition in classical finance, are as follows. First, the returns of noise traders may exceed the returns of arbitrageurs for a length of time, precisely because of noise trader risk. Second, several anomalies in the financial economics literature can be reconciled by the model, e.g., excessive volatility of stock prices, and long-run under-pricing of stocks. The main assumptions of the model are as follows.

1. (*Basic economic structure*) Consider an overlapping generations model of two-period lived individuals.⁸ There is no population growth. In the first period of their lives, individuals are “young” and in the second period they are “old.” For pedagogical simplicity, the young do not consume, so only the old consume (this is not critical to the results). Individuals are born with an endowment of $y > 0$ units of a single consumption good, which is also the numeraire good whose price in any period, is normalized to 1. There are no bequests and no production decisions to be made, only financial decisions, which we now turn to.
2. (*Financial assets*) The only economic decision taken by a young person is to allocate the endowment, y , between a risky asset and a riskless asset. Both assets earn a one-period return $r > 0$ per unit (in terms of the consumption good). The main difference between the two assets is that the riskless asset is in perfectly elastic supply, while there is only one unit of the risky asset in the aggregate. One unit of the riskless asset can always be converted into one unit of the consumption good (and back), hence, the price of the riskless asset equals 1 every period. On the other hand, because the risky asset is in a fixed supply of 1 unit, its price in period t is denoted by P_t in units of the consumption good.
3. (*Types of investors*) There are two types of young people in any period. (a) *Noise traders* hold noisy beliefs about the returns on the risky asset that have no basis in the fundamentals. (b) *Rational arbitrageurs* hold correct beliefs about the returns on the risky asset based only on the fundamentals, although they take account of the presence of noise traders. Variables pertaining to these two types of traders are superscripted, respectively, by $i = n, a$ (n for noise traders and a for arbitrageurs). There is no fundamental risk in the economy because each of the two assets has an identical return of $r > 0$ per period. The one, and only, source of risk is the noisy beliefs held by noise traders (described below).
The proportion of noise traders and arbitrageurs in the population is, respectively, $\mu \in [0, 1]$ and $1 - \mu \in [0, 1]$. Within each type of investor, the population is completely homogeneous. The size of the total population is normalized to $N = 1$.
4. (*Sequence of moves*) In any time period t , each type of young individual, $i = n, a$, receives an endowment, $y > 0$, and decides to hold $\lambda_t^i \geq 0$ units of the risky asset, which cost P_t per unit, and $y - \lambda_t^i P_t$ units of the riskless asset. The old in any period t hold the 1 unit of the risky asset. Since $N = 1$, the total demand for the risky asset by the young is given by

$$\mu \lambda_t^n + (1 - \mu) \lambda_t^a.$$

⁸ Readers can consult any standard macroeconomics book for the overlapping generations model; see, for instance, Blanchard and Fischer (1989, ch. 3).

There is a fixed supply of one unit of the risky asset, hence, the market clearing condition that determines the price of the risky asset, P_t , in period t , is given by

$$\mu\lambda_t^n + (1 - \mu)\lambda_t^a = 1. \quad (21.4)$$

5. (*Nature of noisy beliefs*) In any time period t , noise traders misperceive the expected price of the risky asset in period $t + 1$ by a misperception term ρ_t that is normally distributed with mean ρ^* and variance η^2 , i.e.,

$$\rho_t \sim N(\rho^*, \eta^2), \forall t. \quad (21.5)$$

Hence, $\rho_t > 0$ implies overestimation (or overconfidence), while $\rho_t < 0$ implies underestimation (or conservatism) of expected returns on the risky asset. This is the only source of risk in the model. ρ_t is realized at the end of period t , hence, its realized value is known at the beginning of period $t + 1$. Arbitrageurs continue to think that the return on the risky asset is r , based on the fundamental values. In making their decisions on how much to invest at time t , arbitrageurs know that the price at time $t + 1$ will be influenced by the demand of noise traders at time $t + 1$. At time t , $\rho_{t+1} \sim N(\rho^*, \eta^2)$, so the investors do not know the realization of ρ_{t+1} . But they do know that this creates an uncertain price for the risky asset at time $t + 1$, P_{t+1} (expressed in terms of the consumption good).

6. (*Wealth and its moments when old*) A young arbitrageur at time t purchases λ_t^a units of the risky asset, which cost $\lambda_t^a P_t$ in units of the period t consumption good. When old, each unit of the risky asset gives a return of r and is sold at a price P_{t+1} . All assets must be sold when old because there are no bequests and individuals only derive direct utility from consumption. The remaining amount, $y - \lambda_t^a P_t$, is spent on the riskless asset. The riskless asset earns a return of r when old and can be sold at a price of 1 (expressed in units of the consumption good). Hence, the total wealth, expressed in units of the consumption good of the arbitrageur, when old, net of the cost, y , is given by

$$w_{t+1}^a = (y - \lambda_t^a P_t)(r + 1) + \lambda_t^a(r + P_{t+1}) - y, \quad (21.6)$$

or,

$$w_{t+1}^a = ry + \lambda_t^a [r + P_{t+1} - P_t(1 + r)]. \quad (21.7)$$

For a random variable at time $t + 1$, x_{t+1} , we use the notation $\bar{x}_{t+1} = E[x_{t+1} | I_t]$, $\sigma_{x_{t+1}}^2 = \text{Var}[x_{t+1} | I_t]$, respectively, to denote the mathematical expectation and the variance of x_{t+1} , conditional on the information set, I_t . In (21.7), the only random variable on the RHS is P_{t+1} . This randomness, from the perspective of time t , is caused by the misperception of period $t + 1$ noise traders as reflected in the term $\rho_{t+1} \sim N(\rho^*, \eta^2)$. Thus, the expected wealth of arbitrageurs is

$$\bar{w}_{t+1}^a = ry + \lambda_t^a [r + \bar{P}_{t+1} - P_t(1 + r)], \quad (21.8)$$

and the variance of their wealth is

$$\sigma_{w_{t+1}^a}^2 = (\lambda_t^a)^2 \sigma_{P_{t+1}}^2, \quad (21.9)$$

where $\bar{P}_{t+1} = E[P_{t+1} | I_t]$ and

$$\sigma_{\bar{P}_{t+1}}^2 \equiv E[(P_{t+1} - \bar{P}_{t+1})^2 | I_t], \quad (21.10)$$

is the variance of the period $t + 1$ price, conditional on the period t information set.

Now consider the second group of investors at time t , the young noise traders. They operate in a similar manner to the arbitrageurs, except that they misperceive the price P_{t+1} by an amount $\rho_t \geq 0$, due to investor sentiment. Recall that ρ_t is known at the end of period t , so it is already contained in the information set of noise traders at time t . Hence, analogous to (21.6), noise traders (mis)perceive their wealth to be

$$w_{t+1}^n = (y - \lambda_t^n P_t)(1 + r) + \lambda_t^n(r + \rho_t + P_{t+1}) - y. \quad (21.11)$$

Note that in (21.11), the actual realization of the random term ρ_t is also denoted by ρ_t in order to conserve on notation. Proceeding as above, noise traders compute their expected wealth to be

$$\bar{w}_{t+1}^n = ry + \lambda_t^n [r + \rho_t + \bar{P}_{t+1} - P_t(1 + r)], \quad (21.12)$$

and their computed variance of wealth is analogous to (21.9), $\sigma_{w_{t+1}^n}^2 = (\lambda_t^n)^2 \sigma_{P_{t+1}}^2$, because ρ_t is known at the end of period t .

7. (Preferences) Noise traders and arbitrageurs have identical *constant absolute risk aversion* (CARA) preferences. Since there is no consumption when young, the utility of any individual depends only on their wealth when they are old. This utility is given by

$$U(w_{t+1}) = -\exp(-2\gamma w_{t+1}),$$

where $2\gamma > 0$ is the coefficient of absolute risk aversion, i.e., $-U''/U'$.

We now use a well-known result in statistics. If x is normally distributed with mean $\bar{x}$ and variance σ_x^2 , then

$$E[-\exp(-\gamma x)] = -\exp\left(-\gamma\left(\bar{x} - \frac{\gamma}{2}\sigma_x^2\right)\right), \gamma > 0.$$

It turns out that w_{t+1} is normally distributed; this arises because the misperception term ρ_{t+1} that influences P_{t+1} in (21.11) is normally distributed. Thus, we can compute the expected utility of an individual with random wealth w_{t+1} and the CARA utility function explicitly as follows.

$$EU(w_{t+1}) = E[-\exp(-2\gamma w_{t+1})] = -\exp\left(-2\gamma\left(\bar{w}_{t+1} - \frac{1}{2}2\gamma\sigma_{w_{t+1}}^2\right)\right). \quad (21.13)$$

Notice that the expectation operator is eliminated from the RHS; we assume that $EU > 0$, otherwise demand for the risky asset collapses. Maximizing $EU(w_{t+1})$ is formally equivalent to maximizing $\bar{w}_{t+1} - \gamma\sigma_{w_{t+1}}^2$. Using (21.8), (21.9), an arbitrageur, $i = a$, maximizes the following objective.

$$\max_{\langle \lambda_t^a \in [0,1] \rangle} EU(\bar{w}_{t+1}^a) \Leftrightarrow \max_{\langle \lambda_t^a \in [0,1] \rangle} \left[ry + \lambda_t^a [r + \bar{P}_{t+1} - P_t(1+r)] - \gamma (\lambda_t^a)^2 \sigma_{P_{t+1}}^2 \right]. \quad (21.14)$$

On account of (21.5), the noise trader misperceives the return on the risky asset. Using (21.12), (21.9) we can derive an expression analogous to (21.14) for the noise trader.

$$\max_{\langle \lambda_t^n \in [0,1] \rangle} EU(\bar{w}_{t+1}^n) \Leftrightarrow \max_{\langle \lambda_t^n \in [0,1] \rangle} \left[ry + \lambda_t^n [r + \bar{P}_{t+1} - P_t(1+r) + \rho_t] - \gamma (\lambda_t^n)^2 \sigma_{P_{t+1}}^2 \right]. \quad (21.15)$$

This completes the description of the model. The optimization problems in (21.14), (21.15) are unconstrained concave optimization problems. From the first order conditions, we can find λ_t^a, λ_t^n .

$$\lambda_t^a = \frac{r + \bar{P}_{t+1} - P_t(1+r)}{2\gamma \sigma_{P_{t+1}}^2}. \quad (21.16)$$

$$\lambda_t^n = \frac{r + \bar{P}_{t+1} - P_t(1+r)}{2\gamma \sigma_{P_{t+1}}^2} + \frac{\rho_t}{2\gamma \sigma_{P_{t+1}}^2}. \quad (21.17)$$

The optimal holding of the risky asset can be positive or negative, i.e., $\lambda_t^i \geq 0$. In particular, $\lambda_t^i < 0$ implies that the investor engages in short selling (or shorting) of the asset. From (21.16), (21.17), the following three comparative static effects are obvious.

1. Each type of investor invests more in the risky asset if its expected return, $r + \bar{P}_{t+1}$, is high.
2. Each type of investor invests less in the risky asset if the conditional variance of the next period price of the risky asset, $\sigma_{P_{t+1}}^2$, is high. Risk averse investors dislike fluctuations in the capital gains of the assets they hold. Higher risk aversion (higher γ) magnifies this effect.
3. Noise traders invest more in the risky asset, the greater the realization of the noise term, ρ_t , i.e., the greater is their degree of bullishness.

To find the price of the risky asset in period t , P_t , substitute the demands for the risky asset in (21.16), (21.17) into the market clearing condition (21.4) to get,

$$P_t = \frac{1}{1+r} \left[r + \bar{P}_{t+1} - 2\gamma \sigma_{P_{t+1}}^2 + \mu \rho_t \right]. \quad (21.18)$$

The model is difficult to solve because of the presence of the expectation term, $\bar{P}_{t+1}$, on the right hand side. However, by restricting attention to steady states we can make progress. In the steady state, the conditional expectations of future prices and their conditional variances are time invariant,

$$\bar{P}_{t+1} = \bar{P}_{t+j}, \text{ and } \sigma_{P_{t+1}}^2 = \sigma_{P_{t+j}}^2; j = 2, 3, \dots \quad (21.19)$$

Rewrite (21.18) at $t+1$

$$P_{t+1} = \frac{1}{1+r} \left[r + \bar{P}_{t+2} - 2\gamma \sigma_{P_{t+2}}^2 + \mu \rho_{t+1} \right]. \quad (21.20)$$

Hence, the conditional variance of the one-period-ahead price in the steady state is

$$\sigma_{\bar{P}_{t+1}}^2 = \left(\frac{\mu}{1+r} \right)^2 \eta^2. \quad (21.21)$$

Take expectations on both sides of (21.20), and using (21.5), (21.19), we get $\bar{P}_{t+1} = \frac{1}{1+r} [r + \bar{P}_{t+1} - 2\gamma\sigma_{\bar{P}_{t+1}}^2 + \mu\rho^*]$. Solving out for $\bar{P}_{t+1}$, we get

$$\bar{P}_{t+1} = \frac{1}{r} [r - 2\gamma\sigma_{\bar{P}_{t+1}}^2 + \mu\rho^*]. \quad (21.22)$$

Substituting (21.21), (21.22) into (21.18) we get the steady state solution to the price of the risky asset in period t ,

$$P_t = 1 + \frac{\mu(\rho_t - \rho^*)}{1+r} + \frac{\mu\rho^*}{r} - \frac{2\gamma}{r} \left(\frac{\mu}{1+r} \right)^2 \eta^2. \quad (21.23)$$

We can infer the following insights from (21.23).

1. If the noise traders behaved like the arbitrageurs (i.e., $\rho_t = 0$, $\rho^* = 0$ and $\eta^2 = 0$), then (21.23) implies that the price of the risky asset equals its fundamental value, i.e., $P_t = 1$. In this case, the risky and the riskless assets become perfect substitutes, so the market values them identically.
2. There is no fundamental risk in the model. However, variations in the degree of misperception of the noise traders, as captured by the size of $\rho_t - \rho^*$ (second term in (21.23)), determines the variation in the price of the risky asset in the steady state. If young noise traders at time t are bullish (high ρ_t), then the price of the risky asset is higher. The mechanism works like a self-fulfilling prophesy. From (21.17), noise traders invest more in the risky asset if they believe that the future price of the risky asset will be higher, and this is directly determined by the realization of ρ_t ; this increased demand, in turn, pushes up the actual price of the risky asset. Conversely, if they are bearish (low ρ_t), then the price of the risky asset is lower. As the second term in (21.23) shows, the impact of the degree of bullishness and bearishness of the noise traders is increasing in the fraction of noise traders, μ .
3. In deciding how much of the risky asset to hold, the young in period t take into consideration, the return $r + P_{t+1}$. The period $t + 1$ price, P_{t+1} , is expected to be higher if the average bullishness of period $t + 1$ noise traders, ρ^* , is high. In this case, the young in period t are willing to pay a higher price, P_t , for the risky asset (third term in (21.23)) because they can sell their risky asset in period $t + 1$ at a higher price.
4. Noise traders might become more or less bullish in period $t + 1$ (recall that $\rho_{t+1} \sim N(\rho^*, \eta^2)$, $\forall t$). This creates uncertainty in the minds of investors in period t about the price of the risky asset, P_{t+1} , and its returns, $r + P_{t+1}$. The fourth term in (21.23) takes account of this risk; the size of this term is directly proportional to the degree of risk aversion of the trader, as captured by the size of γ . Risk averse investors are compensated by a reduction in the current price of the risky asset that is directly proportional to the variance, η^2 , of the misperception term.

We now analyze the difference in the expected relative returns of the noise traders and the arbitrageurs, Δ^e . From (21.8), (21.12), and noting that the actual return on the risky asset is r and

not $r + \rho_t$, the “actual” difference in expected wealth levels of the noise traders and the arbitrageurs is given by

$$\Delta^e = (\lambda_t^n - \lambda_t^a) [r + \bar{P}_{t+1} - P_t(1 + r)]. \quad (21.24)$$

Using (21.16), (21.17), we get

$$\lambda_t^n - \lambda_t^a = \frac{\rho_t (1 + r)^2}{2\gamma\mu^2\eta^2}. \quad (21.25)$$

From (21.25), noise traders invest relatively more in the risky asset if they are optimistic ($\rho_t > 0$). Using (21.22), (21.23), the expected excess return from holding the risky asset relative to the safe asset is given by

$$r + \bar{P}_{t+1} - P_t(1 + r) = \frac{2\gamma\mu^2\eta^2}{(1 + r)^2} - \mu\rho_t. \quad (21.26)$$

In the absence of any noise trader misperception ($\rho_t = 0$, $\eta^2 = 0$), the RHS of (21.26) is zero, so there is no excess return on the risk asset. From (21.24), when expected excess return from holding the risky asset is positive, then by holding relatively more of the risky asset ($\lambda_t^n > \lambda_t^a$), noise traders ensure that their returns are higher, i.e., $\Delta^e > 0$.

Substitute (21.25), (21.26) in (21.24) to get

$$\Delta^e = \rho_t - \frac{\rho_t^2 (1 + r)^2}{2\gamma\mu\eta^2}. \quad (21.27)$$

It is quite possible that noise traders have higher expected returns than arbitrageurs, i.e., $\Delta^e > 0$.⁹ It is clear that a necessary condition for this to happen is that $\rho_t > 0$. When $\rho_t > 0$, noise traders are bullish, so they invest relatively more than arbitrageurs in the risky asset ($\lambda_t^n > \lambda_t^a$). But for the risk undertaken by noise traders to pay off, the risky asset should earn excess returns relative to its fundamentals (the riskless asset always earns its fundamental return). This, in turn, requires that the RHS of (21.26) be positive.

The extent of excess return on the risky asset depends on the variance of ρ_{t+1} , which is given by η^2 . The reason is that if there is considerable uncertainty about the beliefs of noise traders at time $t + 1$, then risk averse arbitrageurs at time t wish to curtail their investment in the risk asset. Noise traders are also risk averse, but their bullishness ($\rho_t > 0$) counteracts their risk aversion to invest in the risky asset. On net, it is possible that the price of the risky asset at time t , P_t (see (21.23)), is reduced. Higher risk aversion of the arbitrageurs (higher γ) magnifies this effect, hence, increasing the excess return on the risky asset.

Thus, paradoxically, in an environment of considerable uncertainty and strong investor sentiment in a direction that is contrary to the fundamentals, noise traders may outperform rational arbitrageurs. It is important, in this argument, that noise traders be bullish, but not too bullish. For if they are too bullish, they end up increasing P_t (see second term in (21.23)) and depressing the return on the risky asset. It is for this reason that for a high enough ρ_t , $\Delta^e < 0$. Conversely when noise traders are bearish, $\rho_t < 0$, they can make a considerable loss. The reason is that on

⁹ Hirshleifer et al. (2003) also give a similar theoretical result.

account of noise trader risk, the risky asset always earns a premium in the presence of risk averse investors. If $\rho_t < 0$, noise traders invest less in the risky asset as compared to the arbitrageurs and miss out on the premium, hence, their returns are relatively lower.

As a final comment, taking expectations on both sides of (21.27), and noting that $E\rho_t^2 = \rho^{*2} + \eta^2$ we get

$$E\Delta^e = \rho^* - \frac{\rho^{*2}(1+r)^2 + \eta^2(1+r)^2}{2\gamma\mu\eta^2}. \quad (21.28)$$

From (21.28), a necessary condition for the average return of noise traders to be relatively higher is that mean misperception is positive, i.e., $\rho^* > 0$. The first term captures the positive effect of bullishness of the beliefs of noise traders in enhancing their profits relative to the arbitrageurs. However, more bullish noise traders buy more of the risky asset, driving up the price P_t , which reduces the return on the risky asset and their relative profits (second term). The fact that under these conditions, the average expected returns of noise traders could be relatively higher, predisposes them to outlast the rational arbitrageurs; exactly the opposite conclusion to that advocated in all principles of economics courses.

21.4 The limits to arbitrage

The vast bulk of arbitrage activity in financial markets is not undertaken by dispersed small investors. Rather, it is undertaken by a relatively small number of professional mutual and pension funds managers who manage portfolios on behalf of millions of dispersed investors. We now consider the implications of this *separation between brains and capital for financial markets* in a simple three-period, partial equilibrium model (Shleifer and Vishny, 1997). Figure 21.4 gives a concise bird's eye view of the model that we now explain in greater detail.

1. (*Basic structure*) There are three time periods, $t = 1, 2, 3$. There is a single perfectly divisible asset whose supply is normalized to 1 unit, and whose fundamental value in any time period




Figure 21.4 A basic description of the structure of the three-period model under professional arbitrage.

is $V > 0$. The price of the asset at time t is p_t . In time period $t = 3$, the price of the stock always equals its fundamental value, i.e., $p_3 = V$. Hence, there is no fundamental risk in the economy.

2. (*Types of investors*) There are three kinds of investors who create a demand for the single asset. (1) Noise traders. (2) Arbitrageurs, whose objective is to maximize period 3 profits, or revenues, because there are no costs. (3) Small investors.

The arbitrageurs do not invest their own funds; one may think of them as professional managers of mutual and pension funds who invest the funds of small investors. Noise traders are of the *pessimistic type*. In each of the first two periods, they think that the value of the single asset is $V - S_t$, $t = 1, 2$. The arbitrageurs know and observe S_1 , but they are uncertain about S_2 , which has the interpretation of a shock to noise trader sentiment in period 2. With probability q , S_2 takes a value $S > S_1$, i.e., the pessimistic misperceptions become worse. With the complementary probability $1 - q$, $S_2 = 0$, i.e., all misperception disappears.

Arbitrageurs and small investors are risk neutral and fully rational (although, small investors rely on the arbitrageurs to invest their funds, perhaps because they do not understand the structure of the problem).

3. (*Direct demands for the asset*) We take a reduced-form approach and directly specify demands for noise traders and the arbitrageurs. The direct monetary, or nominal, demand of noise traders in the first two time periods is

$$D_t^n = V - S_t, t = 1, 2. \quad (21.29)$$

Thus, the real demand of noise traders (in units of the asset) equals

$$\frac{D_t^n}{p_t} = \frac{V - S_t}{p_t}, t = 1, 2. \quad (21.30)$$

Arbitrageurs depend solely on the resources gathered from small investors in order to demand the asset. Denote the cumulative funds at the disposal of the arbitrageurs in time period $t = 1, 2$ by F_t . We assume that $F_1 < S_1$ is exogenously given, while F_2 is endogenous and depends on the decisions made by arbitrageurs in period 1. The endogenously determined nominal demand of arbitrageurs for the asset in period 1, D_1^a , must satisfy

$$D_1^a \leq F_1. \quad (21.31)$$

If $D_1^a = F_1$, then arbitrageurs are said to be *fully invested*. However, arbitrageurs could be worried that the pessimism of noise traders might get worse in period 2. In this case, if they are carrying the asset into period 2, they might have to sell it at a lower price. And so, they might be better-off hedging their portfolio and carrying ahead some “cash” into the second period. The cash can then be used in period 2 to buy the asset whose price at $t = 3$ is known with certainty ($p_3 = V$). Thus, if $D_1^a < F_1$, we say that investors are *not fully invested*. In terms of units of the asset, the real demand for the asset by the arbitrageurs in period 1 is, thus, given by

$$\frac{D_1^a}{p_1}. \quad (21.32)$$

Using (21.30), (21.32) we can set the total demand for the asset $\frac{V-S_1}{p_1} + \frac{D_1^a}{p_1}$ equal to the total supply of 1 unit. This determines the price, p_1 , in the first period, i.e.,

$$\frac{D_1^n}{p_1} + \frac{D_1^a}{p_1} = 1 \Leftrightarrow p_1 = V - S_1 + D_1^a. \quad (21.33)$$

In period 2, there are two possibilities depending on the realization of S_2 .

(a) With probability $1 - q$, $S_2 = 0$, so all noise trader misperceptions disappear. Recall that noise trader misperceptions are only of the pessimistic type, and serve only to lower demand and, thus, the price of the asset. Thus, when $S_2 = 0$, the second period price, p_2 , is the highest possible and the returns to the arbitrageurs from selling the asset are the highest possible. Hence, they liquidate their positions at a gain and hold cash between periods 2 and 3. Thus, second period resources, F_2 , of the arbitrageurs do not contribute anything towards second period demand of the asset, and only the noise traders demand the asset in period 2, thus,

$$1 = \frac{D_2^n}{p_2} = \frac{V}{p_2} \\ \Rightarrow p_2 = V.$$

(b) With probability q , $S_2 = S > S_1$, so the noise traders' pessimistic misperceptions deepen further. From (21.30), the demand by noise traders for the asset falls, so p_2 is low (in the absence of any demand by arbitrageurs, $p_2 = V - S < V$). By assumption $p_3 = V$, so the return to arbitrageurs from investing an extra unit of money in the asset exceeds that from holding cash. We make the assumption that in any period, the total resources available to the arbitrageurs to invest in the risky asset are restricted by the following condition.

$$F_1 < S_1; F_2 < S_2. \quad (21.34)$$

Condition (21.34) ensures that the arbitrageurs never have enough resources to completely negate the mispricing created by noise traders in any of the periods. So in the presence of noise traders' misperceptions, it is always the case that $p_1 \leq V$ (see (21.31), (21.33)) and $p_2 \leq V$ (see (21.35), below). So long as $p_2 \leq V$, the arbitrageurs will put all their second period resources, F_2 , into buying the asset. Thus, equating demand to the one unit supply gives the second period price of the asset as follows:¹⁰

$$1 = \frac{D_2^n}{p_2} + \frac{F_2}{p_2} \Leftrightarrow p_2 = V - S + F_2. \quad (21.35)$$

The condition $F_2 < S_2$ ensures that $p_2 \leq V$.

4. (*Performance based arbitrage*) Small investors ask the arbitrageurs to invest their funds because they do not have the necessary expertise or specialized knowledge to understand the working of financial markets. Furthermore, the arbitrageurs are informed about the

¹⁰ We allow short selling of assets. It is possible that some part of the asset is held by individual investors between periods 1 and 2 and these investors choose not to sell the asset in period 2. However, short selling does not require physical ownership of the asset. Thus, the total availability of the asset is one unit in period 2.

fundamental value of the asset, while the small investors are not. In the absence of any specialized knowledge of either the stock market or the fundamental value of the stock, small investors infer the quality of the arbitrageur simply by the past performance of the arbitrageur.¹¹ This is referred to as *performance based arbitrage* (PBA).

PBA can lead to perverse effects. Suppose, for instance, that noise traders' misperception worsens at $t = 2$, i.e., $S_2 = S$ increases. From (21.35), p_2 falls. Hence, the (gross) return on the period 1 investment made by arbitrageurs, $\frac{p_2}{p_1}$, falls; the net return, r , is $r = \frac{p_2 - p_1}{p_1} = \frac{p_2}{p_1} - 1$, so the gross return $r_g = 1 + r = \frac{p_2}{p_1}$. Observing a fall in the return offered by the arbitrageur between periods 1 and 2, PBA implies that in period 2 small investors withdraw funds from the arbitrageur, reducing the resources, F_2 , available to arbitrageurs to buy the asset. This, paradoxically, happens at a time when the gross return on the asset between periods 2 and 3, $\frac{p_3}{p_2}$, is the highest (because p_3 is fixed and p_2 takes the lowest possible value).

PBA is implemented as follows. The gross return r_g generated by the arbitrageur with borrowed funds from the small investors, between periods 1 and 2, on one dollar, is defined as

$$r_g = \frac{D_1^a}{F_1} \times \frac{p_2}{p_1} + \frac{F_1 - D_1^a}{F_1} \times 1. \quad (21.36)$$

Equation (21.36) is easy to understand. A fraction $\frac{D_1^a}{F_1}$ of the first period resources are invested by the arbitrageur in the asset, on which there is a gross return equal to $\frac{p_2}{p_1}$. The remaining fraction $\frac{F_1 - D_1^a}{F_1}$ is carried into period 2 as cash on which there is no net return. The supply of resources available to the arbitrageur at time $t = 2$, F_2 , is assumed to be given by the following functional relation, $F_2 = \phi(F_1 r_g)$; $\phi' > 1$, $\phi'' \leq 0$. For our purposes, a simple linear form will do, so let $\phi(x) = ax + (1 - a)$ where $a > 1$, so

$$F_2 = aF_1 r_g + (1 - a); a > 1. \quad (21.37)$$

Thus, a greater return between periods 1 and 2, r_g , leads to an increase in funds at the disposal of the arbitrageurs, F_2 , in period 2. These features capture PBA in the model. Substituting r_g from (21.36) in (21.37) we get

$$F_2 = a \left(D_1^a \times \frac{p_2}{p_1} + F_1 - D_1^a \right) + (1 - a). \quad (21.38)$$

5. (*Optimization problem of arbitrageurs*) The objective of the arbitrageurs at time $t = 1$ is to maximize the expected profits at $t = 3$. With no fixed costs of operation of the arbitrageurs, this is equivalent to maximizing period 3 expected revenues. Now consider the two possible events in period 2.

(a) With probability $1 - q$, $S_2 = 0$, $p_2 = V$ and arbitrageurs liquidate their position in the asset and only hold cash between periods 2 and 3. Hence, the expected funds at the disposal of investors at $t = 3$, $F_3|_{S_2=0}$, can be found by substituting $p_2 = V$ in (21.38) to get,

¹¹ See Shleifer and Vishny (1997) for why the use of signaling devices to generate a separating equilibrium, which is revealing of the arbitrageur's quality, is not plausible.

$$F_3 |_{S_2=0} = a \left(D_1^a \times \frac{V}{p_1} + F_1 - D_1^a \right) + (1-a). \quad (21.39)$$

(b) With the complementary probability, q , $S_2 = S > S_1$, and noise trader misperceptions deepen. Second period price p_2 drops to its lowest level and with a fixed $p_3 = V$, the return on holding a unit of the asset between periods 2 and 3 is the highest. Hence, the arbitrageurs put all their second period resources, F_2 , into buying the asset. With nominal funds, F_2 , arbitrageurs buy $\frac{F_2}{p_2}$ units of the asset at $t = 2$. At $t = 3$, each of these units sells for a guaranteed price of $p_3 = V$; this gives a total revenue of $\frac{V}{p_2} F_2$. Using (21.38) to substitute for F_2 (and noting that $p_2 \neq V$), the expected third period revenues, in this case, are,

$$F_3 |_{S_2=S} = \frac{V}{p_2} \left(a \left(D_1^a \times \frac{p_2}{p_1} + F_1 - D_1^a \right) + (1-a) \right). \quad (21.40)$$

Using (21.39), (21.40), the unconditional expected third period revenue of the arbitrageurs is $R = (1-q)F_3 |_{S_2=0} + qF_3 |_{S_2=S}$, or

$$R = (1-q) \left(a \left(D_1^a \times \frac{V}{p_1} + F_1 - D_1^a \right) + (1-a) \right) + q \frac{V}{p_2} \left(a \left(D_1^a \times \frac{p_2}{p_1} + F_1 - D_1^a \right) + (1-a) \right). \quad (21.41)$$

The optimization problem of the arbitrageur is to maximize (21.41) with respect to D_1^a subject to the constraint that¹²

$$D_1^a \leq F_1. \quad (21.42)$$

Since $D_1^a \leq F_1$ and from (21.34) $F_1 < S_1$, thus, from (21.33), $p_1 < V$. With probability $1-q$, $S_2 = 0$, $p_2 = V$. Thus, in this case, $\frac{V}{p_1} - 1 > 0$. With probability q , the misperceptions worsen to $S_2 = S > S_1$ and p_2 drops to its lowest level.

Differentiating (21.41) with respect to D_1^a , we get

$$\frac{\partial R}{\partial D_1^a} = a(1-q) \left(\frac{V}{p_1} - 1 \right) + aq \left(\frac{p_2}{p_1} - 1 \right) \frac{V}{p_2}. \quad (21.43)$$

An increase in the first period demand for the asset by the arbitrageurs, D_1^a , creates two opposing effects on third period revenues. These are captured, respectively, by the first and second terms on the RHS in (21.43). The first term captures the state of the world where noise trader misperception clears up ($S_2 = 0$), in which case the arbitrageurs sell off their asset holdings and make a profit; this is the positive marginal impact of increasing D_1^a . The second term captures the complementary state of the world in which misperceptions worsen ($S_2 = S$), and $\frac{p_2}{p_1} < 1$. In this state, due to PBA, small investors withdraw some of their funds from the arbitrageurs. The second term captures the opportunity cost of these lost funds precisely when the return on these funds is very high; this is the negative marginal impact of increasing D_1^a .

The RHS of (21.43) depends on D_1^a through p_1 . For a corner solution, $D_1^a = F_1$, we require that $\frac{\partial R}{\partial D_1^a} > 0$, thus,

¹² We allow $D_1^a < 0$ through the usual process of shorting an asset.

$$(1 - q) \left(\frac{V}{p_1} - 1 \right) + q \left(\frac{p_2}{p_1} - 1 \right) \frac{V}{p_2} > 0 \Leftrightarrow D_1^a = F_1. \quad (21.44)$$

For the case $D_1^a < F_1$, we need $\frac{\partial R}{\partial D_1^a} = 0$, thus,

$$(1 - q) \left(\frac{V}{p_1} - 1 \right) + q \left(\frac{p_2}{p_1} - 1 \right) \frac{V}{p_2} = 0 \Leftrightarrow D_1^a < F_1. \quad (21.45)$$

A necessary condition for (21.45) to hold is that the second term is negative, i.e., $p_2 < p_1$ so $\frac{p_2}{p_1} - 1 < 0$.

What factors are important for the arbitrageur to choose to be *partially invested* ($D_1^a < F_1$) rather than *fully invested* ($D_1^a = F_1$)? Comparing (21.44), (21.45), and assuming that $\frac{p_2}{p_1} - 1 < 0$, the necessary conditions are the following.

1. High probability of worsening misperceptions in period 2, i.e., high q .
2. Small misperception shock, S_1 , in period 1. From (21.33), this increases p_1 relative to V , so $\frac{V}{p_1}$ is low.
3. S should be relatively high so that p_2 is low; this ensures that $\frac{p_2}{p_1}$ is not too high.

At the other extreme, the arbitrageurs are fully invested ($D_1^a = F_1$); they bet on a price recovery in period 2, following the slump in period 1. This is the case that we predominantly want to focus on below. Substituting $D_1^a = F_1$ in (21.38) we get that

$$F_2 = a \left(F_1 \frac{p_2}{p_1} \right) + 1 - a. \quad (21.46)$$

We can summarize these results in a proposition and leave the proof for the reader.

Proposition 21.8 (Shleifer and Vishny, 1997): (a) There exists a value q^* such that if $q < q^*$, then $D_1^a = F_1$, and if $q > q^*$, then $D_1^a < F_1$.
 (b) When the arbitrageurs are fully invested, $D_1^a = F_1$, then $\frac{dp_1}{dS_1} < 0$, $\frac{dp_2}{dS} < 0$, and $\frac{dp_1}{dS} = 0$.
 However, when not fully invested, $D_1^a < F_1$, then $\frac{dp_1}{dS_1} < 0$, $\frac{dp_2}{dS} < 0$, and $\frac{dp_1}{dS} < 0$.

In Proposition 21.8(b), the main difference between the two cases is the sign of $\frac{dp_1}{dS}$. When not fully invested, an expectation of greater noise trader misperception in period 2 reduces the period 2 price and reduces the return from holding the asset between periods 1 and 2. This reduces period 1 demand of the asset on the part of arbitrageurs, hence reducing the period 1 price, $\frac{dp_1}{dS} < 0$.

From (21.35), we know that $p_2 = V - S_2 + F_2$, hence,

$$p_2 = V - S_2 + a \left(F_1 \frac{p_2}{p_1} \right) + 1 - a \equiv f(p_2); f' > 0, f'' = 0. \quad (21.47)$$

In (21.47), a fixed point of the relation $p_2 = f(p_2)$ is given by $p_2^* = f(p_2^*)$. Stability requires that $f'(p_2^*) = a \frac{F_1}{p_1} < 1$. We illustrate the situation in Figure 21.5, for two cases: $f' < 1$ (curve I) and $f' > 1$ (curve II). There are two fixed points corresponding to each of these curves, p_2^* and p_2^{**} ; p_2^* is stable while p_2^{**} is unstable.




Figure 21.5 Stability analysis in the Shleifer and Vishny (1997) model.

Source: Shleifer, A. and Vishny, R. W. (1992). "Liquidation values and debt capacity: a market equilibrium approach." *The Journal of Finance* 47: 1343–66. © American Finance Association. With permission from John Wiley & Sons.

At the unstable equilibrium, $f'(p_2^{**}) > 1$. Even the slightest decrease in price to any $p_2 < p_2^{**}$ will push the second period price down all the way to $V - S_2$ (which is assumed to be strictly positive in the diagram). The intuition is that in this case, second period prices are too low, thus, the returns on investment between periods 1 and 2 are so low that small investors start withdrawing their resources from the arbitrageurs (this is due to PBA). As this happens, second period price falls further, another round of withdrawals takes place and so on, until the arbitrageurs are left with no funds to invest. Unlike the neoclassical assumptions, there is a complete failure of arbitrage. The only demand for the asset in this case comes from noise traders who demand $V - S_2$.

At the stable equilibrium, $p_2^* > V - S_2$, the arbitrageur does not run out of resources because the second period price is high enough, and gross returns on investment are adequate for small investors who rely on PBA to invest at least some resources with the arbitrageurs.

Implicitly differentiate both sides of (21.47) with respect to S_2 to get,

$$\frac{\partial p_2}{\partial S_2} = \frac{-1}{1 - f'} = \frac{-1}{1 - a \frac{F_1}{p_1}}.$$

At the stable equilibrium, $p_2 = p_2^*$, we know that $f'(p_2) = a \frac{F_1}{p_1} < 1$, hence, $\frac{\partial p_2}{\partial S_2} < -1$. Thus, following a worsening of noise trader misperceptions (increase in S_2), prices fall more than one to one. This occurs because PBA ensures that small investors reduce the flow of funds to arbitrageurs, which leads to a further collapse of demand for the asset. From (21.47), we know that in the absence of any demand for the asset by the arbitrageurs, $p_2 = V - S_2$, so p_2 should fall one to one with a worsening of noise trader misperceptions.

The claim of classical finance theory has been that the role of irrational noise traders can be neglected in the presence of rational arbitrageurs because the arbitrageurs will take trading positions that nullify irrational activity. In other words, classical finance assumes that $\frac{\partial p_2}{\partial S_2} = 0$. However, we have just shown that the presence of these rational arbitrageurs coupled with the

plausible assumption of PBA actually makes matters worse because prices move in the wrong direction by even more (i.e., more than one to one). The reason is that as prices fall, gross return on the investment falls between periods 1 and 2, so small investors start withdrawing their money. When this run on the financial resources of the arbitrageurs takes place, they start to liquidate their positions in the asset.¹³ But as they do so, demand for the asset collapses further and its price falls even more. Hence in extreme situations, the arbitrageurs can be ineffective in bringing prices back to their fundamental values.¹⁴

One implication of these results is that newly formed hedge funds are less likely to succeed in the marketplace relative to older, established ones. The reason is that PBA would bite more for the younger funds (higher value of a), because small investors are unsure of their quality. Hence, as noise trader misperceptions worsen, and gross rates of returns on assets decline, the run on younger funds would be far greater relative to that on established funds. Thus, paradoxically, newer and younger funds lose money exactly at the time when expected returns on assets are very high, but average returns in the past are lower.

The financial events involving the Long Term Capital Management (LTCM) fund in 1998 are also indicative of the mechanism outlined in this section. Russia defaulted on its ruble denominated debt in 1998. This eventually created events in the West that look similar to a large S_2 shock in the model (Shleifer, 2000). The result was that arbitrageurs such as LTCM lost funds precisely when the expected returns on these funds were very high. In the event, LTCM was bailed out by the US Federal Reserve Bank for a temporary reprieve, but many other arbitrageurs were less fortunate.

21.5 Gradual flow of information

We have noted above, the very large trading volume in stock markets that poses a puzzle for EMH. One possible explanation of too large a trading volume is the presence of heterogeneity among investors about their prior beliefs. Despite receiving identical public information, different investors may interpret it differently due to underlying heterogeneity in beliefs about fundamentals, or the relevant economic model (Harris and Raviv, 1993; Kandel and Pearson, 1995). This potentially triggers far too much trade among investors even when the underlying fundamentals do not warrant it. This difference in beliefs can induce investors to trade far more than that warranted by the economic fundamentals.

Bubbles and stock market booms are often preceded by excessive trading volumes that appear unlikely to be driven by economic fundamentals. Examples include the internet bubble during 1997–2000 and the stock market boom of the late 1920s; for fascinating accounts, see Ofek and Richardson (2003), Galbraith (1997), and Kindleberger and Aliber (2005). A leading explanation is that during these periods, there is disagreement among the beliefs of market participants about the nature of future price movements.

An illuminating example of these disparate beliefs is given by Huberman and Regev (2001). They look at the effect of news announcements about a biotechnology firm, *EntreMed*, on its stock prices. News that *EntreMed* had undertaken pathbreaking research into cancer treatment

¹³ Arbitrageurs would liquidate even more quickly if they were risk averse rather than risk neutral, as is assumed here. Hence, the conclusions are even more stark under risk aversion.

¹⁴ The result that in the presence of limited capital, arbitrageurs might liquidate their positions prematurely is also a feature of several other models; see, for instance, Shleifer and Vishny (1992).

appeared in the journal *Nature* in November 1997. The news also appeared in the newspaper *the Times* on page A28 (not a particularly salient page), as well as in other media. This resulted in an increase in the price of EntreMed stocks. However, five months later, on May 3, 1998, *The Times* carried out a front page story (a very salient page) on topical cancer research innovations in which EntreMed's contributions were noted prominently. However, this news was not new. Contrary to the predictions of EMH, the stock price of EntreMed rose from \$12/share to \$52/share within a day.

Hong and Stein (1999, 2007) suggest the following interpretation of this event. There are two kinds of investors, *specialists*, who read specialist scientific journals (such as *Nature*) and *generalists*, who avoid a serious scientific read, but keep a lookout for salient bits of news (such as the front page of *The Times*). The original news in *Nature*, read by the specialists, but not by the generalists, lead to a difference in beliefs about the future prices of EntreMed stocks. The attempts by the specialists to buy the EntreMed stock triggered the initial jump in the stock prices. An important aspect of the behavior of generalists is that they do not draw inferences from the buying activities of the specialists. They do not imagine that perhaps the specialists have some important bit of news that could be relevant for future stock prices. The generalists only wake up to the news when it is presented to them in a very salient form in a non-specialist publication, which, in this case, happened five months later.

If it is true that generalists only respond to very salient information, then the framing of information must play an important role in their decisions. An important corollary is that inattentive investors must respond less when they are distracted. DellaVigna and Pollet (2009) exploit this corollary to show that announcements of earnings on Friday generate less trading activity relative to announcements on other days. One plausible explanation is that *inattentive investors* get distracted by weekend activities, so they do not act sufficiently on the announcement when they next have an opportunity, on the following Monday.

Barber and Odean (2008) show that *attention grabbing stocks* that have experienced abnormal trading/returns or news coverage recently, trade more. Coverage of an earnings announcement in the local newspaper also boosts trading of the associated stock (Engelberg and Parsons, 2011). Overnight buy-sell recommendations in the TV show *Mad Money* also lead to greater trading when the viewership on the show is higher (Engleberg et al., 2011). For the Shanghai stock exchange, Seasholes and Wu (2007) find that there is greater trading in a stock if it hits an upper price limit (an attention grabbing event). In a similar vein, Hirshleifer et al. (2009) find that the immediate stock price and volume response to a firm's earnings surprise is weaker, and post-earnings announcement drift is stronger, when a greater number of earnings announcements by other firms are made on the same day.

Hong and Stein (1999) propose a model of gradual information transmission while Hong et al. (2000) provide the relevant empirical evidence. Hong and Stein (2007) give a simple stylized model that conveys the essentials of their ideas, which we now outline below. The main feature of the model is that biased investors are of two types; each type believes in its incorrect economic model of the world. The purpose of this simple model is to illustrate the following points.

1. Investor heterogeneity has an important impact on temporal prices and returns.
2. Momentum and trading volume are closely related.
3. Even when price equals fundamental value, there can be aggressive trading.

Consider a three-period model, with time $t = 1, 2, 3$. There is a single stock that pays a dividend equal to D at time $t = 3$. The interest rate is normalized to zero and there is no risk premium for holding the stock. The realized dividends at time $t = 3$ obey the "true" economic model,

$$D = s_1 + s_2, \quad (21.48)$$

where s_1, s_2 are two distinct, uncorrelated, public signals whose realization is known only at time $t = 2$. At time $t = 1$, the signals are just known to have zero means and respective variances σ_1^2, σ_2^2 .

There are two kinds of investors, each with an incorrect and incompatible model of the real world.

1. A fraction $\lambda \in [0, 1]$ of investors believes that the model describing dividends is $D = s_1$.
2. The remaining fraction, $1 - \lambda$, of investors believes that the model is $D = s_2$.

Furthermore, at time $t = 2$, the fraction λ only observes s_1 (but not s_2), while the fraction $1 - \lambda$ only observes s_2 (but not s_1). Often, in the short run, individuals observe what they are looking for and miss what they are not looking for because of limited attention and confirmation-bias. At time $t = 3$, everyone observes both signals as well as the true model given in (21.48). The model, thus, incorporates both gradual information transmission, and limited attention. In period 3, everyone learns the true model and observes the actual realizations of the signals, s_1, s_2 .

At time t , the price of the stock is given by P_t . Each investor is willing to pay a price for the stock that equals the expected dividend. The aggregate price P_t is simply a weighted average of the prices expected by the two kinds of investors. Thus, letting E denote the expectation operator with respect to the information set at time $t = 1$,

$$P_1 = \lambda E[s_1] + (1 - \lambda)E[s_2] = 0. \quad (21.49)$$

At time $t = 2$, s_1 is observed by a fraction λ and s_2 by a fraction $1 - \lambda$ of investors, hence,

$$P_2 = \lambda s_1 + (1 - \lambda)s_2. \quad (21.50)$$

At time $t = 3$, s_1, s_2 is observed by all investors and the price of the stock is simply equal to the dividend, hence,

$$P_3 = D = s_1 + s_2. \quad (21.51)$$

Using (21.49), (21.50), the monetary return, R_2 , from holding the stock between periods 1 and 2 is

$$R_2 = P_2 - P_1 = \lambda s_1 + (1 - \lambda)s_2. \quad (21.52)$$

Using (21.50), (21.51), the monetary return, R_3 , between periods 2 and 3 is

$$R_3 = P_3 - P_2 = (1 - \lambda)s_1 + \lambda s_2. \quad (21.53)$$

At $t = 1$, $E[R_2] = E[R_3] = 0$ because $E[s_1] = E[s_2] = 0$. Hence, the perceived covariance at $t = 1$ between the returns R_2 and R_3 is given by

$$\text{Cov}(R_2, R_3) = E[R_2 R_3]. \quad (21.54)$$

Substituting R_2, R_3 from (21.52), (21.53) into (21.54) we get

$$\text{Cov}(R_2, R_3) = E[\lambda(1-\lambda)(s_1^2 + s_2^2)] = \lambda(1-\lambda)(\sigma_1^2 + \sigma_2^2). \quad (21.55)$$

We now give the following reasonable definition of momentum; momentum is higher if there is a higher correlation in returns.¹⁵

Definition 21.10 (*Momentum*): Momentum is defined by $\text{Cov}(R_2, R_3) = \lambda(1-\lambda)(\sigma_1^2 + \sigma_2^2)$.

It is reasonable to assume that, at time $t = 2$, investors engage in greater trading if the difference between the investor's own expectation of the period 3 price and the actual price in period 2 is higher; in this case the perceived expected returns are higher.

At time $t = 2$, a fraction λ of the investors expect that $P_3 = s_1$ (these investors observe s_1 at $t = 2$). The current price is, however, P_2 . If $P_2 \neq s_1$, then these investors see a trading opportunity. If $P_2 > s_1$, then such investors believe that $P_2 > P_3$, so they engage in short selling the asset.¹⁶ If $P_2 < s_1$, then they take a long position. Furthermore, the greater is $|P_2 - s_1|$, the greater is the perceived trading opportunity, thus, the greater is the volume of trade engaged in by these traders.

Analogously, a fraction $1 - \lambda$ of the investors expects that $P_3 = s_2$ (these investors observe s_2 at $t = 2$). They see a trading opportunity if $P_2 \neq s_2$. Averaged across the two types of investors, the trading volume is proportional to

$$V = \lambda |P_2 - s_1| + (1 - \lambda) |P_2 - s_2|. \quad (21.56)$$

Substituting P_2 from (21.50) into (21.56) we get

$$\begin{aligned} V &= \lambda |(1-\lambda)(s_2 - s_1)| + (1-\lambda) |\lambda(s_1 - s_2)|. \\ \Rightarrow V &= \lambda(1-\lambda) |s_1 - s_2|. \end{aligned} \quad (21.57)$$

Definition 21.11 (*Trading volume*): The trading volume is defined by $V = \lambda(1-\lambda) |s_1 - s_2|$.

We can now use this simple framework to derive our results.

1. Heterogeneity in beliefs is crucial to generate momentum in stock prices. From (21.55), momentum is maximized at the unique value $\lambda = 0.5$. Suppose, for example, that $\lambda = 1$. In this case, all investors mistakenly believe that the economic model for dividends is given by $D = s_1$. So the realization of s_2 at date $t = 2$ is ignored by all investors and does not affect P_2 , thus, $P_2 = s_1$. P_3 continues to be given by $P_3 = s_1 + s_2$. Since $P_1 = 0$, thus,

¹⁵ Other explanations for momentum have also been given. For instance, Daniel et al. (1998, 2001) give an explanation based on the distinction between private and public information. Investors are overconfident about private information (relative to public information) that they have worked to generate. Our section on behavioral corporate finance below develops this point. Overconfidence arising from private information can drive prices above fundamental values. Public information has an asymmetric effect. Confirmatory public news strengthens private information, however, contrary public news is slow to shake off wrongly held private beliefs (see attribution-bias in the behavioral corporate finance section below). *Ceteris paribus*, public news may gradually, over time, drive prices back towards fundamentals.

¹⁶ For a theoretical model of gradual flow of information in the presence of short sales constraints, see Palfrey and Wang (2012). An important result of this model is speculative overpricing and asymmetric reaction to good news and bad news.

- $R_2 = s_1$ and $R_3 = s_2$. Since s_1, s_2 are uncorrelated, so are R_2, R_3 and there is no momentum in stock prices. Analogously it is easy to check that when $\lambda = 0$, there is no momentum either. Hence, heterogeneity in beliefs is crucial to generate momentum in stock prices.
2. From (21.57), trading volume is (uniquely) maximized when $\lambda = 0.5$. Thus, trading volume is highest when momentum is highest. A testable prediction is that momentum should be highest for stocks whose trading volume is high; this is borne out by the empirical evidence (Lee and Swaminathan, 2000).
 3. Consider now the case where price equals fundamental value. So suppose that the realizations of s_1, s_2 at date $t = 2$ are such that $s_1 = -s_2$. From (21.48), (21.49), (21.51) we know that in this case $P_1 = P_3 = D = 0$. Hence, the prices at times $t = 1, 3$ are equal to the fundamental value of the stock. However, from (21.57), the trading volume is given by $V = 2\lambda(1 - \lambda)|-s_2| > 0$. Furthermore, V is increasing in the size of $|-s_2|$. In the special case of no heterogeneity, i.e., $\lambda = 0$ or $\lambda = 1$, $V = 0$. Thus, on account of investor heterogeneity (when $\lambda \neq 0$ or $\lambda \neq 1$), there can be large trading volumes, despite the fact that prices equal fundamental values ($P_1 = P_3 = D = 0$).

21.6 Stock market underreaction and overreaction

Suppose that there is a single investor who invests in a security that has a return r_t in period t . In every period, the investor hears some news, x_t , about the company that sells the security. The news can take two possible forms, good (G) and bad (B). Empirical evidence suggests that over short horizons (1–12 months) stock markets *underreact* to news. However, in the long run (3–5 years), stock markets *overreact* to consistent news pointing in the same direction. Thus, an important phenomenon in stock markets is *underreaction* in the short run and *overreaction* in the long run.¹⁷ In this section, we explain underreaction and overreaction by using the representativeness heuristic and conservatism-bias (Barberis et al., 1998). First we need to define the terms underreaction and overreaction formally, using the definitions in Barberis et al. (1998).

Definition 21.12 (*Underreaction*): Stock market underreaction takes place if following good news at time t , the next period return, r_{t+1} , is higher than that following bad news.

$$E[r_{t+1} | x_t = G] > E[r_{t+1} | x_t = B]. \quad (21.58)$$

In classical finance, and as reflected in the EMH, good or bad news is immediately reflected in stock market prices; there are no implications for prices or returns in the next period. Indeed, if the market anticipates such news, its effect is supposed to be already incorporated in the rates of return prior to period t . In contrast, from (21.58), stock markets do not react fully to the good news in period t . This mistake is reflected in a higher expectation of the next period return. This has implications for the EMH in that expectations are correlated with elements of the information set. In other words, stale information can help predict future returns. Thus, if investors are sophisticated, and are aware of (21.58), they can beat the market.

Definition 21.13 (*Overreaction*): Stock market overreaction takes place if following a sequence of good news, the expected future return is relatively lower as compared to that following a sequence of bad news. In other words,

¹⁷ There is extensive empirical evidence, see, for instance, De Bondt and Thaler (1985), Bernard and Thomas (1989), Cutler et al. (1991), Chopra et al. (1992), Fama and French (1992), Jegadeesh and Titman (1993), Lakonishok et al. (1994), La Porta (1996), and Chan et al. (1997).

Table 21.1 States and transition probabilities in the mean reverting model (M).

	$\varepsilon_{t+1} = x$	$\varepsilon_{t+1} = -x$
$\varepsilon_t = x$	p_L	$1 - p_L$
$\varepsilon_t = -x$	$1 - p_L$	p_L

Table 21.2 States and transition probabilities in the trending model (T).

	$\varepsilon_{t+1} = x$	$\varepsilon_{t+1} = -x$
$\varepsilon_t = x$	p_H	$1 - p_H$
$\varepsilon_t = -x$	$1 - p_H$	p_H

$$E[r_{t+1} | x_t = G, x_{t-1} = G, \dots, x_{t-j} = G] < E[r_{t+1} | x_t = B, x_{t-1} = B, \dots, x_{t-j} = B]. \quad (21.59)$$

For a given sequence of good news the market overreacts, so to take account of this, future returns are expected to fall. In this case, high past returns help to predict a lower future return, which again runs counter to the EMH.

To explain (21.58), (21.59), we use the following formal setup. There is a single risk-neutral representative investor whose discount rate is θ . A single security pays out all its earnings as dividends. The *true model* of earnings is a random walk, so the dividends at time t , D_t , are given by

$$D_t = D_{t-1} + \varepsilon_t; \quad \varepsilon_t = \pm x, \quad (21.60)$$

where ε_t is a random term. In any time period, ε_t takes on two possible values, $x > 0$ and $-x$ with equal probability.

The investor mistakenly believes that the true model is either (1) *mean reverting* (M), or (2) *trending* (T). In particular, the correct model, the random walk model, is not even considered to be a possibility by the investor.¹⁸ The investor wishes to figure out which of the two incorrect models best fits the sample. In all other respects, the investor obeys the classical laws of probability. Models M and T are both Markov processes that we now describe; see Tables 21.1 and 21.2. We assume that the transition probabilities p_L, p_H obey the restriction,

$$0 \leq p_L < 0.5 < p_H \leq 1. \quad (21.61)$$

In model M there is a relatively greater likelihood of the outcomes reversing from one period to the next (x followed by $-x$ or vice versa). This is the sense in which it is mean reverting. In model T , on the other hand, there is a relatively greater likelihood of the same outcome in two

¹⁸ In the model in Hong et al. (2007), the true model is multivariate, but investors use a univariate model. Economic agents keep track of errors they make in successive periods. Learning allows them to change from one incorrect model to another incorrect model.

Table 21.3 Markov transition probabilities for the two states of the world.

	$s_{t+1} = M$	$s_{t+1} = T$
$s_t = M$	$1 - \lambda_1$	λ_1
$s_t = T$	λ_2	$1 - \lambda_2$

consecutive periods (x followed by x or $-x$ followed by $-x$). In this sense, model T is a trending model.

The investor knows p_L, p_H and believes that the world potentially switches between the two models in any two contiguous time periods. In particular, investors believe that there is a Markov transition *between the two models* in any two consecutive periods. Denote by $s_t = M, T$ to be the state of the world at time t . Then, the state changes according to the transition matrix shown in Table 21.3.

$\lambda_1 \in [0, 1]$ is the transition probability of switching from model M to model T , while λ_2 is the transition probability of switching from model T to model M in two successive periods. We make the following assumptions on λ_1, λ_2 .

$$\lambda_1 + \lambda_2 < 1 \text{ and } \lambda_1 < \lambda_2. \quad (21.62)$$

The assumptions in (21.62) impose two restrictions. (1) Each of λ_1, λ_2 is individually small. So the investor believes that the probability of a model change is small. (2) The probability of switching from model M to T is lower than the probability of switching in the reverse direction; thus model M is the more persistent model.

In order to value the security, the investor must evaluate the expected returns from holding it. This requires the investor to guess which of the two (incorrect) models generates the earnings. In period t , denote by π_t^M , the belief that the data is generated by model M ; with the complementary probability, $1 - \pi_t^M$, the data is believed to be generated by model T .

The investor observes the earnings shock ε_{t+1} and wishes to calculate the posterior probability, $\pi_{t+1}^M = P(s_{t+1} = M \mid \varepsilon_{t+1}, \varepsilon_t, \pi_t^M)$, that model M generated the earnings in period $t + 1$. Using Bayes' rule,

$$\pi_{t+1}^M = \frac{P(\varepsilon_{t+1} \mid \varepsilon_t, s_{t+1} = M, \pi_t^M) P(s_{t+1} = M)}{P(\varepsilon_{t+1} \mid \varepsilon_t, s_{t+1} = M, \pi_t^M) P(s_{t+1} = M) + P(\varepsilon_{t+1} \mid \varepsilon_t, s_{t+1} = T, \pi_t^M) P(s_{t+1} = T)}, \quad (21.63)$$

where

$$P(s_{t+1} = M) = (1 - \lambda_1) \pi_t^M + \lambda_2 (1 - \pi_t^M) \quad (21.64)$$

$$P(s_{t+1} = T) = \lambda_1 \pi_t^M + (1 - \lambda_2) (1 - \pi_t^M) \quad (21.65)$$

Equation (21.64) is derived as follows. In period t , the investor assigns respective probabilities π_t^M and $1 - \pi_t^M$ that the data is generated by models M and T . If the period t model is M , then there is a probability $1 - \lambda_1$ that one transits to model M in the next period (first term in (21.64)). If the period t model is T , then there is a probability λ_2 that one transits to model M in the next

period (second term in (21.64)). Adding the two terms gives the total probability of transiting to model M in period $t + 1$. (21.65) can be explained in a similar manner.

For any values of $\varepsilon_t = \pm x$, $\varepsilon_{t+1} = \pm x$, we can calculate the RHS of (21.63). Suppose, for instance, that $\varepsilon_t = \varepsilon_{t+1} = x$, i.e., we have two successive x outcomes. Using Tables 21.1 and 21.2, we get that

$$P(\varepsilon_{t+1} = x \mid \varepsilon_t = x, s_{t+1} = M, \pi_t^M) = p_L \text{ and } P(\varepsilon_{t+1} = x \mid \varepsilon_t = x, s_{t+1} = T, \pi_t^M) = p_H \quad (21.66)$$

Substituting (21.64), (21.65), (21.66) in (21.63) we get

$$\pi_{t+1}^M = \frac{p_L [(1 - \lambda_1) \pi_t^M + \lambda_2 (1 - \pi_t^M)]}{p_L [(1 - \lambda_1) \pi_t^M + \lambda_2 (1 - \pi_t^M)] + p_H [\lambda_1 \pi_t^M + (1 - \lambda_2) (1 - \pi_t^M)]}. \quad (21.67)$$

Now suppose that we get two different outcomes in two consecutive periods, say, $\varepsilon_t = -x$, $\varepsilon_{t+1} = x$. Then proceeding as above, the analog of (21.67) is

$$\pi_{t+1}^M = \frac{(1 - p_L) [(1 - \lambda_1) \pi_t^M + \lambda_2 (1 - \pi_t^M)]}{(1 - p_L) [(1 - \lambda_1) \pi_t^M + \lambda_2 (1 - \pi_t^M)] + (1 - p_H) [\lambda_1 \pi_t^M + (1 - \lambda_2) (1 - \pi_t^M)]}. \quad (21.68)$$

Equations (21.67) and (21.68) enable us to derive the following important result.

Proposition 21.9 *Following two consecutive signals of the same sign, i.e., $\varepsilon_t = \varepsilon_{t+1} = x$ or $\varepsilon_t = \varepsilon_{t+1} = -x$ we get that*

$$\pi_{t+1}^M < \pi_t^M. \quad (21.69)$$

However, following two consecutive signals of the opposite sign, i.e., $\varepsilon_t = -x$, $\varepsilon_{t+1} = x$ or $\varepsilon_t = x$, $\varepsilon_{t+1} = -x$, we get that

$$\pi_{t+1}^M > \pi_t^M. \quad (21.70)$$

Proof: We sketch a brief heuristic proof. Using (21.67), when we get two successive x outcomes, define

$$\underline{\Delta}(\pi) = (\pi_t^M - \pi_{t+1}^M) \big|_{\pi_t^M = \pi} = \pi - \frac{p_L [(1 - \lambda_1) \pi + \lambda_2 (1 - \pi)]}{p_L [(1 - \lambda_1) \pi + \lambda_2 (1 - \pi)] + p_H [\lambda_1 \pi + (1 - \lambda_2) (1 - \pi)]}.$$

Using (21.68), when we get two consecutive signals of the opposite sign, define

$$\overline{\Delta}(\pi) = (\pi_{t+1}^M - \pi_t^M) \big|_{\pi_t^M = \pi} = \frac{(1 - p_L) [(1 - \lambda_1) \pi + \lambda_2 (1 - \pi)]}{(1 - p_L) [(1 - \lambda_1) \pi + \lambda_2 (1 - \pi)] + (1 - p_H) [\lambda_1 \pi + (1 - \lambda_2) (1 - \pi)]} - \pi.$$

Let q_t be the probability at time t that the signal came from model M . One can now verify that the shapes of $\underline{\Delta}(\pi)$, $\overline{\Delta}(\pi)$ shown in Figure 21.6 are consistent with the expressions above.




Figure 21.6 Figure corresponding to the proof of Proposition 21.9.

It can be checked that $\underline{\Delta}(\pi), \bar{\Delta}(\pi)$ are both concave in π and when $\pi \in [q_1, q_2]$, $\underline{\Delta}(\pi) \geq 0, \bar{\Delta}(\pi) \geq 0$. We can show that starting at any value of π outside the interval $[q_1, q_2]$ the investor's beliefs will come to lie in this interval. Suppose that the prior belief $\pi < q_1$. If there are two consecutive shocks of the same sign, then $\underline{\Delta}(\pi) < 0$ in this region, thus $\pi_{t+1}^M > \pi_t^M$ and beliefs are revised upwards towards the interval $[q_1, q_2]$. If, on the other hand, there are two consecutive shocks of opposite signs, then $\bar{\Delta}(\pi) > 0$ in this region, so $\pi_{t+1}^M > \pi_t^M$. Thus, beliefs are again revised upwards towards the interval $[q_1, q_2]$. One can derive a similar result in the case $q_2 < \pi$. Thus, provided that the investor is given sufficient time to update beliefs, there is no loss in generality in concentrating exclusively in the interval $[q_1, q_2]$ because starting with any prior beliefs outside this interval, with probability one, beliefs come to lie in this interval. In the interval $[q_1, q_2]$, if consecutive shocks have the same sign, then $\underline{\Delta}(\pi) > 0 \Rightarrow \pi_t^M > \pi_{t+1}^M$, as in (21.69). If, however, shocks have opposite signs, then $\bar{\Delta}(\pi) > 0 \Rightarrow \pi_{t+1}^M > \pi_t^M$ which verifies the claim in (21.70). ■

Proposition 21.9 shows that individuals revise their beliefs about the model M downwards if they observe two successive outcomes that are identical. Conversely if there are two successive outcomes that are of the opposite sign, then they revise their beliefs about model M upwards (and about model T downwards). This is an example of the representativeness heuristic. People are too quick to alter their beliefs based on a small sample.

Consider the following simulation exercise. Let $\pi_0^M = 0.5$, $\varepsilon_0 = x$, $p_L = 1/3$, $p_H = 3/4$, $\lambda_1 = 0.1$, $\lambda_2 = 0.3$. Then over the next 20 time periods, using the results derived so far, the evolution of beliefs that the model is mean reverting, π_t^M , is given in Table 21.4 for a sequence of realizations of ε_t , $t = 1, \dots, 20$.

Consider the time periods 1 to 4. During this period, the shocks alternate in sign. To the investor, this is suggestive of model M , despite the fact that the underlying model is a random walk. Hence, the investor's beliefs in model M go up very rapidly (from 0.5 to 0.94). On the other hand, during time periods 10 to 14, there are five successive and identical outcomes, which to the investor, is suggestive of model T . Hence, the beliefs of the investor that the model is M go down very rapidly (from 0.94 to 0.36). Table 21.4 also illustrates both the results in Proposition 21.9.

The price of the security is determined as follows. There is a single representative investor. The price, P_t , of the security in period t is the present discounted value of expected future dividends, $D_t, E_t D_{t+1}, \dots$

Table 21.4 Simulations for the Barberis et al. (1998) model.

t	ε_t	π_t^M	t	ε_t	π_t^M
0	x	.50	11	x	.74
1	$-x$	.80	12	x	.56
2	x	.90	13	x	.44
3	$-x$	.93	14	x	.36
4	x	.94	15	$-x$	.74
5	x	.74	16	x	.89
6	$-x$	.89	17	x	.69
7	$-x$	.69	18	$-x$	.87
8	x	.87	19	x	.92
9	$-x$	.92	20	x	.72
10	x	.94			

$$P_t = \left[\frac{D_t}{1+\theta} + \frac{E_t D_{t+1}}{(1+\theta)^2} + \dots \right]. \quad (21.71)$$

Using the true model, (21.60), i.e., $D_t = D_{t-1} + \varepsilon_t$; $\varepsilon_t = \pm x$, we know that

$$\begin{aligned} D_{t+1} &= D_t + \varepsilon_{t+1}, \text{ so } E_t D_{t+1} = D_t, \\ D_{t+2} &= D_{t+1} + \varepsilon_{t+2} = D_t + \varepsilon_{t+1} + \varepsilon_{t+2}, \text{ so } E_t D_{t+2} = D_t. \end{aligned}$$

In general, $E_t D_{t+j} = D_t, j = 1, 2, 3, \dots$. Thus (21.71) implies

$$\begin{aligned} P_t &= \frac{D_t}{1+\theta} \left[1 + \frac{1}{(1+\theta)} + \frac{1}{(1+\theta)^2} + \dots \right]. \\ \Rightarrow P_t &= \frac{D_t}{1+\theta} \left(\frac{1}{1 - \frac{1}{(1+\theta)}} \right) = \frac{D_t}{\theta}. \end{aligned} \quad (21.72)$$

Thus, if the investor knew the correct model of the economy, then the price of the security is given in (21.72). However, for an investor who mistakenly thinks that the dividends are generated by either a mean reverting or a trending model, Barberis et al. (1998, Proposition 1, p. 324) show that the price is given instead by

$$P_t = \frac{D_t}{\theta} + \varepsilon_t (c_1 - c_2 \pi_t^M), \quad (21.73)$$

where c_1, c_2 are constants that depend on the parameters of the model. Comparing (21.72), (21.73), misperception of the true model adds a bias-term, $\varepsilon_t (c_1 - c_2 \pi_t^M)$, to the correct price. For suitable values of c_1, c_2 one can generate underreaction in the short run and overreaction in the longer run. We briefly provide some intuition for the restrictions on the values of c_1, c_2 ; for details, see Barberis et al. (1998).

In the short run, following good news, there is, on average, insufficient adjustment; see (21.58). Thus, the price does not, on average, adjust sufficiently upwards (short-run underreaction). When

investors know the true model and there is no underreaction, $P_t = \frac{D_t}{\theta}$. However, in the presence of short-run underreaction and model misperception, P_t is given by (21.73). This short-run underreaction implies that $\varepsilon_t(c_1 - c_2\pi_{avg}^M) < 0$ (where π_{avg}^M is the average value of π) or equivalently,

$$c_1 < c_2\pi_{avg}^M. \quad (21.74)$$

Analogously, for P_t to exhibit long-run overreaction following a sequence of good news we know that the expected future price is lower. We know (see Table 21.4) that following a sequence of good news, the belief in model M, π_t^M , is quite low; let this low value be π_{low}^M . Then, for overreaction, we need P_t to exceed its fundamental value $\frac{D_t}{\theta}$, or equivalently

$$c_1 - c_2\pi_{low}^M > 0. \quad (21.75)$$

From (21.74), (21.75) we get the necessary restrictions on c_1, c_2 ,

$$c_2\pi_{low}^M < c_1 < c_2\pi_{avg}^M. \quad (21.76)$$

Condition (21.76) is sufficient to generate both underreaction and overreaction.

There are, however, other explanations of overreaction. Hong et al. (2007) formulate a model in which demands and prices of securities are determined by the interaction between financial advisors and investors. Consider an event that captures technical innovation. Advisors are of two types, tech-savvies (who understand new technology) and old-fogies (who do not understand the new technology, thus, they make downward biased recommendations). In order to signal their type, tech-savvies inflate their forecasts about the new technology, creating a bubble (or overreaction) in the new technology.¹⁹

21.7 Behavioral corporate finance

The subject of *corporate finance* has traditionally examined the interaction between managers of firms and investors in the firm. Implications are then drawn about the various actions taken by the firm, such as the source of finance, e.g., equity versus bonds, merger and acquisition decisions, capital restructuring, and the level of output. Classical finance assumes that both managers and investors are fully rational. Recent developments in *behavioral corporate finance* typically relax one of these two assumptions at a time. In this section we look at one representative model from each of these emerging literatures.²⁰

There are two main reasons for the rapid development and success of behavioral finance (as distinct from behavioral corporate finance). First, there is a generally agreed upon theoretical benchmark, the efficient markets hypothesis, which is refutable. Second, there is a good availability of data that makes it possible to refute the benchmark. However, both these aspects are

¹⁹ Many investors who are *sophisticated* might be informed about the types of the financial advisors. However, the results survive, provided that there is some critical minimum fraction of investors, which is *naïve*.

²⁰ For surveys of behavioral corporate finance, see Baker and Wurgler (2012) and the collection of papers in Baker and Nofsinger (2010).

missing in corporate finance, with the result that developments in this area have been relatively slower, but not insignificant.

21.7.1 *An introduction to managerial biases*

It is likely that the top levels of management in firms (henceforth, CEOs) exhibit several behavioral biases. Further, CEOs might be entrenched in the firm and hard to get rid of, e.g., the market cannot short sell a biased CEO. At the same time, there is a lack of good panel datasets on events that take place within a firm that could enable one to examine the several competing explanations in corporate finance. But why are CEOs likely to make mistakes or have biases that might be difficult to correct? There are several reasons that apply, in particular, to top level management, rather than the lower echelons of the hierarchy (Lazear, 2004).

First, many managerial decisions are of an infrequent nature that do not allow sufficient learning. The decisions are often so complex that the optimal solution is unclear, e.g., decisions to engage in specialized investments, large investments, merger decisions, acquisition decisions, and capital restructuring. Second, there might be no widely acceptable method of evaluating the efficiency of several corporate decisions, e.g., the appropriate evaluation of the long-run effects of mergers. Third, CEOs might be able to protect themselves against market competition. Unlike the standard assumption in the economics of contracts, the board of directors might not always act in a manner that maximizes the shareholder value. Top executives might be able to influence the board, determine their own pay, and be able to extract rents (Bebchuk et al., 2002).

A related question is why do biased CEOs make it to the top? Do not the firms' internal selection processes mitigate against such appointments? Not necessarily. Overconfident managers, for instance, are likely to take more risks and are more likely to leapfrog other managers in a tournament setting when relative performances are compared (Goel and Thakor, 2008). Furthermore, it is conceivable that more confident managers might have other traits that are correlated with the perception of being a good manager. These might include, for instance, the traits of being inspiring, charming, and possessing a positive outlook.

Another reason why the performance of CEOs might be hard to evaluate is because of *attribution-bias*. Outcomes often depend both on the CEO's skills as well as on purely random luck/external factors. However, boards of directors or shareholders might underestimate the effect of luck/external factors. The problem is exacerbated by a *self-serving bias* on the part of the CEOs. When companies do badly, annual reports are likely to pin the blame on external factors that are beyond the control of the CEOs. Conversely, the good performance of companies is attributed to internal factors (Bettman and Weitz, 1983). Bertrand and Mullainathan (2001) find that as oil prices increase, the compensation of executives in oil companies increases. However, oil prices are determined by events that the executives of oil companies have little control over, such as the world demand as well as the strategic policies of the OPEC countries.

Thus, there are good grounds to suspect that top managers can be biased, but is this borne out by the evidence? A growing literature answers this question in the affirmative. In most of this literature, CEO bias is studied using the trait of overconfidence. Overconfidence can mean several things. It could mean, for instance, assigning greater weight to the probability of a good outcome and less weight to poorer outcomes; this can be captured by non-linear probability weighting, as in Part 1 of the book. Or, it could mean an overestimate of the marginal product of one's efforts. One cause of overconfidence is the so-called *better than average effect* (Larrick et al., 2007; Benoit et al., 2014). A very large percentage of individuals rate their skills, e.g., driving skills, competence,

and IQ, better than the average or median values (Larwood and Whittaker, 1977; Svenson, 1981; Alicke, 1985). The persistence of overconfidence over time requires that individuals update information, as it accrues, in a biased and self-serving manner (Langer and Roth, 1975).

Weinstein (1980) suggests two domains in which managers are likely to be overconfident. First, in domains where managers have, or think that they have, control (*illusion of control*). Second, managers are more overconfident in domains where they have a *personal commitment*. Hong and Kubik (2003) examine the relation between accuracy and optimism in the earnings forecasts of security analysts with their job separations. Analysts who are more optimistic, rather than more accurate, are more likely to experience a favorable job separation, e.g., moving up to a high-status brokerage house. Brokerage houses apparently reward optimistic analysts who promote stocks more than relatively more accurate analysts. Other evidence of managerial overconfidence can be found in March and Shapira (1987) and Ben-David et al. (2013).

Entrepreneurs are typically overconfident of the success probability of their startups; an early set of experiments that supported this conclusion were conducted by Camerer and Lovo (1999). Only 5% of startup entrepreneurs believe that their odds are any worse than comparable enterprises and a third believe that their success is assured (Cooper et al., 1988). Based on French data, 56% expect “development” and only 6% of startup entrepreneurs expect “difficulty”; three years on, the respective figures are 38% and 17% (Landier and Thesmar, 2009). Empirically, only half of all startups survive beyond three years (Scarpetta et al., 2002) and the high failure rate among startups is widely reported in the popular press.

Unlike the correlation of some other behavioral phenomena with age, e.g., risk aversion, there is no clear relation between overconfidence and age. A positive relation is found in Hansson et al. (2008) while a negative relation is found in Tournon and Hertzog (2004). For an overview of the literature, see Glaser et al. (2004). Menkhoff et al. (2013) find that different measures of overconfidence need not be highly correlated. They also find that, in their sample, investment advisors are most overconfident while institutional investors are least overconfident.

What are the economic implications arising from biased managerial judgment? Roll (1986) linked managerial overconfidence, or hubris, to the tendency of managers to assign too high a value on firms that they were keen to acquire. So, in this view, managers are subject to a sort of winner’s curse. The successful bidders in mergers bid too high. R&D is an activity with highly uncertain returns. It is not surprising that managerial overconfidence is associated with higher levels of R&D activity and higher patent counts (Hirshleifer et al., 2012). Yet another manifestation of CEO overconfidence is the practice of holding on to the stock of their company for too long, hence, forgoing a profitable opportunity to sell company stock (Meulbroeck, 2001; Hall and Murphy, 2000, 2002). Indeed, this is the proxy used to determine managerial overconfidence in several empirical studies starting with Malmendier and Tate (2005).

Malmendier and Tate (2005) show that overconfident CEOs are more than one and a half times likely to engage in merger activities with other firms relative to the fully rational benchmark. Malmendier and Tate (2008) confirm the relation between overconfident CEOs and mergers. They also find that the market reaction to mergers is smaller when investors perceive CEOs to be more overconfident, suggesting a degree of sophistication in the market. Malmendier et al. (2011) show that more overconfident managers are also more likely to believe that their stocks are undervalued by the market and external finance is overpriced.

One implication of managerial overconfidence is that they believe that the prospects of their company are much better relative to the market perception. Hence, overconfident managers will believe that their stocks are undervalued by the market. For this reason, they will be reluctant to sell equity for investment or for merger activities. If the firm does not have access to adequate

levels of internal finance and/or is incapable of absorbing greater debt, then managerial overconfidence will result in lower levels of investment and merger activity. On the other hand, if sufficient internal finance is available, then overconfidence might lead to excessive levels of investment/mergers (Heaton, 2002; Gervais et al., 2011). In this case, high-powered managerial incentives will make a bad situation worse by making the manager too risk loving.

How can a firm rein in an overconfident manager? High levels of previous debt can reduce the overconfident manager's desire to overinvest. An impartial outside observer, whose consent might be needed by the manager before the board of directors will allow actual investment can also play a positive role. Set against these negative effects, one possible benefit of managerial overconfidence arises in team production (Gervais and Goldstein, 2007). In this setting, a team may contain an overconfident member who overestimates his own marginal product of effort and works harder. Due to complementarities in effort, this reduces the extent of the free-rider problem and increases total output.

Aktas et al. (2015) study the link between narcissism and mergers. Narcissism, a human trait, possibly linked to overconfidence, was measured by the ratio of the number of times that CEOs used first person singular pronouns to the total number of first person pronouns in transcribed speeches. Narcissists were more likely to play the role of acquirers in mergers, negotiate faster, and more likely to have initiated and completed mergers. In their roles as managers of target firms, they are more likely to have successfully secured higher bid premia.

21.7.2 *Rational managers, irrational investors, and merger/acquisition activity*

Three interesting waves of mergers have taken place from the 1960s to the 1990s. The merger movement in the 1960s is sometimes referred to as the *conglomerate wave*. This took the form of mergers between firms belonging to different industries. The *hostile takeover wave* of the 1980s was different in nature from that in the 1960s in at least two respects. The acquirers were mainly financiers and they paid in cash rather than in stock. The merger wave of the 1990s, on the other hand, had similarities with that of the 1960s wave. Mergers took place largely between firms in the same industry, the payment was in stock, and stock market valuations, unlike the 1980s wave, were very high.

Traditional corporate finance largely views mergers as a means to improving efficiency in the presence of industry-specific shocks. By improving efficiency, the prediction is that the profits of the merged firms should increase. However the existing empirical evidence is not able to give a clear answer to this prediction. For instance, what is the appropriate length of time over which profits should increase, and what is the relevant measure of profits? Furthermore, traditional corporate finance cannot explain specific features of the different waves of mergers. For instance, why does payment for a target in a merger sometimes take the form of stocks, but cash at other times? There are also other grounds on which the traditional theory is incompatible with the evidence from the stock market, e.g., the stock market might not react to news of the merger in the predicted manner (Shleifer and Vishny, 2003). Shleifer and Vishny (2003) use behavioral insights to enrich the standard model to take account of such considerations. We now outline this simple, stylized, and reduced-form model of mergers and acquisitions.

Consider two firms, labeled 1 and 2, who own capital stocks K_1 and K_2 , respectively. The stock market assigns values v_1, v_2 per unit, respectively, to these capital stocks, such that $v_1 < v_2$. In the spirit of the models considered in this chapter, it is assumed that the stock market is not efficient. Rather, it is subject to investor sentiment, but its precise source is not modeled here.

Assume that there is some *synergy* between the two firms if they merge; synergy could arise, for instance, from the sharing of information, technology, client lists, and market niches. The combined equity of the two firms, $K_1 + K_2$, is valued at v per unit, so that the total market value of the merged firm is given by

$$V = v \times (K_1 + K_2). \quad (21.77)$$

The following five cases are possible.

$$(a) v = v_1, (b) v < v_1, (c) v = v_2, (d) v_2 < v, (e) v_1 < v < v_2. \quad (21.78)$$

The unit value of the merged firm, v , is either the same or lower than the value of the less valuable of the two merged firms (cases (a) and (b)). Or the market could be very optimistic and value the combined firm the same as the more valuable firm (case (c)), or the market might be extremely euphoric (case (d)).

The most reasonable case seems to be (e), which is the only one that we pursue below.

Definition 21.14 *The no-synergy point is defined to be a value v^* such that*

$$v^* \times (K_1 + K_2) - v_1 K_1 - v_2 K_2 = 0. \quad (21.79)$$

From (21.79), if $v > v^*$, then the synergy level is high enough to ensure that the merged firm is more valuable than two independent firms, and if $v < v^*$, then a merger erodes the value of the two independent firms.

A simple distinction is made between the short run and the long run. In the long run, all assets, whether they belong to any of the independent firms or to the merged firm, are valued by the market at $\bar{v}$ per unit. Thus, the long-run values of firm 1, firm 2, and the merged firm, are, respectively, equal to $\bar{v}K_1$, $\bar{v}K_2$, and $\bar{v}(K_1 + K_2)$. Thus, there are no long-run gains or losses from a merger; any gains/losses only accrue in the short run.

For simplicity, we define a merger between the two firms in one of two ways; as in many modern analyses, the boundaries between mergers and acquisitions are quite blurred. Either firm 1 acquires firm 2, or vice versa. We consider only the case when firm 2, the more valuable firm, acquires firm 1. In this case, we refer to firm 2 as the *acquirer* and firm 1 as the *target*. The market is assumed not to draw any inferences from the method used by one firm to acquire another. Thus, in particular, the beliefs of the market, which are summarized by v_1 , v_2 , v , are not affected by alternative methods of payment, e.g., cash or stock, to acquire the target; this contrasts with the classical case that employs a signaling explanation that is based on the mode of financing (Myers and Majluf, 1984).

Suppose that the acquirer pays a price p per unit of equity to acquire the target, so the total money that is paid to acquire the target is pK_1 . The bargaining process is unmodeled, but some of the interesting possibilities, depending potentially on the relative bargaining strengths of the two firms, are the following.

$$(a) p = v_1, (b) p = v, (c) p < v, (d) v_1 < p < v. \quad (21.80)$$

In case (a), the target receives no *takeover premium*. In case (b), the target has enough bargaining power to receive a price that equals the short-run valuation of the merged firm. Case (c) reflects a somewhat lower bargaining power for the target firm. Case (d) recognizes that both firms have

some bargaining power. In this case, the target firm captures some, but not all, of the benefits of the market's perceived synergy between the two firms. The following proposition is self-evident and stated without proof.

Proposition 21.10 *The short-run effect of the merger on the profit levels of various firms is as follows.*

(a) *Combined market value*

$$v \times (K_1 + K_2) - v_1 K_1 - v_2 K_2. \quad (21.81)$$

(b) *Profits of the target firm*

$$(p - v_1)K_1. \quad (21.82)$$

(c) *Profits of the acquiring firm*

$$v \times (K_1 + K_2) - pK_1 - v_2 K_2 \equiv (v - p)K_1 + (v - v_2)K_2. \quad (21.83)$$

Notice that in (21.81)–(21.83) the appropriate opportunity cost in each case has been taken into account to derive the net benefits. Suppose that $v_1 < p < v$ (case (d) in (21.80)). Then, from (21.82), (21.83) the target certainly benefits from the merger. The acquirer can also potentially benefit from the merger; there is a gain from an increase in the value of the target ($v - p > 0$), but a loss on account of the merged firm being less valuable than the acquirer ($v - v_2 < 0$). If the target has no bargaining power at all, $p = v_1$, then, from (21.82), the target firm derives no net benefit from the merger.

Continue to assume that $v_1 < p < v$ (and recall that we are assuming the plausible case (e) in (21.78), $v_1 < v < v_2$).

The long-run effects on the values of various firms is immediate from the assumptions made above; this is formally stated in Proposition 21.11.

Proposition 21.11 *Suppose that the target is acquired by means of a cash payment. The long-run effect on the target shareholders from being acquired is*

$$(p - \bar{v})K_1. \quad (21.84)$$

The long-run effect on the acquirer shareholders from acquiring the target firm is

$$(\bar{v} - p)K_1. \quad (21.85)$$

From (21.84), (21.85), the long-run effect of the merger on the combined shareholders is

$$(p - \bar{v})K_1 + (\bar{v} - p)K_1 = 0. \quad (21.86)$$

From (21.86), in the long run, the shareholders of the two firms play a zero-sum game when one acquires the other. So why should the acquirer wish to acquire the target? From (21.85), if $\bar{v} > p$, then the acquirer makes a long-run profit at the cost of the target. Hence, short-run undervaluation of the target firm relative to its long-run valuation ($\bar{v} > p$) is a necessary condition for the acquisition. The managers of the target firm would benefit their shareholders in the long run if they can resist takeovers at a price $p < \bar{v}$.

Now consider a merger in which the acquiring firm buys shares (rather than paying cash) of the target firm. In this case, the fraction of the merged firm, λ , owned by the target firm is given by

$$\lambda = \frac{pK_1}{v \times (K_1 + K_2)}, \quad (21.87)$$

i.e., it is the value placed on the target firm in the acquisition, relative to the total value of the merged firm. The long-run value of this share to the shareholders of the target firm is $\bar{v}\lambda (K_1 + K_2)$, which on using (21.87) is

$$\bar{v}\lambda (K_1 + K_2) = \frac{\bar{v}pK_1}{v}. \quad (21.88)$$

The next result is evident from (21.87), (21.88).

Proposition 21.12 *In the long run, the gain to the shareholders of the target firm from being acquired for stock is*

$$\bar{v}\lambda (K_1 + K_2) - \bar{v}K_1 = \bar{v}K_1 \left(\frac{p}{v} - 1 \right). \quad (21.89)$$

Since the game between the acquirer and the target is zero-sum (see Proposition 21.11), the long-run gain to the shareholders of the acquiring firm from acquiring the target for stock is

$$\bar{v}K_1 \left(1 - \frac{p}{v} \right). \quad (21.90)$$

From (21.90), when using stock for acquisition, the shareholders of the acquiring firm gain in the long run only if $p < v$. This is more likely to occur when the acquiring firm has greater bargaining power in the takeover, which drives down the price of the acquisition, p . Conversely, shareholders in the target firm gain in the long run if $p > v$.

Suppose, that two conditions hold. (C1) Profits of the acquiring firm, given in (21.83), are negative, so $(v - p)K_1 + (v - v_2)K_2 < 0$, and (C2) $p < v$. In this case, because of (C1) and Proposition 21.10, the short-run value of the acquiring firm falls when the merger takes place. However, because of (C2) and Proposition 21.12, in the long run, the shareholders of the acquiring firm gain. In the absence of the acquisition, the shareholders would have been even more worse off in the long run. Thus, when the acquiring firm (firm 2) is relatively overvalued ($v < v_2$), and $p < v$ (target is relatively undervalued), then it is worthwhile for the shareholders of the acquiring firm to take a hit in short-term share prices in order to achieve long-run gains. These are the conditions under which theory predicts mergers for stock to take place.

How does the acquisition of firm 1, using stock, compare with the acquisition of an identical amount of capital, i.e., K_1 in the market at some price per unit of capital, P ? Suppose that the market places a value V per unit on the capital acquired by firm 2 from firm 1. Proceeding as in Proposition 21.12, and analogous to (21.90), the long-run gain to the shareholders of firm 2 is given by

$$\bar{v}K_1 \left(1 - \frac{P}{V} \right). \quad (21.91)$$

Comparing (21.90), (21.91), the acquisition of a firm by stock is preferred to acquiring an equivalent amount of capital on the market if

$$\frac{p}{v} < \frac{P}{V}.$$

One expects that the cost of new capital is lower than the cost of acquiring the same capital by buying a firm, i.e., $P < p$. However, it is also likely that $V < v$ (perceived synergies of a merger are likely to be higher than the simple acquisition of new capital). Thus, the result can go either way. In general, mergers do not always dominate investment in new capital or vice versa, hence, we expect to observe both.

Proposition (21.12) raises another interesting issue. By merging, one firm always loses out relative to the other in the long run. Since the acquirer would not acquire if it loses out, this begs the question of why the managers of the target firm would agree to a deal that harms the interests of its long-run shareholders. The reason lies in the different time horizons of the managers. As Shleifer and Vishny (2003) point out, target firms are expected to be those whose managers want to sell out, for reasons of retirement or ownership of illiquid stock options.

Similarly, the managers of the acquirer firm are expected to be ones who are locked into the firm (e.g. relatively younger managers who have a long time left for retirement) or those who have a relatively large and long-term equity stake in the firm. Target managers might be given several kinds of inducements to sell out the interests of their long-term shareholders. These include, generous severance pay, continuation of the managers in top positions in the merged set up, and acceleration of the exercise of stock options. The managers of the acquiring firm benefit by the acquisition of undervalued equity, hence, managers of both firms benefit.

From Proposition 21.11, acquisitions for cash are likely to take place when the target firm is undervalued relative to its long-run value. From Proposition (21.12) we see that acquisitions for stock are likely to take place when $p < v$. This happens when the market has a strong belief in a synergistic merger and the bargaining power of the acquiring firm is much greater than that of the target firm. Furthermore, it helps in this regard if the managers of the target firm are willing to sell-out. Why might the market hold a strong belief in a synergistic merger, particularly when so many mergers fail? It is difficult to believe that investors are entirely rational in this case.

Overall, the results of this model are consistent with a great deal of empirical evidence on mergers. For instance, the prediction that long-run bidder returns from cash acquisitions should be positive (this requires $\bar{v} > p$ in Proposition 21.11) is supported by the evidence (Loughran and Vijh, 1997; Rau and Vermaelen, 1998). The misvaluation motive for mergers, given above, has also been supported by the empirical results in Dong et al. (2006). It has also been found that overvaluation increases the probability of firms becoming acquirers (Ang and Cheng, 2006); condition C1 is more likely to hold when v_2 is high relative to v .

In the face of potential investor sentiment, rational managers in modern behavioral corporate finance theories have the following three main objectives. (1) Maximizing fundamental values. (2) *Catering*, which is an attempt to increase prices beyond fundamental values. (3) *Market timing*, which is a response to temporary mispricing of securities. Thus, managers sell their securities when they perceive that they are overpriced and repurchase them when they perceive that they are underpriced. Indeed, the Shleifer and Vishny (2003) models of mergers may be considered to be an example of a market timing model; acquire the target when it is undervalued and one's own value is overvalued. An important factor in empirically identifying these objectives is to identify not only fundamental value, but also mispriced stocks. This is often a difficult task (Baker and Wurgler, 2012).

Several empirical studies test the catering and timing objective by correlating the investment decisions of managers with proxies for mispricing of stocks. There is a danger that these proxies

might be correlated with fundamental values, but the newer empirical literature has attempted to deal with this problem. For instance, Polk and Sapienza (2009) conjecture that investment made by managers who have a short horizon; or whose assets are more difficult to value; or who are not cash or debt constrained, will be more sensitive to perceived market mispricing. This conjecture is confirmed by the evidence. Similar proxy-based evidence on the relation between investment and mispricing is reported in Chirinko and Schaller (2001) and Gilchrist et al. (2005). Baker et al. (2003) extend a model based on Stein (1996) to predict that equity-dependent firms will exhibit a stronger correlation between investment and mispricing; their empirical results confirm the prediction.

Survey evidence also supports the market timing hypothesis. In surveys, managers report that they typically issue new securities when they perceive an overvaluation of their securities and repurchase when they perceive an undervaluation (Graham and Harvey, 2001). These results are confirmed by Brav et al. (2005) who also point to the growing importance of repurchases relative to dividends. The basic motivation for issuing dividends to take advantage of market timing issues has been found in diverse markets; for instance, for IPO issues, this motivation has been found in Italian firms (Pagano et al., 1998) and in biotech firms (Loughran et al., 1994), for UK data (Marsh, 1982) and US data (Jung et al., 1996; Erel et al., 2012). The timing of opportunistic sales of equities to take advantage of mispricing is particularly targeted at new shareholders through commitment offers that exclude ongoing shareholders (Burch et al., 2004). Evidence suggests that managers are generally successful in earning higher subsequent returns based on their buying and selling decisions that are motivated by market timing issues (Ikenberry et al., 1995, 2000).

Many other examples of the catering objective can be given. These include catering to the market sentiment that favoured diversification mergers in the 1960s conglomerate wave. In the 1960s, news about a diversification merger commanded a huge premium, while by the late 1960s, this sentiment had petered out. In the mid to late 1970s, news about diversification mergers had little or no effect on stock prices, and in the 1980s, it had a negative impact (Matsusaka, 1993; Klein, 2001; Morck et al., 1990). When diversification mergers fell out of fashion, a reverse wave of divestiture and focus followed, with corresponding premia for such activities (Porter, 1987; Kaplan and Weisbach, 1992).

Name changes of firms that reflect the current market sentiment is another form of catering. Examples include the trionics boom of 1959–1962 during which firms made attempts to insert the word “electronics” in their names (Malkiel, 1990). A similar phenomenon was found during the dotcom bubble during which many firms altered their names with an Internet connotation, which had a positive effect on share prices (Cooper et al., 2001). When the dotcom bubble crashed, many firms then moved again to change names that no longer reflected an Internet connotation; again there was a positive effect on share prices of this name change announcement (Cooper et al., 2005).

21.7.3 *Managerial overconfidence with rational investors*

We now use the model of Gervais et al. (2011) to illustrate the second strand of the behavioral corporate finance literature: irrational managers and rational investors. Consider a two-period model; time is denoted by $t = 1, 2$. Suppose that a firm has \$1 in cash at $t = 1$. It can keep the cash, or invest it in a risky project that matures at $t = 2$ with a random payoff π . The random variable π takes one of two values: $\pi = z > 1$ (with probability p), or $\pi = 0$ (with probability $1 - p$). The discount rate and the interest rate are both normalized to zero. The expected net present value of the risky project is negative,

$$pz - 1 < 0. \quad (21.92)$$

Hence, the firm decides not to undertake the project.

Alternatively, the firm can hire a manager to make the investment decision on its behalf. The value added by the manager is that he can observe some signal s of the profitability of the project. With probability $q \in (0, 0.5]$, the signal equals the payoff of the risky project, π ; this captures the skill of the manager. With the complementary probability $1 - q$, the signal s equals a random variable $\tilde{\pi}$, which is independent of π , but has exactly the same probability distribution as π . Thus, the signal s is only useful with probability q .

The manager is overconfident of his skill and overestimates the usefulness of the signal by mentally assigning it a probability $q + h$, where $h \in [0, 0.5]$; the symbol h for overconfidence denotes the “hubris” of the manager (Roll, 1986). Hence, with complementary probability $1 - (q + h)$, the signal is perceived by the overconfident manager to take a value equal to the random variable $\tilde{\pi}$. This is schematically shown in Figure 21.7. Collapsing the decision tree from the end nodes, using the multiplicative law of probabilities, the manager’s signal takes on two possible values, z (good signal) and 0 (bad signal).

Definition 21.15 A measure of the manager’s “skill” is the magnitude of q . A measure of the manager’s “overconfidence” is the magnitude of h .

We now use Bayes’ rule to calculate the probabilities of different realizations of π , conditional on different values of the manager’s signal that takes values either z or 0 .

$$\Pr(\pi = z | s = z) = \frac{\Pr(s = z | \pi = z) \Pr(\pi = z)}{\Pr(s = z | \pi = z) \Pr(\pi = z) + \Pr(s = z | \pi = 0) \Pr(\pi = 0)}. \quad (21.93)$$

From Figure 21.7, the unconditional probabilities are

$$\Pr(\pi = z) = p, \Pr(\pi = 0) = 1 - p.^{21} \quad (21.94)$$

Suppose that $\pi = z$. What is the conditional probability of $s = z$? Recalling that $\pi, \tilde{\pi}$ are independent, Figure 21.7 implies the following calculation.




Figure 21.7 A decision tree depiction of the model of Gervais et al. (2011).

Source: Gervais, Simon, J. B. Heaton, and Terrance Odean (2011), “Overconfidence, compensation contracts, and capitals budgeting,” *Journal of Finance* 66: 1735–77. © American Finance Association. With permission from John Wiley & Sons.

²¹ Note that the unconditional probability $\pi = z, \pi = \tilde{\pi}$ are independent.

$$\Pr(s = z \mid \pi = z) = (q + h) + (1 - q - h)p. \quad (21.95)$$

Similarly, when $\pi = 0$, Figure 21.7 implies the following calculation.

$$\Pr(s = z \mid \pi = 0) = (1 - q - h)p. \quad (21.96)$$

Substitute (21.94) to (21.96) in (21.93) to get

$$\Pr(\pi = z \mid s = z) = p + (q + h)(1 - p) \equiv \bar{c}. \quad (21.97)$$

Proceeding analogously, we get

$$\Pr(\pi = z \mid s = 0) = (1 - q - h)p \equiv \underline{c}. \quad (21.98)$$

For a rational, i.e., unbiased manager, hubris, $h = 0$, and the probabilities stated in (21.97), (21.98) are respectively given by

$$\bar{c}|_{h=0} = p + q(1 - p); \quad \underline{c}|_{h=0} = (1 - q)p. \quad (21.99)$$

From (21.97), (21.98) we get the following result immediately.

Lemma 21.1 $\bar{c}$ is increasing in h, p, q . $\underline{c}$ is decreasing in q, h and increasing in p .

In Lemma 21.1, the most important comparative static result for our purposes is with respect to overconfidence or hubris, h . An increase in hubris increases the probability that the manager assigns to a high payoff, z , on receiving the good signal $s = z$. Conversely, he reduces the probability of a high payoff far too much, on receiving the bad signal, $s = 0$.

The only observable and verifiable event for the firm is the actual realization of the payoff. There are three possible payoffs. If the risky project is not undertaken, the payoff equals the firm's initial cash holding of \$1. If the project is undertaken, then the payoff is a random variable π that has two possible outcomes, $\$z > 1$ and $\$0$. Dropping the dollar sign, we can order these three outcomes as $y_0 = 0, y_1 = 1, y_2 = z$. To incorporate the manager's risk aversion when there are three discrete outcomes, we proceed as follows.

Notice that if the outcome is 0 (failed project) the firm loses its cash balance of \$1 (cost of investment) so, it has nothing to pay the manager with. Hence, in this case, the wage received by the manager must be zero; we denote this by $w_0 = 0$. Denote the wage of the manager contingent on y_1 and y_2 to be, respectively, w_1 and $w_1 + \Delta$, where $\Delta > 0$ is the incremental wage when the outcome improves from $y_1 = 1$ to $y_2 = z$. We may summarize the manager's compensation contract by $\{w_1, \Delta\}$. The manager's utility, u , is assumed to be piecewise linear such that

$$u(w_0) = u(0) = 0; \quad u(w_1) = w_1; \quad u(w_2) = u(w_1 + \Delta) = w_1 + (1 - r)\Delta; \quad 0 < r < 1. \quad (21.100)$$

From (21.100), the incremental utility of Δ is $(1 - r)\Delta$. The size of r , therefore, captures the manager's degree of risk aversion.

THE RATIONAL MANAGER BENCHMARK

Suppose that the manager is rational in the sense of exhibiting no overconfidence, so $h = 0$ and (21.99) applies. The firm uses information on the rational manager's skill (summarized by the parameter q) to decide whether to undertake the project.

Suppose that the rational manager receives the bad signal $s = 0$. From (21.98), (21.99), $\Pr(\pi = z | s = 0) = \underline{c}|_{h=0} = (1 - q)p < p$. Hence, using net present value calculations, the project is undertaken if $\underline{c}|_{h=0} z - 1 > 0$. But $\underline{c}|_{h=0} < p$, hence, using (21.92), the project is never undertaken when a rational manager receives a bad signal.

Now suppose that the rational manager receives the good signal, $s = z$. From (21.97), (21.99), $\Pr(\pi = z | s = z) = \bar{c}|_{h=0} = p + q(1 - p) > p$. Hence, the condition for the project to be undertaken is

$$\bar{c}|_{h=0} z - 1 > 0. \quad (21.101)$$

Substituting $\bar{c}|_{h=0}$ from (21.99) in (21.101) we get that (21.101) is equivalent to the condition:

$$q > 1 - \frac{z - 1}{z(1 - p)} = \bar{q}. \quad (21.102)$$

From (21.101), (21.102), the following result is a straightforward implication.

Proposition 21.13 *Suppose that the manager is rational.*

- (a) *If $q < \bar{q}$, then the project is never undertaken.*
- (b) *If $q > \bar{q}$, then the project is undertaken if and only if the signal is good, i.e., $s = z$.*

Suppose that the skill level of the manager is high enough, in the sense that $q > \bar{q}$ and the manager is rational ($h = 0$). Then the ex-ante gross value, V , of the firm before any signal is received by the manager, is

$$V = \Pr(s = z) (\bar{c}|_{h=0}) z + \Pr(s = 0) 1. \quad (21.103)$$

In (21.103), the first term $\Pr(s = z) (\bar{c}|_{h=0})$ is the joint probability, $\Pr(\pi = z, s = z)$, that the signal and the actual outcome are both equal to z . In this case the project is undertaken and the expected payoff is $\Pr(\pi = z, s = z)z$. However, when the signal is $s = 0$ the project is never undertaken and the firm retains its original cash endowment of \$1; this is the second term.

Using Figure 21.7, $\Pr(s = z) = p(q + h) + p(1 - q - h) = p$; similarly, $\Pr(s = 0) = 1 - p$. Substitute these calculations and (21.99) in (21.103) to get

$$V = pz[p + q(1 - p)] + (1 - p). \quad (21.104)$$

From Proposition 21.13, the firm acts on the manager's signal, s , only if the skill level of the manager is high enough, in the sense that $q > \bar{q}$. Otherwise, if the skill level of the manager is low ($q < \bar{q}$), the manager's signal is ignored. If $q > \bar{q}$, then, from (21.104), the value of the firm is increasing in the skill level of the manager. From (21.102), the critical skill level, $\bar{q}$, is decreasing in z and p . The intuition is that if a successful project is associated with a high reward, z , and/or the probability of its success, p , is too high, then the firm does not need a particularly skillful manager.

THE OVERCONFIDENT MANAGER'S INVESTMENT DECISION

Suppose that the overconfident manager's compensation scheme $\{w_1, \Delta\}$ is taken as given. The firm cannot observe or verify the information possessed by the manager. Hence, it must rely on the manager to make the investment decision.

Proposition 21.14 *Let the manager be overconfident ($h > 0$).*

(a) *Suppose that the manager observes $s = z$. Then the manager decides to undertake the project if*

$$\frac{\Delta}{w_1} \geq \frac{1 - \bar{c}}{\bar{c}(1 - r)} = \frac{(1 - p)(1 - q - h)}{1 - (1 - p)(1 - q - h)} \frac{1}{1 - r} \equiv \bar{\phi}. \quad (21.105)$$

(b) *Suppose that the manager observes $s = 0$. Then the manager decides to undertake the project if*

$$\frac{\Delta}{w_1} \geq \frac{1 - \underline{c}}{\underline{c}(1 - r)} = \frac{1 - p(1 - q - h)}{p(1 - q - h)} \frac{1}{1 - r} \equiv \underline{\phi}. \quad (21.106)$$

Proof: Suppose that the manager observes $s = z$. Using (21.97), he believes that with probability $\bar{c}$, the project generates a good outcome, z , and with probability $1 - \bar{c}$, it generates the bad outcome, 0. The wage of the manager in these two cases is, respectively, $w_1 + \Delta$ and 0. Using (21.100), the expected utility of the manager in this case is

$$\bar{c}[w_1 + (1 - r)\Delta] + (1 - \bar{c})0 = \bar{c}[w_1 + (1 - r)\Delta]. \quad (21.107)$$

If on observing $s = z$, the manager decides not to undertake the project, then he retains the \$1 cash, which is the firm's initial endowment. From (21.100), the manager's wage in this case is w_1 . Thus, on observing $s = z$, the manager invests in the risky project if

$$\bar{c}[w_1 + (1 - r)\Delta] \geq w_1. \quad (21.108)$$

Substituting $\bar{c}$ from (21.97) into (21.108), we get (21.105), as claimed.

Now suppose that the manager receives the signal $s = 0$. Proceeding, as above, the expected utility of the manager from undertaking the project is

$$\underline{c}[w_1 + (1 - r)\Delta] + (1 - \underline{c})0 = \underline{c}[w_1 + (1 - r)\Delta].$$

Comparing this with a wage of w_1 from not investing, the condition for investing requires $\underline{c}[w_1 + (1 - r)\Delta] \geq w_1$. Substituting $\underline{c}$ from (21.98) we get (21.106), as claimed. ■

The following observations can be drawn from Proposition 21.14.

1. From (21.105), (21.106), $\underline{\phi} > \bar{\phi}$. Thus, the barrier to undertaking the project when the manager receives the good news ($s = z$) is lower than when he receives the bad news ($s = 0$).
2. The interpretation of $\frac{\Delta}{w_1}$ is that it is a measure of the convexity of the manager's compensation contract. To see this, plot the outcomes $(0, 1, z)$ on the horizontal axis and the manager's compensation on the vertical axis. The wage corresponding to these three outcomes is 0, w_1 , $w_1 + \Delta$. Clearly an increase in Δ relative to w_1 makes the wage profile

more convex. Suppose now that the manager's risk aversion, as captured by r , increases. This makes the utility function of the manager more concave. From (21.105), (21.106), an increase in r increases both $\underline{\phi}, \bar{\phi}$, so the manager's investment policy becomes more conservative. Hence, to ensure that the project is undertaken, an even larger $\frac{\Delta}{w_1}$ (which is the convexity of the compensation scheme) is called for. In other words, the convexity of the compensation scheme compensates for concavity of the manager's utility function.

3. An increase in the manager's overconfidence, h , reduces $\bar{\phi}$. In this case, the manager's investment policy becomes more aggressive towards investment, following a good signal $s = z$. On the other hand, an increase in h , increases $\underline{\phi}$. Hence, following a bad signal, $s = 0$, the manager's investment policy becomes more conservative. A manager with higher overconfidence places greater reliance on privately observed signals; if the signal is bad then the manager assigns a relatively high posterior to the bad outcome. Thus, risk aversion and overconfidence have different implications for the manager's investment policy; risk aversion increases both $\underline{\phi}, \bar{\phi}$, but overconfidence reduces $\bar{\phi}$ and increases $\underline{\phi}$.

Definition 21.16 (*Underinvestment*): The overconfident manager overinvests (underinvests) if his investment policy is more (less) aggressive relative to the one followed by the firm if it had the same information (i.e., under the rational manager benchmark).

Using Propositions 21.13, 21.14, overinvestment, therefore, implies either of two conditions.

$$(1) q < \bar{q} \text{ and } \frac{\Delta}{w_1} > \bar{\phi} \text{ or } (2) q > \bar{q} \text{ and } \frac{\Delta}{w_1} > \underline{\phi}. \quad (21.109)$$

In (21.109), condition (1) requires that the skill level of the *rational manager* is not high enough for the firm to invest, even when the manager receives a good signal (see Proposition 21.13). However, the *overconfident manager* nevertheless goes ahead and invests because $\frac{\Delta}{w_1} > \bar{\phi}$ (see Proposition 21.14). In condition (2), the skills of the rational manager are high enough for the firm to invest, based on a good signal, but not if the signal is bad (see Proposition 21.13). However, since $\frac{\Delta}{w_1} > \underline{\phi}$, we know, from Proposition 21.14, the overconfident manager will invest if he receives a bad signal.

Using Definition 21.16 and (21.109), underinvestment is characterized by

$$q > \bar{q} \text{ and } \frac{\Delta}{w_1} < \bar{\phi}. \quad (21.110)$$

In this case, the firm would like the rational manager to invest if the signal is good. However, the overconfident manager does not invest.

By looking at the dependence of $\underline{\phi}, \bar{\phi}$ on the parameters of the model, we can draw further inferences about the comparative static results of the model.

Proposition 21.15 *Factors conducive to overinvestment are the following. (a) Small z . (b) Small r . (c) Large $\frac{\Delta}{w_1}$. When $q < \bar{q}$, an increase in the manager's overconfidence, h , creates greater overinvestment. Conversely when $q > \bar{q}$, an increase in the manager's overconfidence, h , creates either lower overinvestment or lower underinvestment.*

Proof: From (21.102), $\bar{q} = 1 - \frac{z-1}{z(1-p)}$, thus, $\bar{q}$ is decreasing in z . Hence, a smaller z increases $\bar{q}$, but does not affect $\underline{\phi}, \bar{\phi}$. From (21.109), if we start with, $\frac{\Delta}{w_1} > \bar{\phi}$, then an increase in $\bar{q}$ makes it more likely that the condition $q < \bar{q}$ will be satisfied, therefore, making overinvestment more

likely. Now consider a reduction in r . This does not affect $\bar{q}$. However, the reduction in r reduces both $\underline{\phi}$, $\bar{\phi}$, which makes it more likely that the conditions for overinvestment in (21.109) are satisfied. It is also straightforward from (21.109) that an increase in $\frac{\Delta}{w_1}$ is more likely to satisfy both conditions for overinvestment. Finally, using (21.105), (21.106), when h increases, $\bar{\phi}$ falls and $\underline{\phi}$ increases. When $q < \bar{q}$, condition (1) in (21.109) is more likely to be satisfied resulting in overinvestment. However, because an increase in h increases ϕ , if $q > \bar{q}$, condition (2) in (21.109) is less likely to hold, so there is less likely to be overinvestment. However, (21.110), the condition for underinvestment, is also less likely to be satisfied because as h increases, $\bar{\phi}$ falls. ■

From Proposition 21.15, when the manager is more risk averse (high r) and receives a smaller incentive payment (low $\frac{\Delta}{w_1}$), then he is more likely to underinvest. Thus, if the firm wishes to remove greater conservatism brought about by an increase in risk aversion, it must convexify the compensation contract even more.

When the manager is not skilled enough, $q < \bar{q}$, then overconfidence creates greater overinvestment. The intuition is that an overconfident manager is too confident of his signal and so places a lower weight on the riskiness of his investment decision. This makes him overinvest. However, when the manager is skilled enough, $q > \bar{q}$, then, both overinvestment and underinvestment fall. In this case, the manager's decisions become closer to the first best with rational managers.

One expects in real life that following a string of successful decisions, the overconfidence of the manager may increase even further (Daniel et al., 1998; Gervais and Odean, 2001). For a fixed compensation scheme $\{w_1, \Delta\}$, this is good news for the firm if the manager is skilled enough ($q > \bar{q}$). But if $q < \bar{q}$, this is bad news for the firm as it causes the overconfident manager to behave in an even more reckless manner relative to a rational manager.

Yet another application is the following. Suppose that a firm has a sequence of projects. The normal expectation is that the firm will take up the projects in a sequential manner. The projects with the highest net present value are undertaken first and so on. Suppose that the success of the initial projects increases overconfidence and that the manager's skill level is not high enough ($q < \bar{q}$). Then the manager is likely to prove detrimental to the firm for subsequent projects. This is because overconfidence induces the manager to take on too much risk, for projects where the net present value is not high enough.

21.7.4 *Behavioral corporate signaling*

In classical corporate finance, dividends are explained by using signaling models (Bhattacharya, 1979; John and Williams, 1985; Miller and Rock, 1985; Guttman et al., 2010). A feature of these and other classical signaling models in finance is that signaling is costly in the sense that it can reduce the value of the firm. However, firms who have better future prospects can take relatively costly actions because this allows them to separate from firms whose prospects are not as good. For classical signaling models used in corporate finance to address a range of issues such as capital restructuring, debt, and IPO underpricing, see Ross (1977), Myers and Majluf (1984), Harris and Raviv (1985), Grinblatt and Hwang (1989), and Allen and Faulhaber (1989).

In contrast, survey evidence shows that executives reject the basic assumptions of the costly signaling models for dividends. These assumptions include the following: dividends destroy firm value; dividends are used instead of cheaper borrowing; dividends are used to forgo profitable investments; dividends are used to signal strength. Indeed, only 4.4% of the executives agreed with the transmission mechanisms used in traditional explanations of dividends based on costly

signaling. Yet most executives agree that dividends are used to provide some sort of a signal to the markets (Brav et al., 2005). They write (p. 45): “However, this information conveyance does not appear to be consciously related to signaling in the academic sense. Managers reject the notion that they pay dividends as a costly signal to convey their firm’s true worth or to purposefully separate their firm from competitors.”

Baker and Wurgler (2011) propose a simple behavioral model of signaling that better explains some of the styled facts associated with issuing dividends. We now consider this model. Suppose that there are two time periods, $t = 1, 2$, and there are two economic actors: the manager of a firm and an investor or holder of equity in the firm. The structure of the model is as follows.

1. The firm and its decisions: The firm is synonymous with a benevolent manager who maximizes the utility of the investor. The cash earnings of the firm at time $t = 1, 2$ are denoted by y_t . The first period earning, y_1 , is private information to the manager. It is common knowledge to both players that y_2 is uniformly distributed over $[0, 2]$. The manager decides on the level of dividends in the first period, d_1 , and the liquidating dividend in the second period, d_2 . The expectation operator in period 1 with respect to the information set of the manager is denoted by E_M .
2. The investor’s preferences: The investor does not know y_1 , but uses first period dividends, d_1 , to form an inference about y_1 . The expectation operator in period 1 with respect to the information set of the investor is denoted by E_I . Thus, the investor’s conditional expectation of the first period cash flow is $E_I[y_1 | d_1]$. The investor has an exogenously given reference point, d^* , for first period dividends. Actual first period dividends, d_1 , constitute the reference point for second period dividends. Define the function

$$v(x) = \begin{cases} 0 & \text{if } x \geq 0 \\ \lambda x & \text{if } x < 0 \end{cases}, \quad (21.111)$$

where $\lambda > 0$ is the parameter of loss aversion.²² Then, the long-term utility of the investor is given by

$$u(d_1, d_2 | d^*) = d_1 + d_2 + v(d_1 - d^*) + v(d_2 - d_1). \quad (21.112)$$

Thus, the investor derives utility from the dividends received in both periods and experiences disutility if the dividends fall below the reference dividend levels in any period.

3. Preferences and constraints of the manager: The main difference from the classical model is that the manager has no control over the cash flow of the firm, y_1, y_2 . There is no agency problem either and the manager cannot engage in costly signaling in the sense of destroying the current value of the firm to signal the strength of the firm. The manager cannot finance dividends from any source (such as new equity or debt) other than current cash flows. Hence, the constraints are:

$$d_1 \in [0, y_1] \text{ and } d_2 = (y_1 - d_1) + y_2. \quad (21.113)$$

We consider a reduced-form approach. The expectation $E_I[y_1 | d_1]$ influences the demand for equity of the firm in the first period and, hence, its price. The manager maximizes a

²² Note that given the formulation for gains, we do not need to specify that $\lambda > 1$ in order to capture loss aversion.

weighted average of these price concerns and the utility of the investor in (21.112). Thus, the objective function of the manager is

$$W = E_M [\theta E_I [y_1 | d_1] + (1 - \theta)u(d_1, d_2 | d^*)], 0 < \theta < 1. \quad (21.114)$$

This completes the model setup. We now solve for the equilibrium. An increase in first period dividends, d_1 , has three effects. First, it leads to an upward revision in $E_I [y_1 | d_1]$ through equilibrium beliefs (see below), which is liked by the manager because it increases the (unmodeled) share prices. Second, an increase in d_1 up to the reference level d^* or beyond, avoids loss aversion in the first period. Third, an increase in d_1 increases the second period reference point, which makes it more likely that u may decrease due to loss aversion in the second period.

Substituting (21.111), (21.112), (21.113) in (21.114) we get

$$W = (1 - \theta)v(d_1 - d^*) + \theta E_I [y_1 | d_1] + (1 - \theta)E_M (y_1 + y_2) + (1 - \theta)E_M v(d_2 - d_1). \quad (21.115)$$

By definition,

$$\begin{aligned} E_M v(d_2 - d_1) &= \lambda \int_0^{2d_1 - y_1} (y_1 + y_2 - 2d_1) dy_2 \\ &= \frac{-\lambda}{2} (2d_1 - y_1)^2 \Big|_{d_1 > \frac{y_1}{2}} + k. \end{aligned} \quad (21.116)$$

where k is a constant of integration,²³ and

$$E_M (y_1 + y_2) = y_1 + 1 \quad (21.117)$$

Substitute (21.116), (21.117) in (21.115) and only writing terms relevant for d_1 we get

$$(1 - \theta)v(d_1 - d^*) + \theta E_I [y_1 | d_1] - (1 - \theta)\frac{\lambda}{2} (2d_1 - y_1)^2 \Big|_{d_1 > \frac{y_1}{2}}. \quad (21.118)$$

The manager chooses d_1 in order to maximize (21.118), conditional on the beliefs of the investors about d_1 . The next proposition describes a separating equilibrium and full set of beliefs in a perfect Bayesian equilibrium.

Proposition 21.16 (Baker and Wurgler, 2012) *Let y^* be the solution to the equation*

$$\frac{\theta}{2} (y^* - d^*) - (1 - \theta)\lambda \left[\left(\frac{\theta}{1 - \theta} \frac{1}{\lambda} \right)^2 - \left(d^* - \frac{y^*}{2} \right)^2 \right] = 0. \quad (21.119)$$

There exists a perfect Bayesian equilibrium (not necessarily unique) in which equilibrium first period dividends are given by

$$d_1 = \begin{cases} y_1 & \text{if } y_1 < d^* \\ d^* & \text{if } d^* < y_1 < y^*, \\ \frac{y_1}{2} + \frac{\theta}{1 - \theta} \frac{1}{\lambda} & \text{if } y^* < y_1 \end{cases} \quad (21.120)$$

²³ The notation $x \Big|_{d_1 > \frac{y_1}{2}}$ means that x is defined for $d_1 > \frac{y_1}{2}$.

and equilibrium beliefs are given by

$$E_I [y_1 | d_1] \begin{cases} d_1 & \text{if } d_1 < d^* \\ \frac{y^* + d^*}{2} & \text{if } d_1 = d^* \\ 2d_1 + \frac{2\theta}{1-\theta} \frac{1}{\lambda} & \text{if } d_1 > \frac{y^*}{2} + \frac{\theta}{1-\theta} \frac{1}{\lambda} \end{cases} . \quad (21.121)$$

From (21.120), firms whose cash flow is lower than the dividend expectations ($y_1 < d^*$) pay out their entire cash flow as dividends in order to minimize loss aversion costs to the investor. Firms who can afford to pay the reference level of dividends or more up to some level y^* , bunch together at the level d^* because of two considerations. First, to eliminate loss aversion costs in the first period. Second, to prevent the buildup of too high a reference level of dividends in the second period (second period reference dividends are d_1); this minimizes expected loss aversion costs in the second period. Finally, firms who have very high levels of incomes ($y^* < y_1$) can afford to pay even higher first period dividends. This increases $E_I [y_1 | d_1]$, which managers desire because of the effects on stock prices. The high income ($y^* < y_1$) also provides a cushion against second period dividends falling short of the second period reference point.

From (21.119), as λ increases and θ decreases, y^* increases. Therefore, from the second row of (21.120), more firms cluster at the reference level of dividends, d^* . The equilibrium described in Proposition 21.16 is also consistent with the partial adjustment model of Lintner (1956). From (21.120), on average, dividends increase less than one for one with an increase in first period income, y_1 .

21.8 Exercises for Part 7

- (Kahneman, 2011) (a) In a sample of 3,141 counties in the United States, counties in which the incidence of kidney cancer is lowest are rural, sparsely populated, and located in the traditionally Republican states in the Midwest, the South, and the West. What inferences can you draw from this statement?
(b) The Bill Gates Foundation poured \$1.7 billion into small schools based on the finding that most successful schools, on average, are small. In Pennsylvania, 6 of the top 50 schools are small (an overrepresentation in the sample by a factor of 4). This data induced the Gates Foundation to invest in creating small schools, sometimes even splitting large schools into smaller schools. Was this money necessarily well spent?
- Rabin (2002) provides the following experimentally observed conditional probabilities when there are two equally possible and mutually exclusive signals (A and B), using data from Table 7 of Rapoport and Budescu (1997).

$$P(A | B) = 58.5\%,$$

$$P(A | AB) = 46.0\%,$$

$$P(A | AAB) = 38.0\%,$$

$$P(A | AAA \dots) = 29.8\%.$$

Is this data consistent with classical probability theory? How might you explain this data? Is this data likely to have come from an underlying *inanimate process* (roulette wheel, throw of a dice, or toss of a coin) or an *animate process* (successful dunks by a player in basketball, or first serve aces by a tennis player)?

- (Clotfelter and Cook, 1989) Gamblers have a tendency to redeem winning lottery tickets for more tickets rather than for cash. What might account for this behavior?
- Kahneman (2011) describes the following findings from Lichtenstein et al. (1978). How might you explain these findings?
 - Strokes cause almost twice as many deaths as all accidents combined, but 80% of respondents judged accidental death to be more likely.
 - Tornadoes were seen as more frequent killers than asthma, although the latter caused 20 times more deaths.
 - Death by lightning strikes was judged to be more likely than death by botulism, even though the latter is 52 times more frequent as a cause of death.
 - Death by accidents was judged to be more than 300 times more likely than death by diabetes, but there are 4 diabetes caused deaths for every death in an accident.
- (Kahneman, 2011) Do you expect the conjunction fallacy to arise in the following problem? Give reasons for your answer.
Which alternative is more probable?
John works in a school.
John is a worker in a school who walks to work.
How does this problem differ from the Linda problem considered above?
- Tversky and Kahneman (1983) gave subjects a choice among one of the following options about the outcome of a match played by the then top tennis player Björn Borg. Please judge the relative probabilities of the choices.

- A. Borg will win the match.
- B. Borg will lose the first set.
- C. Borg will lose the first set but win the match.
- D. Borg will win the first set but lose the match.

72% of the subjects assigned the event B a lower probability than C. Is this outcome consistent with the axioms of probability theory? How can you explain these results?

7. (Hsee, 1998; Kahneman, 2011) Subjects were asked to make a choice between two dinner sets, H and L, shown in Table 7.E1.

Table 7.E1 A summary of the composition of two different dinner sets.

	Set H includes 40 pieces	Set L includes 24 pieces
Dinner plates	8, all in good condition	8, all in good condition
Soup/salad bowls	8, all in good condition	8, all in good condition
Dessert plates	8, all in good condition	8, all in good condition
Cups	8, 2 of them broken	
Saucers	8, 7 of them broken	

Source: Hsee (1998).

In the first condition, subjects had to place a valuation on both dinner sets when they are presented together. The mean willingness to pay (WTP) for set H was \$32.03 and for set L it was \$29.70. In the second condition, subjects had to state their WTP when presented with the two dinner sets separately. In this case, the WTP for L was higher (\$32.69) as compared to H (\$23.25). How might you explain these findings?

8. (Tversky and Kahneman, 1983) Two groups of subjects were given, separately, the problems shown in the separate columns in Table 7.E2.

Table 7.E2 Two different problems given to two groups of subjects in a between-groups design.

First group of subjects	Second group of subjects
A health survey was conducted in a sample of adult males in British Columbia, of all ages and occupations. Please give your best estimates of the following values: What percentage of the men surveyed have had one or more heart attacks? What percentage of the men surveyed are both over 55 years old and have had one or more heart attacks?	A health survey was conducted in a sample of 100 adult males in British Columbia, of all ages and occupations. Please give your best estimates of the following values: How many of the 100 men surveyed have had one or more heart attacks? How many of the 100 men surveyed are both over 55 years old and have had one or more heart attacks?

Source: Kahneman (2011).

65% of the subjects in the first group engaged in the conjunction fallacy. However, only 25% of the subjects in the second group engaged in the conjunction fallacy. What accounts for the difference?

9. (Slovic et al., 2002) Provide explanations for the following.
- Why do entertainers often change their names once they enter the entertainment industry (e.g., the real names of the actors John Denvers, Sandra Dee, and Judy Garland, were, respectively, Henry Deutschendorf, Alexandra Zuck, and Frances Gumm)?
 - Why do movies have background music? After all, can't we understand the events we are watching and the dialog we are hearing without the music?
 - Why are all the models in the mail-order catalog smiling?
 - Why do packages of food products carry all those little blurbs such as: *New*, *Natural*, or *Improved*?
 - Why do traditional cigarette advertisements show images of rugged cowboys and lush waterfalls?
 - Why is genetically modified food promoted as *enhanced* by proponents and *Franken-food* by opponents?
10. (Schwarz, Strack, and Mai, 1991) Consider the following two questions that were asked of subjects in experiments.
- Q1. How happy are you with your life in general?
- Q2. How often do you normally go out on a date?
- When the questions were asked in numerical order the answers were not correlated. However, when they were asked in reverse order, Q2 followed by Q1, then the responses were highly correlated. Which judgment heuristic would you use to explain these results, and why?
11. (Hoff and Pandey, 2006, 2014) In empirical studies conducted on children of different castes in India, children were classified into high and low castes and given a set of mazes to solve in a fixed amount of time. The results are shown in Figure 7.E1. In the *anonymous condition*, high caste children solved 7% more mazes than low caste children. However in *mixed caste groups*, where caste is revealed, the difference rises to 34%. When caste is revealed but the groups are not mixed, then again there is little difference in the number of mazes solved. How would you explain the results? How would you test for the possibility that the




Figure 7.E1 Average number of mazes solved.

Source: Hoff and Pandey (2006) with permission from the American Economic Association.

announcement of castes in mixed groups reduces the self-confidence of low caste participants, thereby reducing their performance?

12. (Englich and Mussweiler, 2001) In a hypothetically constructed case of rape, 42 experienced legal professionals were asked to determine the appropriate sentence, based on carefully prepared case material. After the legal professionals had studied the case material but before they had passed the sentence, they received a call from a journalist who asked if the sentence would be 1 year (first treatment) or 3 years (second treatment). The average sentences passed in the two cases were 25.43 months and 33.38 months. Is this consistent with neoclassical economic theory. If not, how would you explain the result? Can you derive a numerical index of the bias that might have caused the difference in the results in the two treatments?
13. (Wansink et al., 1998) In a field study in Iowa, shoppers in three supermarkets encountered a sales promotion for Campbell soup at about 10% off the regular price. On some days the shelf sign imposed a limit of 12 per person and on other days there was no limit per person. Shoppers purchased an average of 7 cans when the limit was in force and 3.5 when the limit was not in force. Similar findings were made in four studies spanning 89 supermarkets. How can you account for these findings?
14. Can you give reasons for the following two phenomena?
 - (a) Kahneman (2011) notes that we drive too fast when we just come off a highway onto city streets.
 - (b) Consider the following quote adapted from LeBoeuf and Shafir (2006): "A well-intentioned child who turns down exceptionally loud music to meet a parent's demand that it be played at a 'reasonable' volume may fail to adjust sufficiently from a high volume, and may feel that genuine attempts at compromise are being overlooked."
15. (Gigerenzer and Hoffrage, 1995) Give your answer separately to the two problems below before you read beyond the statement of Problem 2.

Problem 1: The probability of breast cancer is 1% for a woman at age forty who participates in routine screening. If a woman has breast cancer, the probability is 80% that she will get a positive mammography. If a woman does not have breast cancer, the probability is 9.6% that she will also get a positive mammography. A woman in this age group had a positive mammography in a routine screening. What is the probability that she actually has breast cancer?

Problem 2: 10 out of every 1,000 women at age forty who participate in routine screening have breast cancer. 8 out of every 10 women with breast cancer will get a positive mammography. 95 out of every 990 women without breast cancer will also get a positive mammography. Here is a new representative sample of women at age forty who got a positive mammography in routine screening. How many of these women do you expect to actually have breast cancer? ___ out of ___

Make sure that you have given your answers to Problems 1 and 2 before you read on. Clearly this test works best in a between-subjects design which is not possible when you do both problems. Do your answers to Problems 1 and 2 tally or differ? If not, can you think why not? Try doing these problems with different sets of friends to induce a between-subjects design. What results do you get in this case?

16. (Stanovich and West, 2000) Imagine that AIDS occurs in 1 in every 1,000 people. Imagine also there is a test to diagnose the disease that always gives a positive result when a person has AIDS. Finally, imagine that the test has a false positive rate of 5%. This means that the test incorrectly indicates that AIDS is present in 5% of the cases where it is actually absent.

Suppose that we choose a person randomly, administer the test, and discover that it yields a positive result (presence of AIDS). What is the probability that the chosen individual actually has AIDS, assuming that we know nothing else about the individual's personal or medical history?

What are the implications of your answer for the problem of base rate neglect?

17. Prove the assertion in (19.41), in Example 19.2.
18. (Kahneman, 2011) One often hears people saying: "Highly intelligent women tend to marry men who are less intelligent than they are." Assuming that this assertion is true, how can one explain this?
19. (Stanovich and West, 2000) Suppose there are four cards lying on a table showing two letters and two numbers (A, D, 3, 7). Subjects are told that (i) each card has a number on one side and a letter on the other, and (ii) the experimenter has the following rule in mind with respect to the four cards: "If there is an A on one side then there is a 3 on the other." Which cards will you turn over to determine whether the experimenter's rule is true or false?
20. (Banerjee and Duflo, 2011) Consider the following description of an empirical study, done on fruit vendors in Chennai, India, from WDR (2015, p. 32): "Each day, the vendors buy fruit on credit to sell during the day. They borrow about 1,000 rupees (the equivalent of \$45 in purchasing parity) each morning at the rate of almost 5% per day and pay back the funds with interest at the end of the day. By forgoing two cups of tea each day, they could save enough after 90 days to avoid having to borrow and would thus increase their incomes by 40 rupees a day, equivalent to about half a day's wages. But they do not do that." How would you explain the behavior of these fruit vendors?
21. (Gabaix and Laibson, 2006) Complete the proof of Proposition 19.5(i),(ii).
22. (Gabaix and Laibson, 2006) Consider the model of shrouded attributes in Section 19.16.2. Show that in a shrouded prices equilibrium, (i) social welfare loss is given by $(1 - \alpha)e$, and (ii) sophisticated consumers are $\bar{p} - e$ units better off than myopic consumers. In the unshrouded equilibrium, show that there is no inefficiency and all consumers (sophisticated and myopic) are equally well-off.
23. Show that the categorization in Example 19.5, $\{(1, 0), (1, 1), \{(0, 0), (0, 1)\}\}$, is optimal on the criterion that among all possible categorizations it minimizes the "sum across categories of the total variation of the mean from each category."
24. (Thaler, 2008) In a set of experiments, the instructions given to the students were: Below you will find four pairs of scenarios. In each case, two events occur in Mr. A's life and one event occurs in Mr. B's life. You are asked to judge whether Mr. A or Mr. B is happier. Would most people rather be A or B? If you think the two scenarios are emotionally equivalent, check "no difference." In all cases, the events are intended to be financially equivalent. The four problems used and the number of responses of each type follow. Explain the choice of the majority of the subjects for each of the four problems.
 1. Mr. A was given tickets to lotteries involving the World Series. He won \$50 in one lottery and \$25 in the other. Mr. B was given a ticket to a single, larger World Series lottery. He won \$75. Who was happier?
56 subjects chose A, 16 chose B, and 15 chose no difference.
 2. Mr. A received a letter from the IRS saying that he made a minor arithmetical mistake on his tax return and owed \$100. He received a similar letter the same day from his state income tax authority saying he owed \$50. There were no other repercussions from either mistake. Mr. B received a letter from the IRS saying that he made a minor arithmetical mis-

take on his tax return and owed \$150. There were no other repercussions from his mistake. Who was more upset?

66 subjects chose A, 14 chose B, and 7 chose no difference.

3. Mr. A bought his first New York State lottery ticket and won \$100. Also, in a freak accident, he damaged the rug in his apartment and had to pay the landlord \$80. Mr. B bought his first New York State lottery ticket and won \$20. Who was happier?

22 subjects chose A, 61 chose B, and 4 chose no difference.

4. Mr. A's car was damaged in a parking lot. He had to spend \$200 to repair the damage. The same day the car was damaged he won \$25 in the office football pool. Mr. B's car was damaged in a parking lot. He had to spend \$175 to repair the damage. Who was more upset?

19 subjects chose A, 63 chose B, and 5 chose no difference.

25. (Thaler, 2008) Provide a mental accounting explanation for the puzzles listed below.
 - (a) Why do so many late night television advertisements for kitchen utensils sell items that (i) have multitude of uses, and (ii) include several bonus items?
 - (b) In buying a new house, why is there such an extensive list of optional add-ons that are quite pricey?
 - (c) Why do we observe manufacturer's recommended retail price on so many items of sale even if the recommended retail price is not followed by the retailer?
26. (Prelec and Loewenstein, 1998) Last year, just after his 25th college reunion, Jones bought a slightly used red two-seater car, financing the purchase with a standard car loan. Although the car performed well enough, he found that suitable driving occasions were less frequent than he had anticipated, and that the thrill of recreating the college years dissipated quickly. The monthly payments became regular reminders of what was in truth a rather expensive indulgence: It was hard to justify each bill with the pleasures of an occasional weekend drive. After a bout of worrying and procrastination, he decided to pay off the entire loan. Writing the check was painful but produced relief, and restored some of the pleasures of driving. Now that he owned the car, he was no longer concerned with driving a certain amount each month, and he let suitable occasions arise spontaneously. He even began to enjoy how the thing looked in the driveway Show how this example can be rationalized in terms of the theoretical model of the red and black of mental accounting in Section 20.5.
27. (Prelec and Loewenstein, 1998) In Section 20.5, we considered the choice between prepaying and postpaying for a vacation. Consider now the case of consumer durables, which is dealt with in Example 20.20. Construct the relevant sequence of hedonic costs and benefits, analogously to Figure 20.1. [Hint: In this case, the benefits (unlike the vacation example) accrue over a much longer period. Hence, at any instant in time, the residual benefits, $\sum_{t>k} u_t$, are high enough, thus, hedonic benefits remain higher relative to the vacation example and never become negative.] Which of the two methods of payment, prepayment or postponement of payment, is preferred?
28. Soman (2001) posed the following question to subjects in an experiment. "Imagine that you recently saw an advertisement on the bulletin board. A literature professor was looking for a research assistant to work for about 15 hours. The payment was in terms of a front row seat to a professional theater performance. On the same bulletin board, a music professor was also looking for a research assistant to work for about 5 hours and this assistant would be paid with a front row ticket to a rock concert by a band that you like. You had recently seen posters for both the theater performance and the rock concert. You think you will like to see both these events, although you expect to like the rock concert more. You work for both professors (15 hours for literature and 5 hours for music) and get paid

with the two tickets to the theater and the rock concert respectively. As you are putting the tickets in your wallet, you notice that both events are scheduled for the same evening and both are at good locations on campus. The tickets are non-transferable, nor can they be exchanged. You can use only one of the tickets, and not the other. Which ticket will you use?" 95.2% of the subjects choose the rock concert.

In a second version of the experiment, subjects were told that they had purchased the tickets. The respective prices are HK\$450 and HK\$150. Now 61.7% of the subjects choose the theater performance. What does this experiment tell us about the monotonicity of sunk-time costs and the difference between sunk-time costs and sunk-dollar costs?

29. Shafir and Thaler (2006) pose the following question to 203 Princeton University students. Try to imagine the following scenario. It is the spring of your sophomore year and you plan to be on campus for your last two years. A student organization has decided to arrange a new concert series featuring just the sort of music you enjoy most. The organizers have taken the unusual step of selling two year subscriptions to this series. The subscription guarantees you great seats to eight concerts, four concerts next year and four concerts the following year. The prices on advance booking are quite attractive: Single concert tickets will cost \$25, but your series only costs \$120 (for an implied per concert price of \$15). You have talked to the organizers and you like the performers they intend to bring to campus, so you decide to buy the subscription.

At the time you make the purchase, how well does each of the following statements capture your feelings? I feel that I (a) just spent \$120, (b) made a \$120 investment, and (c) just saved \$80. Subjects were asked to rank each option from 1 (strongly agree) to 5 (strongly disagree). The mean scores were 3.63, 2.20, 2.65.

This question was followed up by another one. It is now the spring of your senior year and there is one remaining concert in the series to which you subscribed two years ago. In fact, the last concert is a group that has become very popular, and tickets are currently selling for \$50 around campus. You are about to go to the concert.

Which of the following best captures your feeling of the cost to you of attending this concert. (a) \$50 cost (b) \$25 cost (c) \$15 cost (d) \$0 cost (e) \$35 saving. The mean scores were 3.96, 4.37, 2.37, 2.48, 2.40.

The final question posed was the following. It is now the spring of your senior year and there is one remaining concert in the series to which you subscribed two years ago. In fact, the last concert is by a group that has become very popular, and tickets are currently selling for \$50 around campus. Unfortunately, you have lost your ticket, and there is no way of obtaining a replacement ticket from the organizers.

Which of the following best captures your feeling of the cost to you of losing this ticket. (a) \$50 cost (b) \$25 cost (c) \$15 cost (d) \$0 cost. The mean scores were 2.43, 4.06, 3.60, 4.33.

(a) What does this experiment indicate about the investment versus cost option at the time of buying the ticket?

(b) What does this experiment indicate about the historical cost, no cost, and the profit felt at the time when only one concert remains and the ticket is not lost?

(c) How is the answer to (b) changed when one suddenly loses the ticket?

In conjunction, how do (a)–(c) support the title of the Shafir and Thaler (2006) paper: *Invest now, drink later, spend never?*

30. (Shafir and Thaler, 2006) Consider the following problem that illustrates the framing effects arising from *local* versus *global* representations. Imagine that you have a membership to a gym in a nearby town that you travel to for business on a regular basis. You go to this gym

exactly once a week: every Monday night when you're in town. Two options given were: The membership allows you to pay \$20 (non-refundable) per each visit./The membership costs you \$1000 a year (i.e., roughly \$19.23 per visit).

One Monday, just after you have arrived and changed, you receive a phone call that requires you to leave and forgo your exercise that evening. How well does each of the following statements capture your feelings about the cost of the missed workout? Please indicate your response on a scale 1 (strongly agree) to 5 (strongly disagree).

- a. I feel like I just wasted \$20, which is what I paid to work out that evening. (1.98/2.70)
- b. I feel like I wasted nothing, since my visit had already been paid for, and sometimes you just can't take full advantage of things to which you are entitled. (3.77/3.03)
- c. I feel like I wasted something, perhaps time, perhaps money, but no specific amount or measure readily comes to mind. (2.95/2.40)

The pair of numbers in the brackets separated by a forward slash, just after each option, is the actual mean choice when one pays a per visit fee and a per year fee respectively. How does framing alter the results? Should gyms encourage patrons to buy tickets on a per visit basis or a per year basis? Why should anyone buy a per visit basis ticket at all?

31. (Rajagopal and Rha, 2009) Two different scenarios are presented to subjects in experiments in scenario-based experiments. Payments to subjects were given in terms of course credits. Scenario A: Jim is a business management consultant who travels every week on business. Last week, he was booked on the 4:30 p.m. flight to Memphis on Sunday. After he reached the airport at 3:30 p.m., he finds out that his flight is delayed by an hour. Since his first business appointment is scheduled for Monday morning, the delay will not affect his work commitments at all.

Scenario B: Jim is a business management consultant who travels every week on business. Last week, he was booked on the 4:30 p.m. flight to Memphis on Sunday to begin a family vacation with his family who were traveling to Memphis from California. After he reached the airport at 3:30 p.m., he finds out that his flight is delayed by an hour. Since his family is scheduled to join him in Memphis the next day, the delay will not affect their holiday plans at all.

Subjects in Experiments were asked to rate on a scale of 1–7, how upset Jim would be in each of the two scenarios. Letting M_i , $i = A, B$ be the mean numbers representing how upset Jim is in each scenario, the findings were as follows. $M_A = 4.54$ (1.34), $M_B = 3.37$ (1.49), $F(1, 45) = 7.76$, $p < .05$; figures in brackets represent the standard deviations.

How do these results inform our view about the fungibility of time across various uses?

32. Prove Proposition 21.8.

REFERENCES FOR PART 7

- Abeler, J. and Marklein, F. (2013). Fungibility, labels, and consumption. Mimeo, Oxford University.
- Acemoglu, D., Reed, T., and Robinson, J. A. (2013). Chiefs: economic development and elite control of civil society in Sierra Leone. Working Paper 18691, National Bureau of Economic Research, Cambridge, MA.
- Ajzen, I. (1977). Intuitive theories of events and the effects of base-rate information on prediction. *Journal of Personality and Social Psychology* 35: 303–14.
- Akerlof, G. A. and Shiller, R. J., (2008). *Animal Spirits: How Human Psychology Drives the Economy and Why It Matters for Global Capitalism*. Princeton: Princeton University Press.
- Aktas, N., de Bodt, E., Bollaert, H., and Roll, R. (2015). CEO narcissism and the takeover process: from private initiation to deal completion. *Journal of Financial and Quantitative Analysis* in press.
- al-Nowaihi, A. and Dhami, S. (2015). Evidential equilibria: heuristics and biases in static games of complete information. *Games* 6: 637–76.
- Alesina, A., Giuliano, P., and Nunn, N. (2013). On the origins of gender roles: women and the plough. *Quarterly Journal of Economics* 128(2): 469–530.
- Alhakami, A. S. and Slovic, P. (1994). A psychological study of the inverse relationship between perceived risk and perceived benefit. *Risk Analysis* 14: 1085–96.
- Allcott, H. and Mullainathan, S. (2010). Behavior and energy policy. *Science* 327(5970): 1204–05.
- Allcott, H. and Rogers, T. (2014). The short-run and long-run effects of behavioral interventions: experimental evidence from energy conservation. *American Economic Review* 104(10): 3003–37.
- Allen, F. and Faulhaber, G. (1989). Signaling by underpricing in the IPO market. *Journal of Financial Economics* 23: 303–23.
- Alicke, M. D. (1985). Global self-evaluation as determined by the desirability and controllability of trait adjectives. *Journal of Personality and Social Psychology* 49: 1621–30.
- Anderson, J. C., Lowe, D. J., and Reckers, P. M. J. (1993). Evaluation of auditor decisions: hindsight bias effects and the expectation gap. *Journal of Economic Psychology* 14: 711–37.
- Andrade, S. C., Chang, C., and Seasholes, M. S. (2008). Trading imbalances, predictable reversals, and crossstock price pressure. *Journal of Financial Economics* 88: 406–23.
- Ang, J. S. and Cheng, Y. (2006). Direct evidence on the market-driven acquisition theory. *Journal of Financial Research* 29: 199–216.
- Ariely D., Loewenstein G., and Prelec D. (2003). Coherent arbitrariness: stable demand curves without stable preferences. *Quarterly Journal of Economics* 118(1): 73–105.
- Ariely, D., Loewenstein, G., and Prelec, D. (2006). Tom Sawyer and the construction of value. *Journal of Economic Behavior and Organization* 60: 1–10.
- Arkes, H. R. (1991). Costs and benefits of judgment errors: implications for debiasing. *Psychological Bulletin* 110: 486–98.
- Arkes, H. R. and Blumer, C. (1985). The psychology of sunk cost. *Organizational Behavior and Human Decision Processes* 35(1): 124–40.
- Arkes, H. R., Faust, D., Guilmette, T. J., and Hart, K. (1988). Eliminating the hindsight bias. *Journal of Applied Psychology* 73: 305–7.
- Arkes, H. R., Joyner, C. A., Pezzo, M. V., Nash, J. G., Siegel-Jacobs, K., and Stone, E. (1994). The psychology of windfall gains. *Organizational Behavior and Human Decision Processes* 59(3): 331–47.
- Ayton, P. and Fischer, I. (2004). The hot hand fallacy and the gambler's fallacy: two faces of subjective randomness? *Memory and Cognition* 32: 1369–78.
- Ayton, P. and Hardman, D. (1997). Are two rationalities better than one? *Current Psychology of Cognition* 16: 39–51.
- Babcock, L. and Pogarsky, G. (1999). Damage caps and settlement: a behavioral approach. *Journal of Legal Studies* 28: 341–70.
- Baker, H. K. and Nofsinger, J. R. (2010). *Behavioral Finance: Investors, Corporations, and Markets*. Hoboken, NJ: John Wiley & Sons, Inc.
- Baker, M., Nagel, S., and Wurgler, J. (2007). Effect of dividends on consumption. *Brookings Papers on Economics Activity* 1: 277–91.

- Baker, M., Stein, J., and Wurgler, J. (2003). When does the market matter? Stock prices and the investment of equity-dependent firms. *Quarterly Journal of Economics* 118: 969–1006.
- Baker, M. and Wurgler, J. (2011). Dividends as reference points: a behavioral signaling model. Mimeo, NYU.
- Baker, M. and Wurgler, J. (2012). Behavioral corporate finance: a current survey. In G. M. Constantinides, M. Harris, and R. M. Stulz (eds.), *Handbooks in the Economics of Finance, Volume 2*. New York: Elsevier.
- Banerjee, A. and Duflo, E. (2011). *Poor Economics : A Radical Rethinking of the Way to Fight Global Poverty*. Public Affairs.
- Banerjee, A. and Mullainathan, S. (2008). Limited attention and income distribution. *American Economic Review* 98(2): 489–93.
- Barber, B. M., Lee, Y., Liu, Y., and Odean, T. (2007). Is the aggregate investor reluctant to realize losses? Evidence from Taiwan. *European Financial Management* 13: 423–47.
- Barber, B. M., Lee, Y., Liu, Y., and Odean, T. (2009). Just how much do individual investors lose by trading? *Review of Financial Studies* 22: 609–32.
- Barber, B. M., Lee, Y., Liu, Y., and Odean, T. (2014). The cross-section of speculator skill: evidence from Taiwan. *Journal of Financial Markets* 18(c): 1–24.
- Barber, B. M. and Odean, T. (2000). Trading is hazardous to your wealth: the common stock investment performance of individual investors. *Journal of Finance* 55: 773–806.
- Barber, B. M. and Odean, T. (2001). Boys will be boys: gender, overconfidence, and common stock investment. *Quarterly Journal of Economics* 116: 261–92.
- Barber, B. M. and Odean, T. (2008). All that glitters: the effect of attention on the buying behavior of individual and institutional investors. *Review of Financial Studies* 21: 785–818.
- Barber, B. M. and Odean, T. (2013). The behavior of individual investors. Chapter 22 in G. Constantinides, M. Harris, and R. Stulz (eds.), *Handbook of the Economics of Finance*. Elsevier, pp. 1533–69.
- Barber, B. M., Odean, T., and Zhu, N. (2009). Do retail trades move markets? *Review of Financial Studies* 22: 151–86.
- Barberis, N. (2013). The psychology of tail events: progress and challenges. *American Economic Review* 103(3): 611–16.
- Barberis, N. and Huang, M. (2001). Mental accounting, loss aversion, and individual stock returns. *Journal of Finance* 56(4): 1247–92.
- Barberis, N. and Huang, M. (2008). Stocks as lotteries: the implications of probability weighting for security prices. *American Economic Review* 98(5): 2066–100.
- Barberis, N., Huang, M., and Thaler, R. H. (2006). Individual preferences, monetary gambles, and stock market participation: a case for narrow framing. *American Economic Review* 96(4): 1069–90.
- Barberis, N., Shleifer, A., and R. Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics* 49: 307–43.
- Barberis, N. and Thaler, R. H. (2003). A survey of behavioral finance. In G. Constantinides, M. Harris, and R. Stulz (eds.), *Handbook of the Economics of Finance*. Amsterdam: North-Holland, pp. 1051–119.
- Barberis, N. and Xiong, W. (2009). What drives the disposition effect? An analysis of a long-standing preference-based explanation. *Journal of Finance* 64: 751–84.
- Barbey, A. K. and Sloman, S. A. (2007). Base-rate respect: from ecological rationality to dual processes. *Behavioral and Brain Sciences* 30(3): 241–54.
- Bar-Hillel, M. (1980). The base rate fallacy in probability judgements. *Acta Psychologica* 44: 211–33.
- Bar-Hillel, M. and Wagenaar, W. A. (1991). The perception of randomness. *Advances in Applied Mathematics* 12: 428–54.
- Barsky, R. and De Long, B. (1993). Why does the stock market fluctuate? *Quarterly Journal of Economics* 107: 291–311.
- Barth, J., Caprio, G., and Levine, R. (2013). Bank regulation and supervision in 180 Countries from 1999 to 2011. *Journal of Financial Economic Policy* 5(2): 2.
- Basu, K. (1997). Why are so many goods priced to end in nine? And why this practice hurts the producers. *Economics Letters* 54(1): 41–4.
- Basu, K. (2006). Consumer cognition and pricing in the nines in oligopolistic markets. *Journal of Economics and Management Strategy* 15(1): 125–41.

- Bebchuk, L. A., Fried, J. M., and Walker, D. I. (2002). Managerial power and rent extraction in the design of executive compensation. *University of Chicago Law Review* 69(3): 751–846.
- Becker, G. S. (1981). *A Treatise on the Family*. Cambridge, MA, Harvard University Press.
- Benartzi, S. (2001). Excessive extrapolation and the allocation of 401(k) accounts to company stock? *Journal of Finance* 56(5): 1747–64.
- Benartzi, S., Michaely, R., and Thaler, R. H. (1997). Do changes in dividends signal the future or the past? *Journal of Finance* 52: 1007–34.
- Benartzi, S. and Thaler, R. H. (1995). Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics* 110(1): 73–92.
- Benartzi, S. and Thaler, R. H. (2001). Naive diversification strategies in retirement saving plans. *American Economic Review* 91(1): 79–98.
- Benartzi, S. and Thaler, R. H. (2004). Save more tomorrow: using behavioral economics to increase employee savings. *Journal of Political Economy* 112(1): Part 2, S164–S187.
- Benartzi, S. and Thaler, R. H. (2013). Behavioral economics and the retirement savings crisis. *Science* 339: 1152–53.
- Ben-David, I., Graham, J., and Harvey, C. (2013). Managerial miscalibration. *Quarterly Journal of Economics* 128(4): 1547–84.
- Benoit, J. P., Dubra, J., and Moore D. A. (2015). Does the better-than-average effect show that people are overconfident? Two experiments. *Journal of the European Economic Association* 13(2): 293–329.
- Bergman, O., Ellingsen, T., Johannesson, M., and Svensson, C. (2010). Anchoring and cognitive ability. *Economics Letters* 107: 66–68.
- Bernard, T. and Taffesse, A. S. (2014). Aspirations: an approach to measurement with validation using Ethiopian data. *Journal of African Economies* 23(2): 189–224.
- Bernard, V. L. and Thomas, J. K. (1989). Post-earnings announcement drift: delayed price response or risk premium? *Journal of Accounting Research (Supplement)* 1–36.
- Bernard, V. L. and Thomas, J. K. (1990). Evidence that stock prices do not fully reflect the implications of current earnings for future earnings. *Journal of Accounting and Economics* 13: 305–40.
- Bernstein, D. M., Erdfelder, E., Meltzoff, A. N., Peria, W., and Loftus, G. R. (2011). Hindsight bias from 3 to 95 years of age. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 37: 378–91.
- Bertrand, M., Karlan, D., Mullainathan, S., Shafir, E., and Zinman, J. (2010). What's advertising content worth? Evidence from a consumer credit marketing field experiment. *Quarterly Journal of Economics* 125(1): 263–305.
- Bertrand, M. and Morse, A. (2011). Information disclosure, cognitive biases, and payday borrowing. *Journal of Finance* 66(6): 1865–93.
- Bertrand, M. and Mullainathan, S. (2001). Are CEOs rewarded for luck? The ones without principals are. *Quarterly Journal of Economics* 116(3): 901–32.
- Beshears, J., Choi, J. J., Laibson, D., and Madrian, B. C. (2013a). Simplification and saving. *Journal of Economic Behavior and Organization* 95: 130–45.
- Beshears, J., Choi J. J., Fuster, A., Laibson D., and Madrian B. C. (2013b). What goes up must come down? Experimental evidence on intuitive forecasting. *American Economic Review* 103(3): 570–4.
- Betsch, T. and Pohl, D. (2002). Tversky and Kahneman's availability approach to frequency judgement: a critical analysis. In P. Sedlmeier and T. Betsch (eds.), *Etc.: Frequency Processing and Cognition*. Oxford: Oxford University Press, pp. 109–19.
- Bettman, J. R. and Weitz B. A. (1983). Attributions in the board room: causal reasoning in corporate annual reports. *Administrative Science Quarterly* 28(2): 165–83.
- Bewley, T. F. (1998). Why not cut pay? *European Economic Review* 42: 459–90.
- Bhattacharya, S. (1979). Imperfect information dividend policy and the bird in the hand fallacy. *Bell Journal of Economics and Management Science* 10: 259–70.
- Biais, B. and Weber, M. (2009). Hindsight bias, risk perception and investment performance. *Management Science* 55(6): 1018–29.
- Black, F. (1986) Noise. *Journal of Finance* 41(3): 529–43.
- Blanchard, O. and Fischer, S. (1989). *Lectures on Macroeconomics*. Cambridge, MA: MIT Press.
- Bocskocsky, A., Ezekowitz, J., and Stein, C. (2014). The hot hand: a new approach to an old fallacy. Paper presented at the MIT Sloan Sports Analytics Conference.

- Bonini, N., Tentori, K., and Osherson, D. (2004). A different conjunction fallacy. *Mind and Language* 19: 199–210.
- Börgers, T. (1996). On the relevance of evolution and learning to economic theory. *Economic Journal* 106(438): 1274–385.
- Börgers, T. and Sarin, R. (2000). Naïve reinforcement learning with endogenous aspirations. *International Economic Review* 41: 921–50.
- Borges, B., Goldstein, D. G., Ortmann, A., and Gigerenzer, G. (1999). Can ignorance beat the stock market? In G. Gigerenzer, P. M. Todd, and the ABC Research Group (eds.), *Simple Heuristics That Make Us Smart*. New York: Oxford University Press, pp. 59–72.
- Boyer, B., Mitton, T., and Vorkink, K. (2010). Expected idiosyncratic skewness. *Review of Financial Studies* 23(1): 169–202.
- Brase, G. L. (2002a). Which statistical formats facilitate what decisions? The perception and influence of different statistical information formats. *Journal of Behavioral Decision Making* 15: 381–401.
- Brase, G. L. (2002b). Ecological and evolutionary validity: Comments on Johnson-Laird, Legrenzi, Girotto, Legrenzi, and Caverni's (1999) mental-model theory of extensional reasoning. *Psychological Review* 109: 722–8.
- Brase, G. L., Fiddick, L., and Harries, C. (2006). Participant recruitment methods and statistical reasoning performance. *Quarterly Journal of Experimental Psychology* 59: 965–76.
- Brav, A., Graham, J., Harvey, C. R., and Michaely, R. (2005). Payout policy in the 21st century. *Journal of Financial Economics* 77: 483–528.
- Breiman, L., Friedman, J. H., Olshen, R. A., and Stone, C. J. (1993). *Classification and Regression Trees*. New York: Chapman and Hall.
- Brennan, G. and Buchanan, J. (1980). *The Power to Tax: Analytical Foundations of a Fiscal Constitution*. Cambridge: Cambridge University Press.
- Brown, J., Hossain, T., and Morgan, J. (2010). Shrouded attributes and information suppression: evidence from the field. *Quarterly Journal of Economics* 125(2): 859–76.
- Brown, P., Chappel, N., da Silva Rosa, R., and Walter, T. (2006). The reach of the disposition effect: Large sample evidence across investor classes. *International Review of Finance* 6: 42–78.
- Bryant, F. B. and Brockway, J. H. (1997). Hindsight bias in reaction to the verdict in the O. J. Simpson criminal trial. *Basic and Applied Social Psychology* 19: 225–41.
- Buchanan, J. (1967). *Public Finance in Democratic Process: Fiscal Institutions and the Individual Choice*. Chapel Hill: University of North Carolina Press.
- Budescu, D. V. and Rapoport, A. (1994). Subjective randomization in one- and two-person games. *Journal of Behavioral Decision Making* 7: 261–78.
- Burch, T. R., Christie, W. G., and Nanda, V. (2004). Do firms time equity offerings? Evidence from the 1930s and 1940s. *Financial Management* 33: 5–23.
- Busse, M. R., Lacetera, N., Pope, D. G., Silva-Risso, J., and Sydnor, J. R. (2013). Estimating the effect of salience in wholesale and retail car markets. *American Economic Review* 103(3): 575–9.
- Camerer, C. F. (1987). Do biases in probability judgment matter in markets? Experimental evidence. *American Economic Review* 77: 981–97.
- Camerer, C. F. (1989). Does the basketball market believe in the “Hot Hand”? *American Economic Review* 79: 1257–61.
- Camerer, C. F. (1990). Do markets correct biases in probability judgment? Evidence from market experiments. In J. Kagel and L. Green (eds.), *Advances in Behavioral Economics Volume 2*, Northwood, NJ: Ablex Publishers, pp. 125–72.
- Camerer, C. F. (1995). Individual decision making. In J. Kagel and A. E. Roth (eds.), *Handbook of Experimental Economics*. Princeton, NJ: Princeton University Press.
- Camerer, C. F., Babcock, L., Loewenstein, G., and Thaler, R. (1997). Labor supply of New York city cab drivers: one day at a time. *Quarterly Journal of Economics* 112: 407–41.
- Camerer, C. F., Loewenstein, G., and Weber, M. (1989). The curse of knowledge in economic settings: an experimental analysis. *Journal of Political Economy* 97: 1232–54.
- Camerer, C. F. and Lovo, D. (1999). Overconfidence and excess entry: an experimental approach. *American Economic Review* 89: 306–18.
- Campbell, J. Y. and Kyle, A. S. (1993). Smart money, noise trading and stock price behaviour. *Review of Economic Studies* 60(1): 1–34.
- Caplin, A., Dean, M., and Martin, D. (2011). Search and satisficing. *American Economic Review* 101(7): 2899–922.

- Casscells, W., Schoenberger, A., and Grapoys, T. B. (1978). Interpretation by physicians of clinical laboratory results. *New England Journal of Medicine* 299: 999–1000.
- Chan, L., Jegadeesh, N., and Lakonishok, J. (1997). Momentum strategies. *Journal of Finance* 51: 1681–1713.
- Chapman, G. B. and Bornstein, B. H. (1996). The more you ask for, the more you get: anchoring in personal injury verdicts. *Applied Cognitive Psychology* 10: 519–40.
- Chapman, G. B. and Johnson, E. J. (1994). The limits of anchoring. *Journal of Behavioural Decision Making* 7: 223–42.
- Chapman, G. B. and Johnson, E. J. (2002). Incorporating the irrelevant: Anchors in judgments of belief and value. In T. Gilovich, D. Griffin, and D. Kahneman (eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment*. Cambridge: Cambridge University Press, pp. 120–38.
- Charness, G., Karni, E., and Levin, D. (2007). Individual and group decision making under risk: an experimental study of Bayesian updating and violations of first-order stochastic dominance. *Journal of Risk and Uncertainty* 35: 129–48.
- Charness, G., Karni, E., and Levin, D., (2010). On the conjunction fallacy in probability judgment: new experimental. *Games and Economic Behavior* 68: 551–56.
- Cheema, A. and Soman, D. (2006). Malleable mental accounting: the effect of flexibility on the justification of attractive spending and consumption decisions. *Journal of Consumer Psychology* 16: 33–44.
- Chen, G., Kim, K., Nofsinger, J., and Rui, O. (2007). Trading performance, disposition effect, overconfidence, representativeness bias, and experience of emerging market investors. *Journal of Behavioral Decision Making* 20: 425–51.
- Cherry, T. (2001). Mental accounting and other-regarding behavior: evidence from the lab. *Journal of Economic Psychology* 22: 605–15.
- Chetty, R., Looney, A., and Kroft, K. (2009). Salience and taxation: theory and evidence. *American Economic Review* 99(4): 1145–77.
- Chetty, R. and Saez, E. (2013). Teaching the tax code: earnings responses to an experiment with EITC recipients. *American Economic Journal: Applied Economics* 5(1): 1–31.
- Chirinko, R. S. and Schaller, H. (2001). Business fixed investment and “bubbles”: The Japanese case. *American Economic Review* 91: 663–80.
- Cho, I. K. and Matsui, A. (2005). Learning aspiration in repeated games. *Journal of Economic Theory* 124(2): 171–201.
- Choe, H. and Eom, Y. (2009). The disposition effect and investment performance in the futures market. *Journal of Futures Markets* 29: 496–522.
- Choi, J.J., Laibson, D., and Madrian B.C. (2009). Mental accounting in portfolio choice: evidence from a flypaper effect. *American Economic Review* 99(5): 2085–95.
- Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2002). Defined contribution pensions: plan rules, participant decisions, and the path of least resistance. In J. Poterba (ed.), *Tax Policy and the Economy* 16. Cambridge, MA: MIT Press, pp. 67–114.
- Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2004). For better or for worse: default effects and 401(k) savings behavior. In D. Wise (ed.), *Perspectives in the Economics of Aging*. Chicago: University of Chicago Press, pp. 81–121.
- Chopra, N., Lakonishok, J., and Ritter, J. (1992). Measuring abnormal performance: do stocks overreact? *Journal of Financial Economics* 31: 235–68.
- Chung, K. L. (1968). *A Course in Probability Theory*. New York: Harcourt.
- Clotfelter, C. T. and Cook, P. J. (1989). *Selling Hope: State Lotteries in America*. Cambridge, MA: Harvard University Press.
- Clotfelter, C. T. and Cook, P. J. (1993). The gambler’s fallacy in lottery play, *Management Science*. 39: 1521–25.
- Cohen, L. J. (1981). Can human irrationality be experimentally demonstrated? *Behavioral and Brain Sciences* 4: 37–70.
- Cohen, L. and Frazzini, A. (2008). Economic links and predictable returns. *Journal of Finance* 63: 1977–2011.
- Colman, A. M., Pulford, B. D., Omtzigt, D., and al-Nowaihi, A. (2010). Learning to cooperate without awareness in multiplayer minimal social situations. *Cognitive Psychology* 61: 201–27.
- Cooper, A. C., Woo, C. Y., and Dunkelberg, W. C. (1988). Entrepreneurs’ perceived chances for success. *Journal of Business Venturing* 3: 97–108.
- Cooper, I. and Kaplanis, E. (1994). Home bias in equity portfolios, inflation hedging, and international capital market equilibrium. *Review of Financial Studies* 7: 44–60.

- Cooper, M. J., Dimitrov, O., and Rau, P. R. (2001). A rose.com by any other name. *Journal of Finance* 56: 2371–88.
- Cooper, M. J., Gulen, H., and Rau, P. R. (2005). Changing names with style: mutual fund name changes and their effects on fund flows. *Journal of Finance* 60: 2825–58.
- Cosmides, L. and Tooby, J. (1996). Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty. *Cognition* 58: 1–73.
- Courant, P., Gramlich, E., and Laitner, J. (1986). A dynamic micro estimate of the life cycle model. In H. G. Aaron and G. Burtless (eds.), *Retirement and Economic Behavior* Washington DC: Brookings Institution.
- Coval, J. D., Hirshleifer, D. A., and Shumway, T. (2005). Can individual investors beat the market? Working paper, Harvard University.
- Crocker, J. (1982). Biased questions in judgment of covariation studies. *Personality and Social Psychology Bulletin* 8: 214–20.
- Crosan, R. and Sundali, J. (2005). The gambler's fallacy and the hot hand: empirical data from casinos. *Journal of Risk and Uncertainty* 30: 195–209.
- Cutler, D., Poterba, J., and Summers, L. (1991). Speculative dynamics. *Review of Economic Studies* 58: 529–46.
- Czerlinski, J., Goldstein, D. G., and Gigerenzer, G. (1999). How good are simple heuristics? In G. Gigerenzer, P. M. Todd, and the ABC Research Group. *Simple Heuristics That Make Us Smart*. New York: Oxford University Press.
- Daniel, K., Grinblatt, M., Titman, S., and Wermers, R. (1997). Measuring mutual fund performance with characteristic-based benchmarks. *Journal of Finance* 52: 1035–58.
- Daniel, K., Hirshleifer, D., and Subrahmanyam, A. (1998). Investor psychology and security market under- and overreactions, *Journal of Finance* 53: 1839–85.
- Daniel, K., Hirshleifer, D., and Subrahmanyam, A. (2001). Overconfidence, arbitrage, and equilibrium asset pricing. *Journal of Finance* 56(3): 921–65.
- Daniel, K. and Titman, S. (1997). Evidence on the characteristics of cross-sectional variation in stock returns. *Journal of Finance* 52: 1–33.
- Darley, J. and Gross, P. (1983). A hypothesis-confirming bias in labeling effects. *Journal of Personality and Social Psychology* 44: 20–33.
- Datta, S. and Mullainathan, S. (2014). Behavioral design: a new approach to development policy. *Review of Income and Wealth* 60(1): 7–35.
- D'Avolio, G. (2002). The market for borrowing stock. *Journal of Financial Economics* 66: 271–306.
- De Bondt, W. and Thaler, R. (1985). Does the stock market overreact? *Journal of Finance* 40: 793–808.
- De Bondt, W. and Thaler, R. (1987). Further evidence on investor overreaction and stock market seasonality. *Journal of Finance* 42: 557–81.
- De Giorgi, E. G. and Legg, S. (2012). Dynamic portfolio choice and asset pricing with narrow framing and probability weighting. *Journal of Economic Dynamics and Control* 36(7): 951–72.
- De Long, J. B., Shleifer, A., Summers, L., and Waldmann, R. (1990a). Noise trader risk in financial markets. *Journal of Political Economy* 98: 703–38.
- De Long, J. B., Shleifer, A., Summers, L., and Waldmann, R. (1990b). Positive feedback investment strategies and destabilizing rational speculation, *Journal of Finance* 45: 375–95.
- DellaVigna, S. and Gentzkow, M. (2010). Persuasion: empirical evidence. *Annual Review of Economics* 2: 643–69.
- DellaVigna, S. and Pollet, J. M. (2009). Investor inattention and Friday earnings announcements. *Journal of Finance* 64(2): 709–49.
- Denes-Raj, V. and Epstein, S. (1994). Conflict between intuitive and rational processing: when people behave against their better judgment. *Journal of Personality and Social Psychology* 66: 819–29.
- Denzau, A. T. and North, D. C. (1994). Shared mental models: ideologies and institutions. *Kyklos* 47(1): 3–31.
- Dhar, R. and Zhu, N. (2006). Up close and personal: investor sophistication and the disposition effect. *Management Science* 52: 726–40.
- Diamond, P. and Hausmann, J. (1984). Individual retirement and saving behavior. *Journal of Public Economics* 23: 81–114.
- Dohmen, T., Falk, A., Huffman, D., Marklein, F., and Sunde, U. (2009). Biased probability judgment: Evidence of incidence and relationship to economic outcomes from a representative sample. *Journal of Economic Behavior and Organization* 72: 903–15.
- Dong, M., Hirshleifer, D., Richardson, S., and Teoh, S. H. (2006). Does investor misvaluation

- drive the takeover market? *Journal of Finance* 61: 725–62.
- Dorn, D. and Huberman, G. (2005). Talk and action: what individual investors say and what they do. *Review of Finance* 9: 437–81.
- Dorn, D. and Sengmueller, P. (2009). Trading as entertainment. *Management Science* 55: 591–603.
- Drexler, A., Fischer, G., and Schoar, A. (2014). Keeping it simple: financial literacy and rules of thumb. *American Economic Journal: Applied Economics* 6(2): 1–31.
- Duffee, G. R. (1996). Idiosyncratic variation of Treasury bill yields. *Journal of Finance* 51: 527–51.
- Dupas, P. and Robinson, J. (2013). Why don't the poor save more? Evidence from health savings experiments. *American Economic Review* 103(4): 1138–71.
- Eddy, D. M. (1982). Probabilistic reasoning in clinical medicine: problems and opportunities. In: D. Kahneman, P. Slovic, and A. Tversky. *Judgment Under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press, pp. 249–67.
- Edwards, W. (1968). Conservatism in human information processing. In: B. Kleinmütz (ed.), *Formal Representation of Human Judgment*. Wiley, pp. 17–52.
- Edwards, W. and Tversky, A., (1967). *Decision Making: Selected Readings*. Penguin.
- Einhorn, H. J. and Hogarth, R. M. (1981). Behavioral decision theory: processes of judgment and choice. *Annual Review of Psychology* 32: 53–88.
- Eisenman, R. (1993). Belief that drug usage in the United States is increasing when it is really decreasing: an example of the availability heuristic. *Bulletin of the Psychonomic Society* 31: 249–52.
- Ellison, G. (2005). A model of add-on pricing. *Quarterly Journal of Economics* 120: 585–638.
- Elster, J. (1979). *Ulysses and the Sirens*. Cambridge: Cambridge University Press.
- Engelberg, J. and Parsons, C. (2011). The causal impact of media in financial markets. *Journal of Finance* 66: 67–97.
- Engelberg, J., Sasseville, C., and Williams, J. (2011). Market madness: the case of mad money. *Management Science* 58: 351–64.
- Englich, B. and Mussweiler, T. (2001). Sentencing under uncertainty: anchoring effects in the courtroom. *Journal of Applied Social Psychology* 31: 1535–51.
- Englich, B., Mussweiler, T., and Strack, F. (2006). Playing dice with criminal sentences: the influence of irrelevant anchors on experts' judicial decision making. *Personality and Social Psychology Bulletin* 32(2): 188–200.
- Englmaier, F., Schmöllerz, A., and Stowasser, T. (2014). Price discontinuities in an online market for used cars. Mimeo, University of Munich.
- Epley, N. and Gilovich, T. (2001). Putting adjustment back in the anchoring and adjustment heuristic: differential processing of self-generated and experimenter-provided anchors. *Psychological Science* 12: 391–6.
- Epley, N. and Gilovich, T. (2006). The anchoring-and-adjustment heuristic: why the adjustments are insufficient. *Psychological Science* 17: 311–18.
- Erel I., Julio, B., Kim, W., and Weisbach, M. S. (2012). Macroeconomic conditions and capital raising. *Review of Financial Studies* 25(2): 341–76.
- Erev, I. and Roth, A. E. (2001). On simple reinforcement learning models and reciprocation in the prisoner dilemma game. In: G. Gigerenzer, and R. Selten (eds.), *The Adaptive Toolbox*. Cambridge, MA: MIT Press, pp. 215–32.
- Ericson, K. M. M., White, J. M., Laibson, D., and Cohen, J. D. (2015). Money earlier or later? Simple heuristics explain intertemporal choices better than delay discounting. Forthcoming in *Psychological Science*.
- Evans, J. Handley, S. J., Perham, N. Over, D. E., and Thompson, V. A. (2000). Frequency versus probability formats in statistical word problems. *Cognition* 77: 197–213.
- Fama, E. F. (1965). The behavior of stock market prices. *Journal of Business* 38: 34–105.
- Fama, E. (1970). Efficient capital markets: a review of theory and empirical work. *Journal of Finance* 25: 383–417.
- Fama, E., Fisher, L., Jensen, M., and Roll, R. (1969). The adjustment of stock prices to new information. *International Economic Review* 10: 1–21.
- Fama, E. and French, K. (1992). The cross-section of expected stock returns. *Journal of Finance* 47: 427–65.
- Fama, E. and French, K. (1993). Common risk factors in the returns of bonds and stocks, *Journal of Financial Economics* 33: 3–56.

- Fama, E. and French, K. (1996a). Multifactor explanations of asset pricing anomalies. *Journal of Finance* 51: 55–84.
- Fama, E. and French, K. (1996b). The CAPM is wanted, dead or alive. *Journal of Finance* 51(5): 1947–58.
- Fama, E. and French, K. (2015). A five-factor asset pricing model. *Journal of Financial Economics* 116: 1–22.
- Feng, L. and Seasholes, M. (2005). Do investor sophistication and trading experience eliminate behavioral biases in financial markets? *Review of Finance* 9, 305–51.
- Ferraro P. J. and Price, M. K. (2013). Using nonpecuniary strategies to influence behavior: evidence from a large-scale field experiment. *Review of Economics and Statistics* 95(1): 64–73.
- Festinger, L. (1957). *A Theory of Cognitive Dissonance*. Evanstone, IL: Row, Peterson.
- Festinger, L. and Carlsmith, J. M. (1959). Cognitive consequences of forced compliance. *Journal of Abnormal and Social Psychology* 58: 203–11.
- Figlewski, S. (1979). Subjective information and market efficiency in a betting model. *Journal of Political Economy* 87: 75–88.
- Finkelstein, A. (2009). E-Z tax: tax salience and tax rates. *Quarterly Journal of Economics* 124(3): 969–1010.
- Finucane, M. L., Alhakami, A., Slovic, P., and Johnson, S. M. (2000). The affect heuristic in judgments of risks and benefits. *Journal of Behavioral Decision Making* 13: 1–17.
- Fischhoff, B. (1975). Hindsight foresight: the effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance* 1: 288–99.
- Fischhoff, B. (1982). Debiasing. In D. Kahneman, P. Slovic, and A. Tversky (eds.), *Judgment Under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press, pp. 422–44.
- Fischhoff, B. and Beyth, R. (1975). I knew it would happen—remembered probabilities of once-future things. *Organizational Behavior and Human Performance* 13: 1–16.
- Fischhoff, B., Gonzalez, R. M., Lerner, J. S., and Small, D. A. (2005). Evolving judgments of terror risks: foresight, hindsight, and emotion. *Journal of Experimental Psychology: Applied* 11: 124–39.
- Fischhoff B., Slovic, Lichtenstein S., Read S., and Combs B. (1978). How safe is safe enough? A psychometric study of attitudes toward technological risks and benefits. *Policy Sciences* 9: 127–52.
- Frascara, J. (1999). Cognition, emotion and other inescapable dimensions of human experience. *Visible Language* 33: 74–87.
- Frazzini, A. (2006). The disposition effect and underreaction to news. *Journal of Finance* 61: 2017–46.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic Perspectives* 19: 25–42.
- French, K. R. (2008). The cost of active investing. *Journal of Finance* 63: 1537–73.
- French, K. R. and Poterba, J. M. (1991). International diversification and international equity markets. *American Economic Review* 81: 222–6.
- Friedman, M. (1953). The methodology of positive economics. In *Essays in Positive Economics* Chicago: University of Chicago Press, pp. 3–43.
- Froot, K. and Dabora, E. (1999). How are stock prices affected by the location of trade? *Journal of Financial Economics* 53: 189–216.
- Frydman, C. and Rangel, A. (2014). Debiasing the disposition effect by ...saliency of information about a stock's purchase price. *Journal of Economic Behavior and Organization* 107: 541–52.
- Fryer, R. and Jackson, M. (2008). A categorical model of cognition and biased decision making. *BE Journal of Theoretical Economics* 8(1).
- Fudenberg, D., Levine, D. K., Maniadis, Z. (2012). On the robustness of anchoring effects in WTP and WTA experiments. *American Economic Journal: Microeconomics* 4(2): 131–45.
- Fuster, A., Laibson, D., and Mendel, B. (2010). Natural expectations and macroeconomic fluctuations. *Journal of Economic Perspectives* 24(4): 67–84.
- Gabaix, X. (2014). A sparsity-based model of bounded rationality. *Quarterly Journal of Economics* 129(4): 1661–710.
- Gabaix, X. and Laibson, D. (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics* 121(2): 505–40.
- Galbraith, J. K. (1997). *The Great Crash* (1929). Boston, MA: Houghton Mifflin Company.
- Galinsky, A. D. and Mussweiler, T. (2001). First offers as anchors: the role of perspective-taking and negotiator focus. *Journal of Personality and Social Psychology* 81(4): 657–69.

- Galton, F. (1886). Regression towards mediocrity in hereditary stature. *Journal of the Anthropological Institute of Great Britain and Ireland* 15: 246–63.
- Gao, X. and Lin, T.-C. (2014). Do individual investors treat trading as a fun and exciting gambling activity? Evidence from repeated natural experiments. Forthcoming in *Review of Financial Studies*.
- Gervais, S. and Goldstein, I. (2007). The positive effects of biased self-perceptions in firms. Available at SSRN: <http://ssrn.com/abstract=963249>.
- Gervais, S., Heaton, J. B., and Odean, T. (2011). Overconfidence, compensation contracts, and capital budgeting. *Journal of Finance* 66: 1735–77.
- Gervais, S. and Odean, T. (2001). Learning to be overconfident. *Review of Financial Studies* 14: 1–27.
- Gigerenzer, G. (1991). How to make cognitive illusions disappear: beyond heuristics and biases. *European Review of Social Psychology* 2: 83–115.
- Gigerenzer, G. (1993). The bounded rationality of probabilistic mental models. In K. Manktelow and D. Over (eds.), *Rationality: Psychological and Philosophical Perspectives*. Routledge.
- Gigerenzer, G. (1994). Why the distinction between single-event probabilities and frequencies is important for psychology (and vice versa). In G. Wright and P. Ayton (eds.), *Subjective Probability*. Wiley.
- Gigerenzer, G. (1996a). On narrow norms and vague heuristics: a reply to Kahneman and Tversky. *Psychological Review* 103: 592–6.
- Gigerenzer, G. (1996b). The psychology of good judgment: frequency formats and simple algorithms. *Medical Decision Making* 16: 273–80.
- Gigerenzer, G. and Gaissmaier, W. (2011). Heuristic decision making. *Annual Review of Psychology* 62: 451–85.
- Gigerenzer, G. and Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review* 103: 650–69.
- Gigerenzer, G., Hell, W., and Blank, H. (1988). Presentation and content: the use of base rates as a continuous variable. *Journal of Experimental Psychology: Human Perception and Performance* 4(3): 513–25.
- Gigerenzer, G. and Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review* 102: 684–704.
- Gigerenzer, G., Hoffrage, U., and Kleinbölting, H. (1991). Probabilistic mental models: a Brunswikian theory of confidence. *Psychological Review* 98: 506–28.
- Gigerenzer, G. and Selten, R. (2001). Rethinking rationality. In G. Gigerenzer and R. Selten (eds.), *Bounded Rationality: The Adaptive Toolbox*. Cambridge, MA: MIT Press, pp. 1–12.
- Gigerenzer, G. and Selten, R. (eds.). (2001). *Bounded Rationality: The Adaptive Toolbox*. Cambridge, MA: MIT Press.
- Gigerenzer, G., Swijtink, Z., Porter, T., Daston, L. J., Beatty, J., and Krüger, L. (1989). *The Empire of Chance: How Probability Changed Science and Everyday Life*. Cambridge: Cambridge University Press.
- Gigerenzer, G., Todd, P. M., and the ABC Research Group (1999). *Simple Heuristics That Make Us Smart*. New York: Oxford University Press.
- Gilbert, T., Kogan, S., Lochstoer, L. A., and Ozyildirim, A. (2012). Investor inattention and the market impact of summary statistics. *Management Science* 58(2): 336–50.
- Gilchrist, S., Himmelberg, C., and Huberman, G. (2005). Do stock price bubbles influence corporate investment? *Journal of Monetary Economics* 52: 805–27.
- Gilovich, T., Griffin, D., and Kahneman, D. (2002). *Heuristics and Biases: The Psychology of Intuitive Judgment*. New York: Cambridge University Press.
- Gilovich, T., Vallone, R., and Tversky, A. (1985). The hot hand in basketball: on the misperception of random sequences. *Cognitive Psychology* 17: 295–314.
- Giné, X., Cuellar, C. M., and Mazer, R. K. (2014). Financial dis-information: evidence from an audit study in Mexico. Policy Research Working Paper 6902, World Bank, Washington, DC.
- Gintis, H. (2009). *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences*. Princeton, NJ: Princeton University Press.
- Giroto, V. and Gonzalez, M. (2001). Solving probabilistic and statistical problems: a matter of information structure and question form. *Cognition* 78: 247–76.
- Glaser, M., Langer, T., Reyniers, J., and Weber, M., (2007a). Framing effects in stock market forecasts: the difference between asking for prices and asking for returns. *Review of Finance* 11: 325–57.

- Glaser, M., Langer, T., and Weber, M., (2007b). On the trend recognition and forecasting ability of professional traders. *Decision Analysis* 4: 176–93.
- Glaser, M., Nöth, M., and Weber, M. (2004). Behavioral finance. In D. J. Koehler and N. Harvey (eds.), *Blackwell Handbook of Judgment and Decision Making*. Blackwell, pp. 527–46.
- Glaser, M. and Weber, M. (2007). Overconfidence and trading volume. *Geneva Risk and Insurance Revue* 32: 1–36.
- Goel, A. M. and Thakor, A. V. (2008). Overconfidence leadership selection and corporate governance. *Journal of Finance* 63: 2737–84.
- Goldin, J. and Homonoff, T. (2013). Smoke gets in your eyes: cigarette tax salience and regressivity. *American Economic Journal: Economic Policy* 5(1): 302–36.
- Goldstein, D. G. and Gigerenzer, G. (1999). The recognition heuristic: how ignorance makes us smart. In G. Gigerenzer, P. M. Todd, and the ABC Research Group (eds.), *Simple Heuristics That Make Us Smart*. New York: Oxford University Press, pp. 37–58.
- Gollaher, D. L. (2000). *Circumcision: A History of the World's Most Controversial Surgery*. New York: Basic Books.
- Gong, B. Lei, V., and Deng, P. (2013). Before and after: the impact of a real bubble crash on investors' trading behavior in the lab. *Journal of Economic Behavior and Organization* 95: 186–96.
- Goodwill, A. M., Alison, L. A., Lehmann, R., Francis, A., and Eyre, M. (2010). The impact of outcome knowledge, role, and quality of information on the perceived legitimacy of lethal force decisions in counter-terrorism operations. *Behavioral Sciences and the Law* 28: 337–50.
- Gourville, J. T. and Soman, D. (1998). Payment depreciation: the effects of temporally separating payments from consumption. *Journal of Consumer Research* 25(2): 160–74.
- Graham F. and Isaac, A. G. (2002). The behavioral life-cycle theory of consumer behavior: survey evidence. *Journal of Economic Behavior and Organization* 48(4): 391–401.
- Graham, J. R. and Harvey, C. R. (2001). The theory and practice of corporate finance: Evidence from the field. *Journal of Financial Economics* 60: 187–243.
- Graham, J. R., Harvey, C. R., and Huang, H. (2009). Investor competence, trading frequency, and home bias. *Management Science* 55: 1094–106.
- Green, B. S. and Zwiebel, J. (2013). The hot hand fallacy: cognitive mistakes or equilibrium adjustments? Evidence from baseball. Mimeo, Stanford University.
- Green, D., Jacowitz, K. E., Kahneman, D., and McFadden, D. (1998). Referendum contingent valuation, anchoring, and willingness to pay for public goods. *Resources and Energy Economics* 20: 85–116.
- Greenwood, R. and Vayanos, D. (2010). Price pressure in the government bond market. *American Economic Review* 100: 585–90.
- Grether, D. (1980). Bayes' rules as a descriptive model: the representativeness heuristic. *Quarterly Journal of Economics* 95: 537–57.
- Grether, D. M. (1992). Testing Bayes' rule and the representativeness heuristic: some experimental evidence. *Journal of Economic Behavior and Organization* 17: 31–57.
- Griffin, D. and Tversky, A. (1992). The weighing of evidence and the determinants of confidence. *Cognitive Psychology* 24(3): 411–35.
- Grinblatt, M. and Hwang, C. Y. (1989). Signaling and the pricing of new issues. *Journal of Finance* 44: 393–420.
- Grinblatt, M. and Keloharju, M. (2000). The investment behavior and performance of various investor types: a study of Finland's unique data set. *Journal of Financial Economics* 55: 43–67.
- Grinblatt, M. and Keloharju, M. (2001). What makes investors trade? *Journal of Finance* 56: 589–616.
- Grinblatt, M. and Keloharju, M. (2009). Sensation seeking, overconfidence, and trading activity. *Journal of Finance* 64: 549–78.
- Grinblatt, M., Keloharju, M., and Linnainmaa, J. (2012). IQ, trading behavior, and performance. *Journal of Financial Economics* 55: 43–67.
- Grinblatt M. and Moskowitz, T. J. (2004). Predicting stock price movements from past returns: the role of consistency and tax-loss selling. *Journal of Financial Economics* 71: 541–79.
- Grinblatt, M. and Titman, S. (1989). Mutual fund performance: an analysis of quarterly portfolio holdings. *Journal of Business* 62: 393–416.
- Grossman, S. and Hart, O. (1980). Disclosure laws and takeover bids. *Journal of Finance* 35: 323–34.
- Guilbault, R. L., Bryant, F. B., Brockway, J. H., and Posavac, E. J. (2004). A meta-analysis of research

- on hindsight bias. *Basic and Applied Social Psychology* 26: 103–17.
- Guryan, J. and Kearney, M. S. (2008). Gambling at lucky stores: empirical evidence from state lottery sales. *American Economic Review* 98(1): 458–73.
- Guttman, I., Kadan, O., and Kandel, E., (2010). Dividend stickiness and strategic pooling. *Review of Financial Studies* 23: 4455–95.
- Guyon, N. and Huillery, E. (2014). The aspiration-poverty trap: why do students from low social background limit their ambition? Evidence from France. Mimeo.
- Hall, B. J. and Murphy, K. J. (2000). Optimal exercise prices for executive stock options. *American Economic Review* 90: 209–14.
- Hall, B. J. and Murphy, K. J. (2002). Stock options for undiversified executives. *Journal of Accounting and Economics* 33: 3–42.
- Hansson, P., Rönnlund, M., Juslin, P., Nilsson, L.-G., (2008). Adult age differences in the realism of confidence judgments: overconfidence, format dependence, and cognitive predictors. *Psychology and Aging* 23: 531–44.
- Harley, E. M. (2007). Hindsight bias in legal decision making. *Social Cognition* 25: 48–63.
- Harris, L. and Gurel, E. (1986). Price and volume effects associated with changes in the SandP 500: new evidence for the existence of price pressure. *Journal of Finance* 41: 851–60.
- Harris, M. and Raviv, A. (1985). A sequential signaling model of convertible debt call policy. *Journal of Finance* 40, 1263–281.
- Harris, M. and Raviv, A. (1993). Differences of opinion make a horse race. *Review of Financial Studies* 6(3): 473–506.
- Hastie, R., Schkade, D. A., and Payne, J. W. (1999). Juror judgments in civil cases: hindsight effects on judgments of liability for punitive damages. *Law and Human Behavior* 23: 597–614.
- Hastorf, A. and Cantril, H. (1954). They saw a game: a case study. *Journal of Personality and Social Psychology* 49: 129.
- Hatsopoulos, G. N., Krugman, P. R., and Poterba, J. M. (1989). Overconsumption: The Challenge to U.S. Economic Policy, American Business Conference, Special Report for American Council for Capital Formation Center for Policy Research.
- Haushofer, J. and Fehr, E. (2014). On the psychology of poverty. *Science* 344 (6186): 862–67.
- Heath, C. (1995). Escalation and de-escalation of commitment in response to sunk costs: the role of budgeting in mental accounting, *Organizational Behavior and Human Decision Processes* 62: 38–54.
- Heath, C., Huddart, S., and Lang, M. (1999). Psychological factors and stock option exercise. *Quarterly Journal of Economics* 114: 601–27.
- Heath, C. and Tversky, A. (1991). Preference and belief: ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty* 4: 5–28.
- Heaton, J. B. (2002). Managerial optimism and corporate finance. *Financial Management* 31: 33–45.
- Heilman, C. M., Nakamoto, K., and Rao, A. G. (2002). Pleasant surprises: consumer response to unexpected in-store promotions. *Journal of Marketing Research* 39(2): 242–51.
- Henderson, V. (2012). Prospect theory, liquidation and the disposition effect. *Management Science* 58: 445–60.
- Hens, T. and Vlcek, M. (2011). Does prospect theory explain the disposition effect? *Journal of Behavioral Finance* 12: 141–57.
- Hertwig, R., Benz, B., and Krauss, S. (2008). The conjunction fallacy and the many meanings of and. *Cognition* 108: 740–53.
- Hertwig, R. and Gigerenzer, G. (1999). The conjunction fallacy revisited: how intelligent inferences look like reasoning errors. *Journal of Behavioral Decision Making* 12: 275–305.
- Hertwig, R., Pachur, T., and Kurzenhäuser, S. (2005). Judgments of risk frequencies: tests of possible cognitive mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 31: 621–42.
- Hirshleifer, D., Lim, S., and Teoh, S. (2009). Driven to distraction: extraneous events and underreaction to earnings news. *Journal of Finance* 64: 2289–325.
- Hirshleifer, D., Low, A., and Teoh, S. H. (2012). Are overconfident CEOs better innovators? *Journal of Finance* 67: 1457–98.
- Hirshleifer, D. A., Subrahmanyam, A., and Titman, S. (2003). Feedback and the success of irrational investors. Dice Center Working Paper No. 2003–22.
- Hoff, K. and Pandey, P. (2006). Discrimination, social identity, and durable inequalities. *American Economic Review* 96(2): 206–11.

- Hoff, K. and Pandey, P. (2014). Making up people: the effect of identity on performance in a modernizing society. *Journal of Development Economics* 106: 118–31.
- Hogg, R. V., McKean, J. W., and Craig, A. T. (2005). *Introduction to Mathematical Statistics*, 6th edition, Prentice Hall.
- Holbrook, R. and Stafford, F. (1971). The propensity of consume separate types of income: a generalized permanent income hypothesis. *Econometrica* 39: 1–21.
- Hong, H. and Kubik, J. D. (2003). Analyzing the analysts: career concerns and biased earnings forecasts. *Journal of Finance* 58: 313–52.
- Hong, H., Lim, T. and Stein, J. (2000). Bad news travels slowly: size, analyst coverage, and the profitability of momentum strategies. *Journal of Finance* 55: 265–95.
- Hong, H. and Stein, J. (1999). A unified theory of underreaction, momentum trading, and overreaction in asset markets. *Journal of Finance* 54: 2143–84.
- Hong, H. and Stein, J. (2007). Disagreement and the stock market. *Journal of Economic Perspectives* 21: 109–28.
- Hong, H., Stein, J., and Yu, J. (2007). Simple forecasts and paradigm shifts. *Journal of Finance* 62: 1207–42.
- Hossain, T. and Morgan, J. (2006). . . . Plus shipping and handling: revenue (non) equivalence in field experiments on eBay. *B. E. Journal of Economic Analysis and Policy: Advances in Economic Analysis and Policy* 6(2): 1–27.
- Hotz, V. J. and Scholz, J. K. (2003). The earned income tax credit. In R. Moffitt (ed.), *Means-Tested Transfer Programs in the United States*. Chicago: University of Chicago Press, pp. 141–98.
- Hsee, C. K. (1998). Less is better: when low-value options are valued more highly than high-value options. *Journal of Behavioral Decision Making* 11(2): 107–21.
- Huberman, G. and Regev, T. (2001). Contagious speculation and a cure for cancer: a non-event that made stock prices soar. *Journal of Finance* 56: 387–96.
- Hurd, M. D. (1987). Savings of the elderly and desired bequests. *American Economic Review* 77: 298–312.
- Hvidkjaer, S. (2008). Small trades and the cross-section of stock returns. *Review of Financial Studies* 21: 1123–51.
- Ikenberry, D., Lakonishok, J., and Vermaelen, T. (1995). Market underreaction to open market share repurchases. *Journal of Financial Economics* 39: 181–208.
- Ikenberry, D., Lakonishok, J., and Vermaelen, T. (2000). Stock repurchases in Canada: Performance and strategic trading. *Journal of Finance* 55: 2373–98.
- Ippolito, R. A. (1989). Efficiency with costly information: a study of mutual fund performance, 1965–1984. *Quarterly Journal of Economics* 104(1): 1–23.
- Jacowitz, K. E. and Kahneman, D. (1995). Measures of anchoring in estimation tasks. *Personality and Social Psychological Bulletin* 21: 1161–66.
- Jegadeesh, N. and Titman, S. (1993). Returns to buying winners and selling losers: implications for stock market efficiency. *Journal of Finance* 48: 65–91.
- Jensen, M. C. (1978). Some anomalous evidence regarding market efficiency. *Journal of Financial Economics* 6(2/3): 95–101.
- Jensen, R. (2012). Do labor market opportunities affect young women's work and family decisions? experimental evidence from India. *Quarterly Journal of Economics* 127(2): 753–92.
- Jensen, R. and Oster, E. (2009). The power of TV: cable television and women's status in India. *Quarterly Journal of Economics* 124(3): 1057–94.
- John, K. and Williams, J. (1985). Dividends dilution, and taxes: a signalling equilibrium. *Journal of Finance* 40: 1053–70.
- Jørgensen, C. B., Suetens, S., and Tyran, J.-R. (2011). Predicting lotto numbers. CEPR Discussion Papers 8314.
- Jovanovic, B. (1982). Truthful disclosure of information. *Bell Journal of Economics* 13: 36–44.
- Jung, K., Kim, Y. C., and Stulz, R. M. (1996). Timing, investment opportunities, managerial discretion, and the security issue decision. *Journal of Financial Economics* 42: 159–85.
- Kahneman, D. (1981). Who shall be the arbiter of our intuitions? *Behavioral and Brain Sciences* 4: 339–40.
- Kahneman, D. (2000). A psychological point of view: violations of rational rules as a diagnostic of mental processes (commentary on Stanovich and West). *Behavioral and Brain Sciences* 23: 681–3.

- Kahneman, D. (2003). Maps of bounded rationality: psychology for behavioral economics, *American Economic Review* 93(5): 1449–75.
- Kahneman, D. (2011). *Thinking Fast and Slow*. New York: Farrar, Strauss, Giroux.
- Kahneman, D. and Frederick, S. (2005). A model of heuristic judgment. In K. J. Holyoak and R. G. Morrison (eds.), *The Cambridge Handbook of Thinking and Reasoning*. Cambridge: Cambridge University Press, pp. 267–93.
- Kahneman, D. and Klein, G. (2009). Conditions for intuitive expertise: a failure to disagree. *American Psychologist* 64: 515–26.
- Kahneman, D. and Knetsch, J. (1993). Anchoring or shallow inferences: the effect of format, unpublished manuscript, University of California, Berkeley.
- Kahneman, D., Slovic, P., and Tversky, A. (eds.). (1982). *Judgement Under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.
- Kahneman, D. and Tversky, A. (1972). On prediction and judgment. *Oregon Research Institute Bulletin* 12(4).
- Kahneman, D. and Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica* 47(2): 263–91.
- Kahneman, D. and Tversky, A. (1982). *Judgement Under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press.
- Kahneman, D. and Tversky, A. (1984). Choices, values, and frames. *The American Psychologist* 39: 341–50.
- Kahneman, D. and Tversky, A. (1996). On the reality of cognitive illusions: a reply to Gigerenzer's critique. *Psychological Review* 103: 582–91.
- Kahneman, D. and Tversky, A. (2000). *Choices, Values, and Frames*. Cambridge: Cambridge University Press.
- Kamin, K. A. and Rachlinski, J. J. (1995). Ex post-ex ante: Determining liability in hindsight. *Law and Human Behavior* 19: 89–104.
- Kandel, E. and Pearson, N. D. (1995). Differential interpretation of public signals and trade in speculative markets. *Journal of Political Economy* 103: 831–72.
- Kaniel, R., Liu, S., Saar, G., and Titman, S. (2012). Individual investor trading and return patterns around earnings announcements. *Journal of Finance* 67: 639–80.
- Kaplan, S. N. and Weisbach, M. (1992). The success of acquisitions: evidence from divestitures. *Journal of Finance* 47: 107–38.
- Karandikar, R., Mookherjee, D., Debraj, R., and Vega-Redondo, F. (1998). Evolving aspirations and cooperation. *Journal of Economic Theory* 80(2): 292–331.
- Karlan, D., Morten, M., and Zinman, J. (2012). A personal touch: text messaging for loan repayment. Working Paper 17952, National Bureau of Economic Research, Cambridge, MA.
- Kaustia, M. (2004). Market-wide impact of the disposition effect: Evidence from IPO trading volume. *Journal of Financial Markets* 7: 207–35.
- Kaustia, M. (2010). Disposition effect. In H. K. Baker and J. R. Nofsinger (eds.), *Behavioral Finance: Investors, Corporations, and Markets*. Hoboken, NJ: John Wiley and Sons, pp. 171–89.
- Kelly, E. K. and Tetlock, P. C. (2013). How wise are Crowds? Insights from retail orders and stock returns. *Journal of Finance* 68: 1229–65.
- Kindleberger, C. P. and Aliber, R. Z. (2005). *Manias, Panics, and Crashes: A History of Financial Crises*. 5th edition. Hoboken, NJ: John Wiley and Sons, Inc.
- King, M. (2005). Monetary Policy: Practice Ahead of Theory, The Mais Lecture 2005: Speech by the Governor, *Bank of England Quarterly Bulletin*.
- King, M. A. and Dicks-Mireaux L. D. L. (1982). Asset holdings and the life cycle model. *Economic Journal* 92: 247–67.
- Kirchler, E. and Maciejovsky, B. (2002). Simultaneous over- and underconfidence: evidence from experimental asset markets. *Journal of Risk and Uncertainty* 25: 65–2585.
- Klayman, J. and Ha, Y. W. (1987). Confirmation, disconfirmation, and information in hypotheses testing. *Psychological Review* 94: 211–28.
- Klein, P. G. (2001). Were the acquisitive conglomerates inefficient? *RAND Journal of Economics* 32: 745–61.
- Koehler, J. J. (1996). The base rate fallacy reconsidered: descriptive, normative and methodological challenges. *Behavioral and Brain Sciences* 19: 1–53.
- Korniotis, G. M. and Kumar, A. (2009). Do older investors make better investment decisions? *Review of Economics and Statistics* 93: 244–65.
- Korniotis, G. M. and Kumar, A. (2013). Do portfolio distortions reflect superior information or

- psychological biases? *Journal of Financial and Quantitative Analysis* 48(1): 1–45.
- Korvost, M. and Damian, M. F. (2008). The differential influence of decades and units on multi-digit number comparison. *Quarterly Journal of Experimental Psychology* 61(8): 1250–64.
- Krishnamurty, A. (2002). The bond/old bond spread. *Journal of Financial Economics* 66: 463–506.
- Kumar, A. (2009). Hard-to-value stocks, behavioral biases, and informed trading. *Journal of Financial and Quantitative Analysis* 44: 1375–401.
- Kuran, T. and Sunstein, C. (1999). Availability cascades and risk regulation. *Stanford Law Review* 51(4): 683–768.
- Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica* 53: 1315–36.
- La Ferrara, E., Chong, A., and Duryea, S. (2012). Soap operas and fertility: evidence from Brazil. *American Economic Journal: Applied Economics* 4(4): 1–31.
- La Porta, R. (1996). Expectations and the cross-section of returns. *Journal of Finance* 51: 1715–1742.
- La Porta, R., Lakonishok, J., Shleifer, A., and Vishny, R. (1997). Good news for value stocks: further evidence on market efficiency. *Journal of Finance* 49: 1541–78.
- Lacetera, N., Pope, D. G., and Sydnor, H. (2012). Heuristic thinking and limited attention in the car market. *American Economic Review* 102(5): 2206–36.
- Lakonishok, J., Shleifer, A., and Vishny, R. W. (1992). The impact of institutional trading on stock prices. *Journal of Financial Economics* 32: 23–44.
- Lakonishok, J., Shleifer, A., and Vishny, R. (1994). Contrarian investment, extrapolation and risk. *Journal of Finance* 49: 1541–78.
- Lamont, O. and Thaler, R. (2003). Can the market add and subtract? Mispricing in tech stock carve-outs. *Journal of Political Economy* 111: 227–68.
- Landier, A. and Thesmar, D. (2009). Financial contracting with optimistic entrepreneurs. *Review of Financial Studies* 22: 117–150.
- Landsberger, M. (1966). Windfall income and consumption: comment. *American Economic Review* 56: 534–45.
- Langer, E. J. and Roth, J. (1975). Heads I win tails it's chance: the illusion of control as a function of the sequence of outcomes in a purely chance task. *Journal of Personality and Social Psychology* 32: 951–5.
- Larrick, R. P. (2004). Debiasing. In D. Koehler and N. Harvey (eds.), *Blackwell Handbook of Judgment and Decision Making*. Oxford: Blackwell, pp. 316–37.
- Larrick, R. P., Burson, K. A., and Soll, J. B. (2007). Social comparison and confidence: when thinking you're better than average predicts overconfidence (and when it does not). *Organizational Behavior and Human Decision Processes* 102: 76–94.
- Larwood, L. and Whittaker, W. (1977). Managerial myopia: self-serving biases in organizational planning. *Journal of Applied Psychology* 62: 194–8.
- Lazear, E. P. (2004). Balanced skills and Entrepreneurship. *American Economic Review* 94(2): 208–11.
- LeBoeuf, R. A. and Shafir, E. (2006). The long and short of it: physical anchoring effects. *Journal of Behavioral Decision Making* 19: 393–406.
- LeBoeuf, R. and Shafir, E. (2009). Anchoring on the here and now in time and distance judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 35: 81–93.
- Lee, C. M. C., Shleifer, A., and Thaler, R. H. (1991). Investor sentiment and the closed-end fund puzzle. *Journal of Finance* 46: 76–110.
- Lee, C. and Swaminathan, B. (2000). Price momentum and trading volume. *Journal of Finance* 55(5): 2017–69.
- Levi, I. (1983). Who commits the base rate fallacy? *Behavioral and Brain Sciences* 6: 502–6.
- Levin, L. (1998). Are assets fungible? Testing the behavioral theory of life-cycle savings. *Journal of Economic Behavior and Organization* 36: 59–83.
- Leviston, Z., Walker, I., and Morwinski, S. (2013). Your opinion on climate change might not be as common as you think. *Nature Climate Change* 3: 334–7.
- Lewis, K. (1999). Trying to explain home bias in equities and consumption. *Journal of Economic Literature* 37: 571–608.
- Lichtenstein, S. and Slovic, P. (1971). Reversals of preferences between bids and choices in gambling decisions. *Journal of Experimental Psychology* 89: 46–55.
- Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M., and Combs, B. (1978). Judged frequency of

- lethal events. *Journal of Experimental Psychology: Human Learning and Memory* 4: 551–78.
- Linnainmaa, J. T. (2010). Do limit orders alter inferences about investor performance and behavior? *Journal of Finance* 65: 1473–506.
- Lintner, J. (1956). Distribution of incomes of corporations among dividends, retained earnings and taxes. *American Economic Review* 46: 97–113.
- Locke, E. A. and Latham, G. P. (1990). *A Theory of Goal Setting and Task Performance*. Englewood Cliffs, NJ: Prentice-Hall.
- Loewenstein, G. F., Weber, E. U., Hsee, C. K., and Welch, E. S. (2001). Risk as feelings. *Psychological Bulletin* 127: 267–86.
- Lopes, L. L. (1981). Performing competently. *Behavioral and Brain Sciences* 4: 343–44.
- Lopes, L. L. (1991). The rhetoric of irrationality. *Theory and Psychology* 1: 65–82.
- Lord, C. G., Ross, L., and Lepper, M. R. (1979). Biased assimilation and attitude polarization: the effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology* 37(11): 2098–109.
- Lord, C. G. and Taylor, C. A. (2009). Biased assimilation: effects of assumptions and expectations on the interpretation of new evidence. *Social and Personality Psychology Compass* 3(5): 827–41.
- Loughran, T. and Ritter, J. R. (1995). The new issues puzzle. *Journal of Finance* 50: 23–52.
- Loughran, T., Ritter, J., and Rydqvist, K. (1994). Initial public offerings: international insights. *Pacific-Basin Finance Journal* 2: 165–99.
- Loughran, T. and Vijh, A. M. (1997). Do long-term shareholders benefit from corporate acquisitions? *Journal of Finance* 52: 1765–90.
- McKenzie, C. R. M., Liersch, M. J., and Yaniv, I. (2008). Overconfidence in interval estimation: What does expertise buy you? *Organizational Behavior and Human Decision Processes* 107: 179–91.
- Mackie, G. (1996). Ending footbinding and infibulation: a convention account. *American Sociological Review* 61(6): 999–1017.
- MacKinlay, A. C. (1995). Multifactor models do not explain deviations from the Capital Asset Pricing Models. *Journal of Financial Economics* 38: 3–28.
- Maćkowiak, B. and Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review* 99: 769–803.
- McLeod, P. and Dienes, Z. (1996). Do fielders know where to go to catch the ball, or only how to get there? *Journal of Experimental Psychology: Human Perception and Performance* 22: 531–43.
- McLeod, P., Read, N., and Dienes, Z., (2006). The generalized optic acceleration cancellation theory of catching. *Journal of Experimental Psychology* 32(1): 139–48.
- McNeil, B. J., Pauker, S. G., Sox, H. C., Jr., and Tversky, A. (1982). On the elicitation of preferences for alternative therapies. *New England Journal of Medicine* 306: 1259–62.
- Malkiel, B. (1990). *A Random Walk Down Wall Street*. New York: W. W. Norton.
- Malmendier, U. and Tate, G. (2005). CEO overconfidence and corporate investment. *Journal of Finance* 60: 2661–700.
- Malmendier, U. and Tate, G. (2008). Who makes acquisitions? CEO overconfidence and the market's reaction. *Journal of Financial Economics* 89: 20–43.
- Malmendier, U. and Tate, G. (2009). Superstar CEOs. *Quarterly Journal of Economics* 124(4): 1593–638.
- Malmendier, U., Tate, G., and Yan, J. (2011). Overconfidence and early-life experiences: the effect of managerial traits on corporate financial policies. *Journal of Finance* 66: 1687–733.
- Mani, A., Mullainathan, S., Shafir, E., and Zhao, J. (2013). Poverty impedes cognitive function. *Science* 341(6149): 976–80.
- March, J. G. and Shapira, Z. (1987). Managerial perspectives on risk and risk taking. *Management Science* 33: 1404–418.
- Margolis, H. (1982). *Selfishness, Altruism and Rationality*. Cambridge: Cambridge University Press.
- Margolis, H. (1987). *Patterns, Thinking, and Cognition*. Chicago: University of Chicago Press.
- Marsh, P. (1982). The choice between equity and debt: an empirical study. *Journal of Finance* 37: 121–44.
- Marsh, T. A. and Merton, R. C. (1986). Dividend variability and variance bound tests for the rationality of stock prices. *American Economic Review* 76(3): 483–98.
- Martignon, L. (2001). Comparisons between simple heuristics and optimization. In G. Gigerenzer and R. Selten (eds.), *Bounded Rationality: The Adaptive Toolbox (Proceedings of the Dahlem Workshop)*. Boston, MA: MIT Press, pp. 147–71.

- Matsusaka, J. G. (1993). Takeover motives during the conglomerate merger wave. *RAND Journal of Economics* 24: 357–79.
- Meehl, P. E. and Rose, A. (1955). Antecedent probability and the efficiency of psychometric signs, patterns or cutting scores. *Psychological Bulletin* 52: 194–216.
- Mehra, R. and Prescott, E. C. (1985). The equity premium: a puzzle. *Journal of Monetary Economics* 15: 145–61.
- Mengel, F. (2012). On the evolution of coarse categories. *Journal of Theoretical Biology* 307: 117–24.
- Menkhoff, L., Schmeling, M., and Schmidt, U. (2013). Overconfidence, experience, and professionalism: an experimental study. *Journal of Economic Behavior and Organization* 86: 92–101.
- Mercier, H. and Sperber, D. (2011). Why do humans reason? *Behavioral and Brain Sciences* 34: 57–111.
- Messer, W. S. and Griggs, R. A. (1993). Another look at Linda. *Bulletin of the Psychonomic Society* 31: 193–96.
- Metzger, M. (1984). Biases in betting: an application of laboratory findings. *Psychological Reports* 56: 883–8.
- Meulbroeck, L. K. (2001). The efficiency of equity-linked compensation: understanding the full cost of awarding executive stock options. *Financial Management* 30(2): 5–30.
- Michaely, R., Thaler, R., and Womack, K. (1995). Price reactions to dividend initiations and omissions. *Journal of Finance* 50: 573–608.
- Milgrom, P. (1981). Good news and bad news: representation theorems and applications. *Bell Journal of Economics* 12: 380–91.
- Milkman, K. L. and Beshears, J. (2009). Mental accounting and small windfalls: evidence from an online grocer. *Journal of Economic Behavior and Organization* 71: 384–94.
- Miller, J. B. and Sanjurjo, A. (2015). Surprised by the gambler's and hot hand fallacies? A truth in the law of small numbers. IGER working paper no. 552.
- Miller, M. H. and Rock, K. (1985). Dividend policy under asymmetric information. *Journal of Finance* 40: 1031–51.
- Mitchell, O. S. and Utkus, S. P. (2003). The role of company stock in defined contribution plans. In O. S. Mitchell and K. Smetters (eds.), *The Pension Challenge: Risk Transfers and Retirement Income Security*. Philadelphia: University of Pennsylvania Press.
- Mitchell, O. S. and Utkus, S. P. (2004). The role of company stock in defined contribution plans. In O. Mitchell and K. Smetters (eds.), *The Pension Challenge: Risk Transfers and Retirement Income Security*. Oxford: Oxford University Press, pp. 33–70.
- Mohlin, E. (2014). Optimal categorization. Forthcoming in *Journal of Economic Theory*.
- Morck, R., Shleifer, A., and Vishny, R. W. (1990). Do managerial objectives drive bad acquisitions? *Journal of Finance* 45: 31–48.
- Moro, R. (2009). On the nature of the conjunction fallacy. *Synthese* 171: 1–24.
- Moskowitz, T. and Wertheim, L. J. (2011). Scorecasting: *The Hidden Influences behind How Sports Are Played and Games Are Won*. New York: Crown Publishing.
- Mullainathan, S. (2002). A memory-based model of bounded rationality. *Quarterly Journal of Economics* 117(3): 735–74.
- Mullainathan, S. and Shafir, E. (2013). *Scarcity: Why Having Too Little Means So Much*. New York: Times Books.
- Mullen, B., Atkins, J. L., Champion, D. S., et al. (1985). The false consensus effect: a meta-analysis of 115 hypothesis tests. *Journal of Experimental Social Psychology* 21: 262–83.
- Munro, G. D., Ditto, P. H., Lockhart, L. K., Fagerlin, A., Gready, M., and Peterson, E. (2002). Biased assimilation of sociopolitical arguments: Evaluating the 1996 U.S. presidential debate. *Basic and Applied Social Psychology* 24: 15–26.
- Mussweiler, T. and Schneller, K. (2003). What goes up must come down—how charts influence decisions to buy and sell stocks. *The Journal of Behavioral Finance* 4(3): 121–30.
- Mussweiler, T. and Strack, F. (1999). Hypothesis-consistent testing and semantic priming in the anchoring paradigm: a selective accessibility model. *Journal of Experimental Social Psychology* 35: 136–64.
- Mussweiler, T. and Strack, F. (2000a). The use of category and exemplar knowledge in the solution of anchoring tasks. *Journal of Personality and Social Psychology* 78: 1038–52.
- Mussweiler, T. and Strack, F. (2000b). Numeric judgement under uncertainty: the role of knowledge in anchoring. *Journal of Experimental Social Psychology* 36: 495–518.
- Mussweiler, T., Strack, F., and Pfeiffer, T. (2000). Overcoming the inevitable anchoring effect: considering the opposite compensates for selec-

- tive accessibility. *Personality and Social Psychology Bulletin* 26: 1142–50.
- Myers, S. C. and Majluf, N. S. (1984). Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics* 13: 187–221.
- Navarro, A. D. and Fantino, E. (2009). The sunk time effect: an exploration. *Journal of Behavioral Decision Making* 22: 252–70.
- Nickerson, R. S. (1998). Confirmation bias: a ubiquitous phenomenon in many guises. *Review of General Psychology* 2(2): 175–220.
- Northcraft, G. B. and Neale, M. A. (1987). Experts, amateurs, and real estate: an anchoring-and-adjustment perspective on property pricing decisions. *Organizational Behaviour and Human Decision Processes* 39: 84–97.
- Nunn, N. (2008). The long-term effects of Africa's slave trades. *Quarterly Journal of Economics* 123(1): 139–76.
- Nunn, N. and Wantchekon, L. (2011). The slave trade and the origins of mistrust in Africa. *American Economic Review* 101(7): 3221–52.
- Odean, T. (1998). Are investors reluctant to realize their losses? *Journal of Finance* 53: 1775–98.
- Odean, T. (1999). Do investors trade too much? *American Economic Review* 89: 1279–98.
- Oechssler, J., Roider, A., and Schmitz, P. W. (2009). Cognitive abilities and behavioral biases. *Journal of Economic Behavior and Organization* 72: 147–52.
- Ofek, E. and M. Richardson (2003). Dot-com mania: market inefficiency in the internet sector. *Journal of Finance* 58: 1113–37.
- Okada, E. M. and Hoch, S. J. (2004). Spending time versus spending money. *Journal of Consumer Research* 31: 313–23.
- Oppenheimer, D. M. (2003). Not so fast! (and not so frugal!): rethinking the recognition heuristic. *Cognition* 90: B1–B9.
- Orbell, S. and Sheeran, P. (2000). Motivational and volitional processes in action initiation: a field study of the role of implementation intentions. *Journal of Applied Social Psychology* 30: 106–43.
- Pachur, T., Hertwig, R., and Steinmann, F. (2012). How do people judge risks: availability heuristic, affect heuristic, or both? *Journal of Experimental Psychology: Applied* 18: 314–30.
- Pagano, M., Fabio, P., and Zingales, L. (1998). Why do companies go public? An empirical analysis. *Journal of Finance* 53: 27–64.
- Palfrey, T. R. and Wang, S. W. (2012). Speculative overpricing in asset markets with information flows. *Econometrica* 80(5): 1937–76.
- Payne, J. W. and Bettman, J. R. (2001). Preferential choice and adaptive strategy use. In G. Gigerenzer and R. Selten (eds.), *Bounded Rationality: The Adaptive Toolbox*. MIT Press, pp. 123–45.
- Payne, J. W., Bettman, J. R., and Johnson, E. J. (1993). *The Adaptive Decision Maker*. New York: Cambridge University Press.
- Pennington, D. C. (1981). The British fireman's strike of 1977/78: an investigation of judgments in foresight and hindsight. *British Journal of Social Psychology* 20: 89–96.
- Peski, M. (2011). Prior symmetry, similarity-based reasoning, and endogenous categorization. *Journal of Economic Theory* 146(1): 111–40.
- Pinker, S. (1997). *How The Mind Works*. New York: Norton.
- Plous, S. (1989). Thinking the unthinkable: the effects of anchoring on likelihood estimates of nuclear war. *Journal of Applied Social Psychology* 19: 67–91.
- Plous, S., (1991). Biases in the assimilation of technological breakdowns: do accidents make us safer? *Journal of Applied Social Psychology* 21: 1058–82.
- Pogarsky, G. and Babcock, L. (2001). Damage caps, motivated anchoring and bargaining impasse. *Journal of Legal Studies* 30: 143–59.
- Polk, C. and Sapienza, P. (2009). The stock market and corporate investment: a test of catering theory. *Review of Financial Studies* 22(1): 187–217.
- Porter, M. E. (1987). From competitive advantage to corporate strategy. *Harvard Business Review* 65: 43–59.
- Poterba, J. M. (2003). Employer stock and 401(k) plans. *American Economic Review* 93: 398–404.
- Powdthavee, N. and Riyanto, Y. E. (2015). Would you pay for transparently useless advice? A test of boundaries of beliefs in the folly of predictions. Forthcoming in *Review of Economics and Statistics*.
- Prelec, D. and Loewenstein, G. (1998). The red and the black: mental accounting of savings and debt. *Marketing Science* 17(1): pp. 4–28.
- Quattrone, G. A. (1982). Overattribution and unit formation: when behavior engulfs the person. *Journal of Personality and Social Psychology* 42: 593–607.

- Quattrone, G. A., Lawrence, C. P., Finkel, S. E., and Andrus, D. C. (1981). Explorations in anchoring: the effects of prior range, anchor extremity, and suggestive hints. Mimeo, Stanford University, Stanford, CA.
- Rabin, M. (2002). Inference by believers in the law of small numbers. *Quarterly Journal of Economics* 117: 775–816.
- Rabin, M. and Schrag, J. (1999). First impressions matter: a model of confirmatory bias. *Quarterly Journal of Economics* 114(1): 37–82.
- Rabin, M. and Vayanos, D. (2010). The gambler's and hot-hand fallacies: theory and applications. *Review of Economic Studies* 77: 730–78.
- Rabin, M. and Weizsäcker, G. (2009). Narrow bracketing and dominated choices. *American Economic Review* 99(4): 1508–43.
- Rachlinski, J. J. (1998). A positive psychological theory of judging in hindsight. *University of Chicago Law Review* 65: 571–625.
- Rajagopal, P. and Rha, J.-Y. (2009). The mental accounting of time. *Journal of Economic Psychology* 30: 772–81.
- Rapoport, A. and Budescu, D. V. (1992). Generation of random binary series in strictly competitive games. *Journal of Experimental Psychology* 121: 352–64.
- Rapoport, A. and Budescu, D. V. (1997). Randomization in individual choice behavior. *Psychological Review* 104: 603–17.
- Rau, P. R. and Vermaelen, T. (1998). Glamour, value and the post-acquisition performance of acquiring firms. *Journal of Financial Economics* 49: 223–53.
- Ray, D. (2006). Aspirations, poverty, and economic change. In A. Banerjee, R. Bénabou, and D. Mookherjee (eds.), *Understanding Poverty*. Oxford: Oxford University Press.
- Read, D. and Loewenstein, G. (1995). Diversification bias: explaining the discrepancy in variety seeking between combined and separated choice. *Journal of Experimental Psychology: Applied* 1: 34–49.
- Read, D., Loewenstein, G., and Rabin, M. (1999). Choice bracketing. *Journal of Risk and Uncertainty* 19(1–3): 171–97.
- Redelmeier, D. A. and Shafir, E. (1995). Medical decision making in situations that offer multiple alternatives. *JAMA—Journal of the American Medical Association—U.S. Edition* 273(4): 302–5.
- Reinganum, M. R. (1983). The anomalous stock market behavior of small firms in January. *Journal of Financial Economics* 12: 89–104.
- Robbins, J. M. and Krueger, J. I. (2005). Social projection to ingroups and outgroups: a review and meta-analysis. *Personality and Social Psychology Review* 9(1): 32–47.
- Rockenbach, B. (2004). The behavioral relevance of mental accounting for the pricing of financial options. *Journal of Economic Behavior and Organization* 53: 513–27.
- Roese, N. J. and Maniar, S. D. (1997). Perceptions of purple: counterfactual and hindsight judgments at Northwestern Wildcats football games. *Personality and Social Psychology Bulletin* 23: 1245–53.
- Roll, R. (1986). The hubris hypothesis of corporate takeovers. *Journal of Business* 59: 197–216.
- Roll, R. (1988). R^2 . *Journal of Finance* 43: 541–66.
- Rosenthal, L. and Young, C. (1990). The seemingly anomalous price behavior of Royal Dutch/Shell and Unilever N.V./PLC. *Journal of Financial Economics* 26: 123–41.
- Ross L., Greene D., and House, P. (1977). The false consensus effect: an egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology* 13: 279–301.
- Ross S. (1977). The determination of financial structure: the incentive-signaling approach. *Bell Journal of Economics* 8: 23–40.
- Roth, B. and Voskort, A. (2014). Stereotypes and false consensus: how financial professionals predict risk preferences. *Journal of Economic Behavior and Organization* 107: 553–65.
- Rozeff, M. S. and Kinney, W. R. Jr. (1976). Capital market seasonality: the case of stock returns. *Journal of Financial Economics* 3: 379–402.
- Rubinstein A. (1998). *Modeling Bounded Rationality*. Cambridge, MA: MIT Press.
- Rubinstein, M. (2001). Rational markets: yes or no? The affirmative case. *Financial Analysts Journal* (May–June) 57(3): 15–29.
- Saks, M. J., Hollinger, L. A., Wissler, R. L., Evans, D. L., and Hart, A. J. (1997). Reducing variability in civil jury awards. *Law and Human Behavior* 21(3): 243–56.
- Samuelson, P. (1963). Risk and uncertainty: a fallacy of large numbers. *Scientia* 98: 108–13.
- Scarpetta, S., Hemmings, P., Tressel, T., and Jaejoon, W. (2002). The role of policy and

- institutions for productivity and firm dynamics: evidence from micro and industry data. OECD working paper No. 329.
- Scharfstein, D. S. and Stein, J. C. (1990). Herd behavior and investment. *American Economic Review* 80: 465–79.
- Schelling, T. (1984). Self command in practice, in policy and in a theory of rational choice. *American Economic Review* 74: 1–11.
- Schwarz, N. and Bohner, G. (2001). The construction of attitudes. In A. Tesser and N. Schwarz (eds.), *Blackwell Handbook of Social Psychology: Intraindividual Processes*. Oxford: Blackwell, pp. 436–57.
- Schwarz, N., Strack, F., Hilton, D., and Naderer, G. (1991). Base rates, representativeness, and the logic of conversation: the contextual relevance of irrelevant information. *Social Cognition* 9: 67–84.
- Schwarz, N., Strack, F., and Mai, H. P. (1991). Assimilation and contrast effects in part-whole question sequences: a conversational logic analysis. *Public Opinion Quarterly* 55: 3–23.
- Seasholes, M. S. and Wu, G. (2007). Predictable behavior, profits, and attention. *Journal of Empirical Finance* 15: 590–610.
- Selten R. (1998). Aspiration adaptation theory. *Journal of Mathematical Psychology* 42(2): 191–214.
- Selten R. (2001). What is bounded rationality? In G. Gigerenzer, and R. Selten (eds.), *Bounded Rationality: The Adaptive Toolbox. Dahlem Workshop Report* Cambridge, MA: MIT Press, pp. 1–12.
- Selten, R., Pittnauer, S. and Hohnisch, M. (2012). Dealing with dynamic decision problems when knowledge of the environment is limited: an approach based on goal systems. *Journal of Behavioral Decision Making* 25: 443–57.
- Seru, A., Shumway, T., and Stoffman, N. (2010). Learning by trading. *Review of Financial Studies* 23: 705–839.
- Shafir, E. and Thaler, R. H. (2006). Invest now, drink later, spend never: the mental accounting of delayed consumption. *Journal of Economic Psychology* 27: 694–712.
- Shah, A. K., Mullainathan, S., and Shafir, E. (2012). Some consequences of having too little. *Science* 338: 682–5.
- Shapira, Z. and Venezia, I. (2001). Patterns of behavior of professionally managed and independent investors. *Journal of Banking and Finance* 25: 1573–87.
- Shapiro, C. (1995). Aftermarkets and consumer welfare: making sense of Kodak. *Antitrust Law Journal* 63: 483–511.
- Shefrin, H. and Statman, M. (1984). Explaining investor preference for cash dividends. *Journal of Financial Economics* 13: 253–82.
- Shefrin, H. and Statman, M. (1985). The disposition to sell winners too early and ride losers too long. *Journal of Finance* 40: 777–90.
- Shefrin, H. and Statman, M. (2000). Behavioral portfolio theory. *Journal of Financial and Quantitative Analysis* 35(2): 127–51.
- Shefrin, H. and Thaler, R. H. (1988). The behavioral life-cycle hypothesis. *Economic Inquiry* 26: 609–43.
- Shefrin, H. M. and Thaler, R. (2004). Mental accounting, saving and self control. Chapter 14 in Camerer C. (ed.) *Advances in Behavioral Economics*. Princeton: Princeton University Press.
- Shiller, R. (1981). Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* 71: 421–36.
- Shiller, R. (2003). From efficient markets theory to behavioral finance. *Journal of Economic Perspectives* 17(1): 83–104.
- Shiller, R. (2015). *Irrational Exuberance*. 3rd edition. Princeton, NJ: Princeton University Press.
- Shleifer, A. (1986). Do demand curves for stocks slope down? *Journal of Finance* 41: 579–90.
- Shleifer, A. (2000). *Inefficient Markets: An Introduction to Behavioral Finance*. Oxford: Oxford University Press.
- Shleifer, A., Mullainathan, S., and Schwartzstein, J. (2008). Coarse thinking and persuasion. *Quarterly Journal of Economics* 123(2): 577–619.
- Shleifer, A. and Vishny, R. (1992). Liquidation values and debt capacity: a market equilibrium approach. *Journal of Finance* 47(4): 1343–66.
- Shleifer, A. and Vishny, R. (1997). The limits of arbitrage. *Journal of Finance* 52: 35–55.
- Shleifer, A. and Vishny, R. (2003). Stock market driven acquisitions. *Journal of Financial Economics* 70: 295–312.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics* 69: 99–118.
- Simon, H. A. (1956). Rational choice and the structure of environments. *Psychological Review* 63: 129–38.

- Simon, H. A. (1957). *Models of Man*. New York: Wiley.
- Simon, H. A. (1978). Rational decision-making in business organizations. Nobel Memorial Lecture, December 8, 1978.
- Simon, H. A. (2000). Bounded rationality in social science: today and tomorrow. *Mind and Society* 1(1): 25–39.
- Simonson, I. (1990). The effect of purchase quantity and timing on variety-seeking behavior. *Journal of Marketing Research* 27(2): 150–62.
- Sims, C. A. (1998). Stickiness. *Carnegie-Rochester Conference Series on Public Policy* 49: 317–56.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics* 50(3): 665–90.
- Sims, C. A. (2006). Rational inattention: beyond the linear-quadratic case. *American Economic Review Papers and Proceedings* 96(2): 158–63.
- Skinner, J. (1989). Housing wealth, and aggregate saving. *Regional Science and Urban Economics* 19: 305–24.
- Soman, S. A., Over, D., Slovak, L., and Stibel, J. M. (2003). Frequency illusions and other fallacies. *Organizational Behavior and Human Decision Processes* 91: 296–309.
- Slovic, P. (1995). The construction of preferences. *American Psychologist* 50: 364–71.
- Slovic, P., Finucane, M., Peters, E., and MacGregor, D. G. (2002). The affect heuristic. In T. Gilovich, D. Griffin, and D. Kahneman (eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment* New York: Cambridge University Press, pp. 397–98.
- Slovic, P., Fischhoff, B., and Lichtenstein, S. (1980). Facts and fears: understanding perceived risk. In R. Schwing and W. Albers (eds.), *Societal Risk Assessment: How Safe is Safe Enough*. Plenum Press, pp. 181–214.
- Slovic, P., MacGregor, D. G., Malmfors, T., and Purchase, I. F. H. (1999). *Influence of Affective Processes on Toxicologists' Judgments of Risk (Report No. 99-2)*. Eugene, OR: Decision Research.
- Slovic, P., Monahan, J., and MacGregor, D. G. (2000). Violence risk assessment and risk communication: the effects of using actual cases, providing instruction, and employing probability versus frequency formats. *Law and Human Behavior* 24: 271–96.
- Snyder, M. and Swann, W. B. (1978). Hypothesis-testing processes in social interaction. *Journal of Personality and Social Psychology* 36: 1202–12.
- Solan L., Rosenblatt T., and Osherson D. (2008). False consensus bias in contract interpretation. *Columbia Law Review* 108(5): 1272–300.
- Solnik, B. and Zuo, L. (2011). A global equilibrium asset pricing model with home preference. *Management Science* 58: 273–92.
- Soman, D. (2001). The mental accounting of sunk time costs: why time is not like money. *Journal of Behavioral Decision Making* 14: 169–85.
- Soman, D. and Cheema, A. (2001). The effect of windfall gains on the sunk cost effect. *Marketing Letters* 12(1): 49–60.
- Soman, D. and Cheema, A. (2004). When goals are counter-productive: the effects of violation of a behavioral goal on subsequent performance. *Journal of Consumer Research* 31(1): 52–62.
- Spiess, D. K. and Affleck-Graves, J. (1995). Underperformance in long-run stock returns following seasoned equity offerings. *Journal of Financial Economics* 38: 243–68.
- Stango, V. and Zinman, J. (2015). Limited and varying consumer attention: evidence from shocks to the salience of penalty fees. Forthcoming in *Review of Financial Studies*.
- Stanovich, K. E. (1999). *Who is Rational? Studies of Individual Differences in Reasoning*. Mahwah, NJ: Erlbaum.
- Stanovich, K. E. (2012). On the distinction between rationality and intelligence: implications for understanding individual differences in reasoning. In K. Holyoak and R. Morrison (eds.), *The Oxford Handbook of Thinking and Reasoning*. New York: Oxford University Press, pp. 343–65.
- Stanovich, K. E. and West, R. F. (2000). Individual differences in reasoning: implications for the rationality debate [Target article and commentaries]. *Behavioral and Brain Sciences* 23: 645–726.
- Stanovich, K. E. and West, R. (2002). Individual differences in reasoning: implications for the rationality debate? In T. Gilovich, D. Griffin, and D. Kahneman (eds.), *Heuristics and Biases: The Psychology of Intuitive Judgment*. New York: Cambridge University Press, pp. 421–40.
- Stanovich, K. E. and West, R. F. (2007). Natural Myside bias is independent of cognitive ability. *Thinking and Reasoning* 13(3): 225–47.
- Stanovich, K. E. and West, R. F. (2008a). On the failure of cognitive ability to predict Myside and

- one-sided thinking biases. *Thinking and Reasoning* 14(2): 129–67.
- Stanovich, K. E. and West, R. F. (2008b). On the relative independence of thinking biases and cognitive ability. *Journal of Personality and Social Psychology* 94(4): 672–95.
- Stein, J. C. (1996). Rational capital budgeting in an irrational world. *Journal of Business* 69: 429–55.
- Stigler, G. J. (1961). The economics of information. *Journal of Political Economy* 69: 213–25.
- Stigler, G. J. (1987). *The Theory of Price*, 4th edition. New York, Macmillan Publishing.
- Stolarz-Fantino, S., Fantino, E., Zizzo, D. J., and Wen, J. (2003). The conjunction fallacy: New evidence for robustness. *American Journal of Psychology* 116: 15–34.
- Strack, F., Martin, L. L., and Schwarz, N. (1988). Priming and communication: the social determinants of information use in judgments of life-satisfaction. *European Journal of Social Psychology* 18: 429–42.
- Strack, F. and Mussweiler, T. (1997). Explaining the enigmatic anchoring effect: mechanisms of selective accessibility. *Journal of Personality and Social Psychology* 73: 437–46.
- Subrahmanyam, A. (2008). Behavioral finance: a review and synthesis. *European Financial Management* 14: 12–29.
- Summers, L. and Carroll, C. (1987). Why is the U. S. Saving Rate So Low? *Brookings Papers on Economic Activity* 2: 607–35.
- Sunstein, C. R. (1996). Social norms and social roles. *Columbia Law Review* 96(4): 903–68.
- Svenson, O. (1981). Are we all less risky and more skillful than our fellow drivers? *Acta Psychologica* 47: 143–48.
- Taleb, N. N. (2007). *The Black Swan: The Impact of the Highly Improbable*. Random House.
- Tentori, K., Bonini, N., and Osherson, D. (2004). The conjunction fallacy: a misunderstanding about conjunction? *Cognitive Science* 28: 467–77.
- Tentori, K. and Crupi, V. (2012). On the conjunction fallacy and the meaning of and, yet again: A reply to Hertwig, Benz, and Krauss (2008). *Cognition* 122: 123–34.
- Tentori, K., Crupi, V., and Russo, S. (2013). On the determinants of the conjunction fallacy: probability vs. inductive confirmation. *Journal of Experimental Psychology: General* 142: 235–55.
- Terrell, D. (1994). A test of the gambler's fallacy: evidence from pari-mutuel games. *Journal of Risk and Uncertainty* 8: 309–17.
- Terrell, D. (1998). Biases in assessments of probabilities: new evidence from greyhound races. *Journal of Risk and Uncertainty* 17: 151–66.
- Terrell, D. and Farmer, A. (1996). Optimal betting and efficiency in parimutuel betting markets with information costs. *Economic Journal* 106: 846–68.
- Tetlock, P. E. (2002). Cognitive biases in path-dependent systems: theory driven reasoning about plausible pasts and probable futures in world politics. In T. Gilovich, D. W. Griffin, and D. Kahneman (eds.), *Inferences, Heuristics and Biases: New Directions in Judgment Under Uncertainty*. New York: Cambridge University Press.
- Tetlock, P. E. (2006). *Expert Political Judgment: How Good Is It? How Can We Know?* Princeton, NJ: Princeton University Press.
- Tetlock, P. E. (2010). Second thoughts on expert political judgment. Reply to symposium on Expert political judgment: How good is it? How can we know? *Critical Review* 22(4): 467–88.
- Thaler, R. H. (1980). Toward a positive theory of consumer choice. *Journal of Economic Behavior and Organization* 1: 39–60.
- Thaler, R. (1985). Mental accounting and consumer choice. *Marketing Science* 4: 199–214. Republished (2008) in *Marketing Science*.
- Thaler, R. H. (1989). Inter-industry wage differentials. *Journal of Economic Perspectives* 3: 181–93.
- Thaler, R. H. (1999). Mental accounting matters. *Journal of Behavioral Decision Making* 12: 183–206.
- Thaler, R. H. (2008). Mental accounting and consumer choice. *Marketing Science* 27(1): 15–25.
- Thaler, R. H. (2015). *Misbehaving: The Making of Behavioral Economics*. New York: W. W. Norton.
- Thaler, R. H. and Benartzi, S. (2007). The behavioral economics of retirement savings behavior. AARP Public Policy Institute Working Paper 2007–02.
- Thaler, R. and Johnson, E. (1990). Gambling with the house money and trying to break even: the effects of prior outcomes on risky choice. *Management Science* 36: 643–60.
- Thaler, R. H. and Shefrin, H. M. (1981). An economic-theory of self-control. *Journal of Political Economy* 89(2) 392–406.

- Thaler, Richard H. and Sunstein, C. R. (2008). *Nudge: Improving Decisions about Health, Wealth, and Happiness*. New Haven, CT: Yale University Press.
- Todd, P. M. (2001). Fast and frugal heuristics for environmentally bounded minds. In G. Gigerenzer and R. Selten (eds.), *Bounded Rationality: The Adaptive Toolbox*. Cambridge, MA: MIT Press.
- Todd, P. M. and Gigerenzer, G. (2000). Précis of simple heuristics that make us smart. *Behavioral and Brain Sciences* 23: 727–41.
- Touron, D. R. and Hertzog, C. (2004). Distinguishing age differences in knowledge, strategy use, and confidence during strategic skill acquisition. *Psychology and Aging* 19: 452–66.
- Trope, Y. and Bassok, M. (1982). Confirmatory and diagnosing strategies in social information gathering. *Journal of Personality and Social Psychology* 43: 22–34.
- Tversky, A. (1972). Elimination by aspects: a theory of choice. *Psychological Review* 79: 281–99.
- Tversky, A. and Gilovich T. (1989a). The cold facts about the “hot hand” in basketball. *Chance* 2: 16–21.
- Tversky, A. and Gilovich T. (1989b). The hot hand: statistical reality or cognitive illusion? *Chance* 2: 31–4.
- Tversky, A. and Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin* 76: 105–10.
- Tversky, A. and Kahneman, D. (1973). Availability: a heuristic for judging frequency and probability. *Cognitive Psychology* 5(2): 207–32.
- Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: heuristics and biases. *Science* 185: 1124–31.
- Tversky, A. and Kahneman, D. (1980). Causal schemas in judgments under uncertainty. In M. Fishbein (ed.), *Progress in Social Psychology*, Hillsdale, NJ: Erlbaum, pp. 49–72.
- Tversky, A. and Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science* 211(4481): 453–8.
- Tversky, A. and Kahneman, D. (1983). Extensional vs. intuitive reasoning: the conjunction fallacy in probability judgment. *Psychological Review* 90: 293–315.
- Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4): 297–323.
- Wansink, B., Kent, R. J., and Hoch, S. J. (1998). An anchoring and adjustment model of purchase quantity decisions. *Journal of Marketing Research* 35(1): 71–81.
- Wason, P. C. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology* 20: 273–81.
- Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology* 39(5): 806–20.
- Wertenbroch, K. (1998). Consumption self-control by rationing purchase quantities of virtue and vice. *Marketing Science* 4: 317–37.
- Wetherick, N. E. (1993). Human rationality. In: K. Manktelow and D. Over (eds.), *Rationality: Psychological and Philosophical Perspectives*. Routledge.
- Wiederholt, M. (2010). Rational inattention. In S. M. Durlauf and L. E. Blume (eds.), *The New Palgrave Dictionary of Economics*. Online edition. Palgrave Macmillan.
- Wilson, T. D., Houston, C., Etling, K. M., and Brekke, N. (1996). A new look at anchoring effects: Basic anchoring and its antecedents. *Journal of Experimental Psychology: General* 4: 387–402.
- Woodford, M. (2012). Inattentive valuation and reference-dependent choice. Mimeo, Columbia University.
- World Development Report (2015). *Mind, Society, and Behaviour*. Washington DC: The World Bank.
- Wright, W. F. and Anderson, U. (1989). Effects of situation familiarity and financial incentives on use of the anchoring and adjustment heuristic for probability assessment. *Organizational Behavior and Human Decision Processes* 44: 68–82.
- Wurgler, J. and Zhuravskaya, K. (2002). Does arbitrage flatten demand curves for stocks? *Journal of Business* 75: 583–608.
- Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist* 35(2): 151–75.
- Zaltman, G. (1997). Rethinking market research: putting people back in. *Journal of Marketing Research* 34: 424–37.

PART 8

Behavioral Welfare Economics

INTRODUCTION TO PART 8

In neoclassical economics, rational individuals are considered to be the best judges of their own welfare, and their choices reflect informed judgment of their welfare. One can then use revealed preference arguments to derive an individual's preferences based on observed choices. The findings from behavioral economics show that individuals suffer from a range of biases and misperceptions. Here, and throughout the book, the terms *biases* and *misperceptions* only make sense, relative to the rational benchmark in neoclassical economics. There should be no presumption that, in any absolute sense, the actual behavior of humans should either be termed as a bias or a misperception. In the presence of biases and misperceptions, individuals can no longer be considered to be making informed judgments of their welfare. This calls for a modification of the methods of neoclassical welfare economics, because classical revealed preference analysis does not apply. These issues form the subject matter of Part 8.

Libertarianism and *paternalism* are loaded words. Economists have typically used these terms to mean, respectively, respect for individual choices, and interference with individual choices in a direction that improves welfare. Relative to behavioral interventions, neoclassical forms of paternalism are sometimes referred to as *heavy-handed paternalism*. Examples include banning of certain substances, such as narcotics; protection of the vulnerable against usury laws; health and safety regulations; warnings on cigarette packs; and regulation of the “approval process” for drugs.

The main approach advocated in behavioral economics is *soft paternalism*. This is an umbrella term that nests several kinds of interventions, such as *libertarian paternalism*, *asymmetric paternalism*, and *light paternalism*. The main thrust of these approaches is to design interventions, or *nudges*, that improve the welfare of boundedly rational individuals, but impose minimal costs on fully rational individuals. A leading example is *default options* that address a range of potential biases of boundedly rational individuals; these biases include, but are not limited to, present-biased preferences, temptations, and limited attention. Appropriate defaults improve the decisions of boundedly rational individuals, but impose small, if any, costs on rational individuals.

We also consider a range of other important topics in this part that we briefly allude to now.

We examine the case for *mindless economics* versus the case for *mindful economics*. This debate is about the very scope of economics and addresses the question: Should economics be limited to choice data alone, or are non-choice data admissible too? The reader will not be surprised to learn that this book strongly advocates the use of choice and non-choice data, such as survey evidence, experimental evidence, and neuroeconomic and biometric evidence. The extreme view that economics ought to be restricted to choice data and revealed preference arguments only is not subscribed to in this book.

We consider formal models of choice-based welfare. While these directions are promising, they also show up the limitations of the approach, even when the choice-based approaches are stretched to their limits. Choice-based models must address the issue of choices that depart from those expected under the rational benchmark. In a leading model, one deals with this issue by trimming-away the anomalous choices. However, such trimming-away necessitates the use of either non-choice data, or the invocation of a welfare criteria for trimming the choices, which is what one is trying to construct in the first place. We also consider the extension of this framework to issues of limited attention.

The limitations of the choice-based approach, despite commendable progress in this direction, have prompted some to focus on preferences rather than choice. We also consider the development of alternative frameworks such as the *contractarian approach* in which interventions take the form of directly providing information to individuals about their behavioral biases.

In the last section, we give readers a flavor of the rapidly developing literature in behavioral public economics. We discuss how we might formulate traditional public finance concepts such as tax incidence and the excess burden of taxes in the presence of behavioral factors, such as limited attention. We also briefly comment on reduced-form approaches to behavioral welfare economics that nest a range of behavioral explanations.

None of the frameworks considered in Part 8 will persuade all readers of the book. This, despite scholarship of a high quality, flags up the inherent and deep difficulty of dealing with the underlying issues that straddle neoclassical economics, behavioral economics, philosophy, and ethics.

CHAPTER 22

Behavioral Welfare Economics

22.1 Introduction

Neoclassical welfare economics follows the liberal tradition in economics. The primitives in the model are unobserved individual *preferences* over a set of outcomes. Observed *choices* reflect underlying preferences. In neoclassical theory, policymakers and society respect observed choices because they reflect a *considered judgment* of one's welfare. For this reason, these preferences are sometimes known as *normative preferences*. The neoclassical approach is encapsulated in *revealed preference analysis*. Individuals make choices based on normative preferences and the aim is to see if the observed choices can be rationalized by some underlying utility function (Samuelson, 1948; Houthakker, 1950; Afriat, 1967).

Neoclassical economics allows for the regulation of the behavior of individuals in some cases. In the behavioral economics literature, neoclassical paternalism is sometimes referred to as *heavy-handed paternalism* (Loewenstein and Haisley, 2008). Not everyone will agree on this terminology, but I use it simply as a convenient way of describing the neoclassical benchmark. Its main motivation is that individuals may impose externalities on others, or even on themselves, that they can be made to internalize by appropriately designed regulation or *Pigouvian taxes*.¹ Examples include compulsory helmets for drivers of two-wheelers, seat belts in cars, mandatory third party vehicle insurance, and certificates of roadworthiness of vehicles. These regulations are often justified by a lack of adequate awareness of the risks on the part of individuals, or by the externalities imposed on others from careless behavior. In other cases, individuals might not provide inadequately towards a public good. In this case, the neoclassical solution is either government provision, financed by taxes, or Lindahl taxes. There are also regulations on firms in markets that depart from the perfectly competitive ideal, as encapsulated in, say, antitrust laws and environmental regulations.

The evidence from behavioral economics contests the neoclassical view that individuals always make considered judgments of their welfare. A large body of evidence suggests that there is a wedge between revealed and normative preferences. The wedge may be exacerbated by complexity, long horizons, and limited personal experience (Beshears et al., 2008). The relevant evidence is discussed throughout the book so here we give only a brief outline and suppress the references as far as possible.

Individuals exhibit a range of judgment heuristics and biases that are reflected in actual choices; see Chapter 19 for details. Examples include anchoring, representativeness, law of small numbers,

¹ The interested reader can pursue these arguments in any standard book on public economics (Atkinson and Stiglitz, 1980; Stiglitz, 2000) or in microeconomic theory (Mas-Collel et al., 1995).

base rate underweighting, affect heuristic, hindsight bias, conjunction fallacy, confirmation bias, ignoring regression to the mean, and false consensus effect. In other cases, individuals might not have well-defined preferences and choices may be unduly influenced by informationally equivalent frames and contexts. Individuals might also make important decisions, such as buying consumer durables, or marriage/divorce, in an emotional hot state of the mind that they regret later, once they are in a cold state. Individuals can sometimes exhibit self-control problems on account of present-biased preferences that lead to suboptimal outcomes such as inadequate savings for retirement, obesity, procrastination, and drug use. Furthermore, many individuals report that they would prefer to save more, or prefer the median portfolio to their own, and others pay to kick a bad habit, indicating that they recognize their choices to be suboptimal (Camerer et al., 2003; Beshears et al., 2008).

Data from the US shows that many individuals do not utilize the full limits of their 401(k) pension plans. There is also inadequate diversification of portfolios. For instance, individuals often use the $1/n$ heuristic to allocate their portfolio, where n is the number of options available. This skews their pension portfolios towards equities, particularly company equity, which can have adverse consequences, as it did for Enron employees. Individuals often overutilize payday loans. These are loans that are repaid when the next paycheck arrives; the short-term interest rates on these loans are excessively steep. Individuals often suffer from limited attention. For instance, they react more to sales taxes that are more salient, and there are discontinuous drops in the prices of used cars at 10K thresholds.

Given this evidence, what welfare or normative significance can one attach to individual choices, and how can we modify neoclassical methods to take account of observed choices? This forms the main subject matter of this part of the book.

The main thrust in *behavioral welfare economics* has been on reconstructing the informed underlying judgment of individuals. This has given rise to a range of suggestions under the umbrella term of *soft paternalism* that we describe in more detail below. Camerer et al. (2003) view the relaxation of normative preferences in welfare economics as a natural progression of the relaxation of other restrictive assumptions in economics, such as perfect competition, perfect information, certainty, etc.

We begin, in Section 22.2, by outlining a simple framework of analysis (Dalton and Ghosal, 2011, 2016; Dalton et al., 2016). Individuals derive utility from actions and psychological states. A feedback function maps actions to psychological states. Classical decision makers, who have no behavioral biases or misperceptions, take account of the feedback. However, behavioral decision makers ignore the feedback, hence, their revealed choices differ from normative choices and their welfare level is also lower. We explore a range of policy interventions that can raise the welfare levels of behavioral decision makers to the same level as neoclassical decision makers. These include *direct paternalism* (directly choosing an action for the decision maker), *indirect paternalism* (using Pigouvian taxes/subsidies), and *soft paternalism* (appropriate nudges that impose minimal costs on rational individuals but confer maximal benefits to boundedly rational individuals).

In contrast to *heavy-handed paternalism* in neoclassical economics, the generic form of paternalism in behavioral economics is known as *soft paternalism* or *light paternalism*. In Section 22.3, we consider a range of proposals that have been made under this umbrella term. These include *libertarian paternalism* (Thaler and Sunstein, 2003, 2009; Sunstein and Thaler, 2003), *asymmetric paternalism* (Camerer et al., 2003), and *light paternalism* (Loewenstein and Ubel, 2008).

At the heart of most proposals under the ambit of soft paternalism is the need to distort individual choices as little as possible, and use as little coercion as possible. The idea is that the mistakes,

or misperceptions, of boundedly rational individuals will be *nudged* in the direction of a beneficial correction by a *benevolent planner* or a *choice architect*. At the same time, the minimum intrusion associated with soft paternalism is designed to impose little or no cost on the actions of the fully rational individuals.

Let us consider an example. A *choice architect* or *benevolent planner* might arrange food items in a cafeteria in a certain order. Boundedly rational consumers who pay inordinate attention to items placed at eye level, could benefit from the placement of healthy options at eye level; this is an example of a *nudge*. Neoclassically rational consumers are not influenced by the arrangement of items, so little cost is imposed on them. The intervention of the choice architect is a form of paternalism. However, ultimately it is the consumer who exercises choice; in this sense, the proposal respects *consumer sovereignty*, which is a form of liberalism. Hence, one is led to the term libertarian-paternalism. We consider the effects of several concrete policy proposals in this tradition; these include the role of default options, deadlines, and cooling-off periods.

These proposals are by no means perfect and dissenting views have been given (Sugden, 2008, 2013). Under soft paternalism, the choice architect proposes welfare improving nudges to individuals, based on designing choices, *as judged best by the individuals* under perfect information, unlimited cognitive abilities, and complete self-control. While the preference for median portfolios rather than their own, and checking into rehab to kick a habit, provide tantalizing evidence of individuals judging their choices to be poor, it might not always be clear what individuals judge to be in their best interests.

To operationalize light paternalism, one needs to specify what a mistake/misperception is. Furthermore, in the light of these mistakes/misperceptions, how should we judge the outcomes of light-paternalistic policies? The latter question requires us to have some behavioral welfare criteria. The reason that a traditional welfare criterion is not suitable is that actual choices might not be indicative of an individual's best interests. Thus, we cannot directly use existing measures of welfare, such as equivalent variation and/or compensating variation, to determine the direction in which individual welfare is improving. As Loewenstein and Haisley state (p. 12): "Clearly, it doesn't make sense to assess whether someone is committing an error using a measure that is premised on the assumption that people don't commit errors."

In the typical case considered in economics, consumers have greater information about their own type, for instance, their demand for a good or their health status, relative to the firm. In Section 22.4, we consider the opposite possibility. Namely, that firms are better informed about the consumer's usage of their product, say, energy usage, perhaps because consumers exhibit limited attention (Kamenica et al., 2011). In particular we focus on the *RECAP regulation* suggested by Thaler and Sunstein (2009). Among other things, it requires firms to improve the transmission of information to consumers about their own usage of a product.

As an example, energy firms may be required to install home power meters for their clients, allowing consumers to monitor their own usage. We show that, under exogenous prices, this regulation can reduce the expenditure of consumers, and increase their welfare. However, this regulation may also induce consumers to switch firms. If endogenous prices are allowed, it is not possible to say whether prices will increase or decrease, although intuition suggests that the switching decisions will put downward pressures on prices.

Section 22.5 considers the appropriate scope of economics. An extreme position holds that the only admissible data in economics is *choice data*, and the scope of economics is solely to infer preferences from observed choices (Gul and Pesendorfer, 2008). In this view, non-choice data, such as survey data, self-reported measures of well-being, neuroeconomic data, data based on pupil dilation or skin conductance, or pulse rates, and data generated by mouselab software, are

considered outside the scope of economics. Indeed, in the most extreme view, even experimental data may be viewed as non-choice data, or at least, minimally, with great suspicion. In this view, economists should restrict themselves to observing pairs of quantities and prices, remaining agnostic about the reasons for choice. A second component of the Gul–Pesendorfer position is that the neoclassical framework is flexible enough to account for unexpected observations.

Most behavioral economists do not subscribe to the Gul–Pesendorfer position; certainly this book does not. The edited volume by Caplin and Schotter (2008) gives a range of views on this matter and we describe some selected views. Definitions are inherently subjective, and one need not subscribe to the one suggested by Gul–Pesendorfer. Furthermore, the record of non-choice data in estimating economic parameters of interest, and in illuminating diverse economic phenomena, has been impressive (Camerer, 2008). Misusing the flexibility of the neoclassical model to take account of unexpected observations by introducing ad-hoc auxiliary assumptions, produces *wacky models* (Köszegi and Rabin, 2008a; Camerer, 2008). It is possible to show that by taking account of behavioral misperceptions, one may be able to recover underlying preferences from choices (Köszegi and Rabin, 2008a,b). Experiments have been hugely influential in exploring parameter spaces that may not be observed in actual choice data, and in allowing for treatments that are difficult to observe outside the experimental lab; this has allowed for stringent testing of economic theory (Schotter, 2008). The introductory chapter considers many of these issues in greater detail.

In Section 22.6, we consider *choice-based welfare criteria* with a focus on the Bernheim–Rangel proposal (Bernheim and Rangel, 2009; Salant and Rubinstein, 2008; Rubinstein and Salant, 2012). One begins with a domain of sets of outcomes, and *ancillary conditions* or *frames*, over which individuals have preferences that need not be complete or transitive. If the domain is trimmed-down to a “welfare relevant domain”, then one can define preference relations that are complete but possibly intransitive. However, intransitivity does not preclude a welfare analysis. The preference relations on the trimmed-down domain allow one to define criteria for welfare improvement. Under the assumption that each ancillary condition corresponds to a different self of the same individual, and frame-independent choices, one can relate the preference relations to the neoclassical Pareto criterion. Rubinstein and Salant (2012) depart from the Bernheim–Rangel framework by proposing a welfare criterion based on preference orderings that are frame dependent. They explain inconsistent choices by invoking switches in preference orderings between frames.

The developments in choice-based measures of welfare under behavioral biases are promising, but they also illustrate the great difficulty in making progress in this direction. For instance, in the Bernheim–Rangel model, it is not clear who will trim-down the original domain to a welfare relevant domain, or indeed how it could be done. Presumably, in order to do so, one will need to use some non-choice data or invoke some other welfare criteria, which runs into the danger of an infinite regress of welfare criteria. This has led some to make calls for advocating welfare criteria based on preferences rather than choice (Loewenstein and Haisley, 2008).

Section 22.7 takes account of *limited attention* in the choice-based approach to revealed preference (Masatlioglu et al., 2012). For instance, consumers might simply choose between the two cheapest suppliers, ignoring all others; or in choosing a car, potential buyers might first focus on a car with the best fuel efficiency, and then choose the cheapest car in this category, ignoring all others. In the Bernheim–Rangel approach, an alternative x is unambiguously preferred to y , if y is never chosen when x is available. However, under limited attention, in some choice situations, the individual might choose y simply because he does not pay attention to x . We jointly consider the twin phenomena of *revealed preference* and *revealed inattention* and propose an appropriate modification of the weak axiom of revealed preference under these conditions. We close by

briefly considering the relation with related approaches (Cherepanov et al., 2013; Manzini and Mariotti, 2007).

Section 22.8 outlines the *contractarian approach* (Sugden, 2013) that draws on an older literature in normative economics (Buchanan, 1968, 1975, 1986). This approach criticizes the paternalism advocated in neoclassical economics, and under soft paternalism, because proposals addressed to a social planner lack a valid *addressee*. The contractarian approach advocates addressing recommendations directly to the individuals who exhibit the relevant biases; for instance, through a public information campaign that disseminates information about the harmful effects of smoking.

The recommendations made by a contractarian appeal to the self-interest of individuals in improving their own welfare, rather than to the general or larger interests of a group of people. Nudges are not allowed in this approach if individuals choose not to be nudged. However, providing information directly to individuals under the contractarian approach may also be viewed as a nudge. Furthermore, it is not clear who will decide how the relevant information will be framed; this is particularly important if framing effects are important, or unclear. Thus, the contractarian approach also suffers from some of the same drawbacks that it finds wanting in other theories.

Section 22.9 continues our discussion of *behavioral public finance* from Part 7. It considers welfare issues and public policy in the presence of behavioral factors, particularly, limited attention. Section 22.9.1 considers the problem of computing the relative *tax incidence* on producers and consumers of an increase in the sales tax under limited attention (Chetty et al., 2009). Under limited attention, producers bear a lower burden of the tax increase and the tax incidence falls relatively more on the consumers. While the final outcome is similar to the case of a fall in the elasticity of demand, tax elasticity and limited attention have different effects that we explain.

Section 22.9.2 derives the formulae for the *excess burden* of taxation under limited attention (Chetty et al., 2009). We show that the excess burden increases with the square of a parameter, θ , such that lower values of θ denote higher inattention. When there are no income effects, then limited attention may reduce the excess burden of taxation, which is welfare improving. Finally, in Section 22.9.3, we consider some reduced-form approaches to behavioral welfare analysis. We mainly focus on models that make a distinction between *decision utility* and *experienced utility* (Mullainathan et al., 2012). We also briefly discuss other promising approaches that extend classical optimal tax theory by incorporating behavioral factors such as limited attention (Farhi and Gabaix, 2015).

22.2 Fixing basic ideas

In this section, we introduce some of the issues in behavioral welfare economics, using a simple motivating framework.

22.2.1 The basic framework

Consider the simple theoretical framework in Dalton and Ghosal (2011). Suppose that an individual chooses from the set of actions $A = [\underline{a}, \bar{a}] \subset \mathbb{R}$ and experiences the set of possible psychological states $\Theta = [\underline{\theta}, \bar{\theta}] \subset \mathbb{R}$. Examples of psychological states include moods, aspirations, temptations, reference points, motivations, and degree of self-confidence.²

² The framework is easily extended to the case where A and Θ are subsets of a finite dimensional Euclidean space.

Psychological states are modeled using the theory of *reciprocal determinism* (Bandura, 1977). This implies that psychological states influence the actions of the individual, but are also, in turn, influenced by actions through a feedback function. For instance, more confident/adventurous individuals are more likely to try out a bungee jump. Once they have made the jump (action), their confidence and sense of adventure might be further increased (feedback).

Formally, we assume that the individual has well-defined preferences and the utility function is given by $u : A \times \Theta \rightarrow \mathbb{R}$; u is twice continuously differentiable in both arguments up to order 2, and has continuous cross partial derivatives. The *feedback function* that maps actions to psychological states is given by $f : A \rightarrow \Theta$; f is twice continuously differentiable. We assume that

$$\text{For all } a \in A, \exists \theta \in [\underline{\theta}, \bar{\theta}] : \theta = f(a). \quad (22.1)$$

Definition 22.1 A consistent decision state is a pair (a, θ) , where $a \in A$ and $\theta \in \Theta$, such that $\theta = f(a)$.

Thus, for $(a, \theta) \in A \times \Theta$, the decision maker's utility is $u(a, \theta)$ and for a consistent decision state, $(a, \theta) = (a, f(a))$, utility is $u(a, f(a))$.

Formal models of behavioral welfare economics typically consider two kinds of decision makers. *Classical* decision makers conform to the assumptions of neoclassical economics. *Behavioral* decision makers, in most behavioral welfare models, typically exhibit some behavioral departures relative to classical decision makers. The main distinction between the two decision makers in the model, in this section, is the degree of awareness of the feedback function.

Classical decision makers take account of f , and solve the following *classical decision problem* (CDP).

$$a^c \in \arg \max_{a \in A} v(a) \equiv u(a, f(a)). \quad (22.2)$$

v is a continuous function on a compact interval, hence, a maximum exists, although there is no guarantee that it is unique.

Behavioral decision makers do not take account of the feedback function, f . The solution to a *behavioral decision problem* (BDP) is a pair $(a^b, \theta^b) \in A \times \Theta$ that solves two conditions.³

1. For a given psychological state, θ , the chosen action is optimal, i.e.,

$$a^b \in \arg \max_{a \in A} u(a, \theta). \quad (22.3)$$

u is a continuous function on a compact interval, so a maximum exists.

2. The psychological state is influenced by the feedback rule in equilibrium, i.e.,

$$\theta^b = f(a^b), \quad (22.4)$$

although this is not recognized in the BDP.

³ For an axiomatic characterization of the choice theoretic foundations of the CDP and the BDP and their welfare implications, see Dalton and Ghosal (2016).

For pedagogical reasons, we assume that the global maximum lies in the interior of $[a, \bar{a}]$ for both kinds of decision makers. Our assumptions do not guarantee the existence of a unique equilibrium because neither u nor v are strictly concave. If there are several maxima, then the global maximum is the upper bound of this set. For pedagogical reasons, we shall only be interested in the global maximum below. We use the notation u_1 and u_2 , respectively, for partial derivatives of $u(a, f(a))$ with respect to the first and the second arguments. First and second order partial derivatives of f are written, respectively, as f' and f'' .

The first order condition for an interior optimum to the CDP, evaluated at the optimal values, is

$$u_1(a^c, f(a^c)) + u_2(a^c, f(a^c))f'(a^c) = 0. \quad (22.5)$$

The first order condition and the equilibrium condition for an interior optimum to a BDP are given by

$$u_1(a^b, f(a^b)) = 0 \text{ and } \theta^b = f(a^b). \quad (22.6)$$

In the first order condition, the behavioral decision maker does not take account of the feedback rule between actions and psychological states. The welfare criteria is defined next.

Definition 22.2 Let $(a_1, \theta_1), (a_2, \theta_2), (a, \theta) \in A \times \Theta$. The pair (a_1, θ_1) dominates the pair (a_2, θ_2) if $u(a_1, \theta_1) > u(a_2, \theta_2)$. A consistent decision state (a, θ) is optimal if there is no other consistent decision state that dominates it.

Consider the solution to the BDP (a^b, θ^b) in (22.6), so $\theta^b = f(a^b)$. Using (22.5), let a' be the solution to

$$u_1(a', f(a^b)) + u_2(a', f(a^b))f'(a^b) = 0. \quad (22.7)$$

Then, it follows that

$$u(a^b, \theta^b) < u(a', f(a^b)) \leq u(a^c, f(a^c)). \quad (22.8)$$

The first inequality in (22.8) follows from (22.7), while the second follows from the optimality (but also possibly non-uniqueness) of the optimal solution $(a^c, f(a^c))$ to the CDP. Hence, the solution to the BDP is welfare dominated by the solution to the CDP. Hence, one might think of designing policy interventions to improve the decisions taken by behavioral decision makers.

It is useful to distinguish between the following classes of policy interventions. The first two play a major role in neoclassical economics that does not take into account the existence of behavioral decision makers; yet, many such interventions appear designed for a world that has behavioral decision makers. The third plays a prominent role in behavioral welfare economics.

1. *Direct paternalism*: A third party may directly choose an action for the behavioral decision makers that forces them to take account of the feedback effect. Alternatively, the third party may make these actions legally binding. In neoclassical economics, direct paternalism takes the form of legal requirements such as wearing seat belts; minimum driving age; minimum age for various kinds of consent; prohibition on using mobile phones while driving; purchase of compulsory car insurance; and the definition and prohibition of publicly obscene behavior.

2. *Indirect paternalism*: This takes the form of policy interventions, typically through changes in prices, that induce decision makers to internalize externalities caused to others, or to themselves. Examples are vehicle road taxes based on CO₂ emissions, emissions taxes on factories, subsidies to charitable giving, and child tax credits.
3. *Libertarian paternalism or soft paternalism*: These interventions (see Section 22.3, below) impose low costs on classical decision makers. However, a third party, typically a benevolent planner, controls the menu of options, for instance, defaults, faced by behavioral decision makers. This is a form of *nudging* and embodies various degrees of paternalism, depending on how intrusive the nudges are. The behavioral decision makers are, however, free to choose an option from the menu given, for instance, they can override the default option without the threat of legal sanctions that arises under direct paternalism. This aspect of the intervention respects individual liberty.

Hence, these policies are said to exhibit *libertarian paternalism*. Given that individuals ultimately choose their own actions, despite facing a paternalistic third party, this class of interventions is also sometimes given the generic name of *soft paternalism*. Dalton and Ghosal (2011) use this term to include psychological therapies such as cognitive behavioral therapy and expert advice.

In the context of the simple model used in this section, the nudge takes the form of a choice by the planner of some $\bar{\theta} \in \Theta$, where $\bar{\theta}$ may be thought as a reference point, or a default option.

Insofar as behavioral interventions are about making recommendations, for instance, to nudge consumers, it is worth asking who the relevant *addressee* of such recommendations is. In neoclassical economics one does not typically ask this question. The normative benchmark addressee is a benevolent social planner. This is also the approach taken in most of modern behavioral welfare economics such as under soft paternalism as well as in choice-based or revealed preference approaches; see Sections 22.6 and 22.7, below.

In contrast, Sugden (2008, 2013), in advocating the *contractarian approach* (see Section 22.8, below) questions the usefulness of the social planner benchmark on the grounds that this is a purely hypothetical construct (i.e., lacks a valid addressee). His approach is to respond to behavioral biases by identifying directly the relevant addressee, i.e., the individual who exhibits these biases. This takes the form of policy interventions that provide information “directly” to individuals who “choose” to be nudged.

Once we move away from the normative benchmark of a social planner, the relevant addressee for behavioral recommendations could be an executive agency, a regulatory body, or any policy making department in the government. Such an entity, which we simply term as a policymaker, might have the following two kinds of informational disadvantages, singly, or in conjunction.

Definition 22.3

- (a) *Under Type 1 asymmetric information, the policymaker is uncertain about the degree of behavioral biases of individuals. We capture this by assigning respective probabilities $\mu_1 \in [0, 1]$ and $1 - \mu_1$ that the decision maker solves a CDP and a BDP.*
- (b) *Under Type 2 asymmetric information, the policymaker is uncertain about the correct preferences of the decision maker but knows the type, classical or behavioral, of the decision maker. The policymaker assigns a probability $\mu_2 \in [0, 1]$ to the correct utility function (v for a CDP and u for a BDP). But with probability $1 - \mu_2$, the policymaker has the following incorrect beliefs about the decision maker's preferences.*

- (bi) If the decision maker is classical, then the policymaker believes that preferences are given by $\tilde{v}(a) \equiv \tilde{u}(a, \tilde{f}(a))$, where $\tilde{v} : A \rightarrow \mathbb{R}$ is twice continuously differentiable.
- (ii) If the decision maker is behavioral, then the policymaker believes that preferences are given by $\hat{u}(a, \theta)$, where $\hat{u}$ is twice continuously differentiable in both arguments at least up to order two and has continuous cross partial derivatives.

We now consider the solution under various regimes.

22.2.2 Direct paternalism

Under direct paternalism, the policymaker directly chooses an action for the individual under one of the two kinds of information asymmetries in Definition 22.3. Consider first Type 1 asymmetric information. The objective function of the policymaker is

$$a_1^d \in \arg \max_{a \in A} \mu_1 u(a, f(a)) + (1 - \mu_1) u(a, \theta), \quad (22.9)$$

and the consistency requirement is

$$\theta_1^d = f(a_1^d). \quad (22.10)$$

The first order condition to the optimization problem in (22.9) is

$$\mu_1 \left[u_1 \left(a_1^d, f(a_1^d) \right) + u_2 \left(a_1^d, f(a_1^d) \right) f'(a_1^d) \right] + (1 - \mu_1) u_1(a_1^d, \theta_1^d) = 0. \quad (22.11)$$

Jointly solving (22.10) and (22.11) we get the solution (a_1^d, θ_1^d) ; a solution always exists because the objective function is continuous over a compact interval. In the special cases of $\mu_1 = 1$, and $\mu_1 = 0$, respectively, we get the solution under CDP and BDP.

Consider Type 2 asymmetric information in which the planner believes that the decision maker solves a CDP, but is uncertain of the exact preferences (Definition 22.3(bi)). The objective function of the policymaker is

$$a_2^d \in \arg \max_{a \in A} \mu_2 u(a, f(a)) + (1 - \mu_2) \tilde{u}(a, \tilde{f}(a)), \quad (22.12)$$

subject to the consistency requirement

$$\theta_2^d = f(a_2^d). \quad (22.13)$$

The first order condition to the problem in (22.12) is

$$\begin{aligned} \mu_2 \left[u_1 \left(a_2^d, f(a_2^d) \right) + u_2 \left(a_2^d, f(a_2^d) \right) f'(a_2^d) \right] + (1 - \mu_2) \left[\tilde{u}_1 \left(a_2^d, \tilde{f}(a_2^d) \right) \right. \\ \left. + \tilde{u}_2 \left(a_2^d, \tilde{f}(a_2^d) \right) \tilde{f}'(a_2^d) \right] = 0. \end{aligned} \quad (22.14)$$

Jointly solving (22.13), (22.14), we get the solution (a_2^d, θ_2^d) . If there is no misperception of preferences ($\mu_2 = 1$) then the solution is identical to the CDP, (a^c, θ^c) .

Interventions under direct paternalism are more attractive if μ_1 is low (i.e., it is more likely that the decision maker is solving a BDP), and if μ_2 is high (the policymaker has greater knowledge

of the underlying preferences). At the other extreme (high μ_1 and low μ_2), direct paternalism may do more harm than good.

22.2.3 Indirect paternalism

Under indirect paternalism, the policymaker is the Stackelberg leader and imposes taxes/subsidies, t , on the actions taken by the decision maker; t is also known as a Pigouvian tax. The behavioral decision maker observes t , takes it into account, and then chooses his optimal action. Thus, the BDP in this case is

$$a^i \in \arg \max_{a \in A} u(a, \theta) - ta,$$

where t is the tax ($t > 0$) or subsidy ($t < 0$) per unit of action, a , imposed by a policymaker. The solution is given by the analog of (22.6):

$$u_1(a^i, f(a^i)) - t = 0 \text{ and } \theta^i = f(a^i). \quad (22.15)$$

In the absence of any information asymmetries, the policymaker can implement the solution to the CDP. Comparing (22.5) and (22.15), and choosing

$$t = -u_2(a^c, f(a^c))f'(a^c), \quad (22.16)$$

ensures that the solution is $(a^i, \theta^i) = (a^c, \theta^c)$. The tax/subsidy gives the correct incentives to the decision maker to internalize the own-produced externality of ignoring the feedback effect.

We now consider the presence of asymmetric information. We distinguish between two kinds of solutions. The solution, which arises under asymmetric information on the part of the policymaker (the ex-ante solution) and the solution that arises when the asymmetric information is resolved (the ex-post solution). For instance, the policymaker may assign probabilities that the decision maker to be regulated is classical or behavioral (asymmetric information). Yet, once the appropriate regulations are announced, the actual decision maker, either classical or behavioral, takes actions under the given regulations (resolution of asymmetric information).

Proposition 22.1 (Dalton and Ghosal, 2011): Consider Type 1 asymmetric information (Definition 22.3a).

- (a) Ex-ante solution: The following level of Pigouvian tax, $t = t_1$, ensures the ex-ante solution is identical to that in the CDP, (a^c, θ^c) ,

$$t_1 = (1 - \mu_1) u_1(a^c, \theta^c). \quad (22.17)$$

- (b) The ex-post solution: The ex-post solution, i.e., once the decision maker turns out to be a classical or behavioral decision maker, is not necessarily identical to (a^c, θ^c) .

Proof: (a) Consider Type 1 asymmetric information. The objective function of the policymaker is

$$a_1^i \in \arg \max_{a \in A} \mu_1 [u(a, f(a)) - t_1 a] + (1 - \mu_1) [u(a, \theta) - t_1 a] \quad (22.18)$$

subject to the consistency requirement

$$\theta_1^i = f(a_1^i). \quad (22.19)$$

The first order condition to this optimization problem is

$$\mu_1 [u_1(a_1^i, f(a_1^i)) + u_2(a_1^i, f(a_1^i))f'(a_1^i) - t_1] + (1 - \mu_1) [u_1(a_1^i, \theta_1^i) - t_1] = 0, \quad (22.20)$$

which can be rewritten as

$$u_1(a_1^i, f(a_1^i)) + u_2(a_1^i, f(a_1^i))f'(a_1^i) = \frac{1}{\mu_1} [t_1 - (1 - \mu_1) u_1(a_1^i, \theta_1^i)]. \quad (22.21)$$

Setting $t_1 = (1 - \mu_1) u_1(a^c, \theta^c)$ ensures that the RHS of (22.21) equals zero. Comparing with (22.5) and using (22.19), the solution is $(a_1^i, \theta_1^i) = (a^c, \theta^c)$, as claimed. This is the ex-ante solution, prior to the resolution of the asymmetric information faced by the policymaker.

(b) The Stackelberg leader, the policymaker, announces the Pigouvian tax given in (22.17), under uncertainty about the type of the decision maker: classical or behavioral. Having observed the tax, the decision maker, or Stackelberg follower, who privately knows his type to be either classical or behavioral, chooses an action $a \in A$. The ex-post solution, say $(\tilde{a}, \tilde{\theta})$, is unlikely to equal (a^c, θ^c) . For instance, the optimization problem of a behavioral decision maker is given by

$$\tilde{a} \in \arg \max_{a \in A} u(a, \theta) - [(1 - \mu_1) u_1(a^c, \theta^c)] a,$$

subject to the consistency condition $\tilde{\theta} = f(\tilde{a})$. A similar argument can be applied if, ex-post, the decision maker solves a CDP. ■

The difference between the ex-ante and ex-post solutions reflects the distortions created by an indirect policy intervention. From (22.17), the size of the intervention depends on the size of μ_1 . The size of the distortion diminishes as $\mu_1 \rightarrow 1$, i.e., the closer the decision maker turns out to be a classical one.

22.2.4 Libertarian paternalism

Under libertarian paternalism, boundedly rational individuals are *nudged* in a beneficial direction by a benevolent planner and minimum costs are imposed on the fully rational individuals. We consider a simple example that illustrates some of the implications of libertarian paternalism. Suppose that the action set $A = \{a_l, a_h\}$, where $a_l < a_h$ and a_l is the status quo. To fix ideas, think of a_l as the action “stay at home” and a_h as the action “go to college.” The set of psychological states is given by $\Theta = \{\theta_l, \theta_h\}$; let us interpret these states as the aspiration levels. The utility function of the individual is given by

$$\begin{cases} u(a_l, \theta) = 0, \theta \in \{\theta_l, \theta_h\} \\ u(a_h, \theta_l) < 0, u(a_h, \theta_h) > 0 \end{cases} \quad (22.22)$$

Thus, staying at home gives utility, normalized to zero, irrespective of the psychological state. Going to college is a relatively better option under the good psychological state, but not if the psychological state is bad (e.g., the individual cannot study properly if depressed, and fails, but

must pay fees and other college expenses). Denote the optimal choice of an action under libertarian paternalism by a^* .

Behavioral decision makers ignore the feedback between actions and psychological states. Hence, under the assumptions made so far, they choose as follows.

$$\begin{cases} u(a_l, \theta_l) > u(a_h, \theta_l) & \Rightarrow a^* = a_l \\ u(a_l, \theta_h) < u(a_h, \theta_h) & \Rightarrow a^* = a_h \end{cases} \quad (22.23)$$

In contrast, classical decision makers take account of the feedback function, f . Assume that

$$\theta_l = f(a_l) < \theta_h = f(a_h). \quad (22.24)$$

A classical decision maker realizes that if he went to college, his aspiration level will be raised (see (22.24)), thus, going to college is a dominant action for such a decision maker.

In contrast, for behavioral decision makers, the decision is critically dependent on the initial aspiration levels (see (22.23)). Those who have low aspiration levels, $\theta = \theta_l$, choose not to go to college, so $a = a_l$. In contrast, those with high aspirations, $\theta = \theta_h$, do go to college, so $a = a_h$. Since $u(a_l, \theta_l) = 0$, and $u(a_h, \theta_h) > 0$, if the initial aspiration level of a behavioral decision maker is θ_l , a nudge to the decision maker that either raises aspiration levels to θ_h , or makes going to college a *default option*, is welfare improving. This does not alter the optimal decision of the classical decision makers but improves the decision of the behavioral decision makers. This is the basic idea behind libertarian paternalism. For an interpretation of θ as aspirations in a model that explains poverty arising from aspiration failures, see Dalton et al. (2016).

22.2.5 Effect of other policies

In neoclassical economics, providing more information to decision makers, or expanding their set of actions typically improves the welfare of individuals. In the present framework, this need not be the case because behavioral decision makers ignore the feedback between actions and psychological states. Exercise 5 illustrates these issues.

22.3 Soft paternalism

In economics, *paternalism* may be taken as any infringement of *consumer sovereignty* that is intended to improve the well-being of economic actors. For instance, parents might choose healthy snacks for children and limit access to the TV, or the kinds of websites they can visit on the Internet. Paternalism has been viewed with suspicion in neoclassical economics. However, as Jolls et al. (1988, p. 1545) note in their *anti-antipaternalistic approach*: “No axiom demonstrates that people make choices that serve their best interests; this is a question to be answered based on evidence.”

Soft paternalism is an umbrella term that is associated with several kinds of behavioral welfare frameworks that advocate limited and light interventions, for instance, the choice of default options. These interventions impose little or no costs on classically rational individuals but intend to improve the welfare of boundedly rational individuals by suggestions or *nudges*, while allowing them the liberty to choose their actions, e.g., overriding the default option.

22.3.1 *Libertarian paternalism*

The main target for soft paternalism is boundedly rational individuals. The interventions are designed to influence the choices of such individuals for the better; in this sense the interventions are *paternalistic*. Yet, the interventions do not infringe the liberty of the individuals to make the eventual choices; in this sense the interventions are liberal. Thaler and Sunstein (2003, 2009) and Sunstein and Thaler (2003) argue that their *libertarian-paternalism* proposal is not an oxymoron. In their proposal, Sunstein and Thaler take it that they are advising a *planner* or a *choice architect* whose job is to decide on the options to be presented to an individual. The planner decides on the menu of choices for individuals, based on his judgment about what options are in their best interests, as would be judged by the individuals if one could invoke their underlying rational selves.

Sunstein and Thaler (2003) note that most economists are libertarians, a position associated with the respect for freedom of individual choice. They argue that the unfavorable response to paternalism in economics and law, which they call *dogmatic paternalism* (p. 1163), is based on one false assumption and two misconceptions that we describe now.

The false assumption is that (Sunstein and Thaler, 2003, p. 1163): “almost all people, almost all of the time, make choices that are in their best interest or at the very least are better, by their own lights, than the choices that would be made by third parties.” A range of evidence in behavioral economics that runs contrary to this assumption is highlighted in this book; it is also briefly alluded to in the introduction. Thaler and Sunstein mainly refer to the evidence on context dependence of preferences and the instability of preferences. Within behavioral economics, this is uncontroversial. Furthermore, individual choices are influenced by the individual’s psychological state and by self-control problems. Finally, the choices made by individuals are sometimes not consistent with the choices that they endorse in their role as consumers. For instance, this is the case when individuals visit a rehab centre to kick a habit or prefer the median portfolio of a savings plan to their own (Sunstein and Thaler, 2003, p. 1169).

If a choice architect or planner has to act in a paternalistic manner then he needs to discover what is best for an individual. Sunstein and Thaler (2003, p. 1162) treat individuals as not acting in their best interests if they make “decisions that they would change if they had complete information, unlimited cognitive abilities, and no lack of self-control.” Sugden (2008, 2013) terms this as the *informed desire criterion*. How does the choice architect figure out what is best for the individual? The choice architect could use a cost–benefit analysis or employ insights from psychology to figure out when more options will be beneficial (Sunstein and Thaler, 2003, p. 1166). Or, the choice architect could use simple rules of thumb and experiment as in automatic defaults and the Save More Tomorrow plan (Sunstein and Thaler, 2003, p. 1190).

The two misconceptions, in Thaler and Sunstein’s view, are the following (Sunstein and Thaler, 2003, p. 1163). First, there is a viable alternative to paternalism. Second, paternalism necessarily involves coercion. The behavioral biases and misperceptions of individuals show that they can potentially benefit from the “steer” provided by the choice architect (subject to the caveats about the choice architect’s role discussed below). Furthermore, paternalism need not be coercive. The essence of libertarian paternalism is that the choice architect respects the freedom of choice of the individual, once a menu of choices is offered to the individual. This view is not universally agreed upon. For instance, a literature in Law and Economics defends the rationality assumption and views paternalism with suspicion on account of regulatory capture and the bounded rationality of bureaucrats (Posner, 1986).

In Thaler and Sunstein's view, a policy is paternalistic if it alters the choices of individuals in a manner which makes them "better off". The term better off is not formally defined. Furthermore, in this view, primacy is given to preferences rather than choices, for choice might, for a variety of reasons, not faithfully reflect preferences. If this is the case and the two misconceptions are accepted, then Thaler and Sunstein (2003, p. 175) conclude that: "we can abandon the less interesting question of whether to be paternalistic or not and turn to the more constructive question of how to choose among paternalistic options." Let us clarify these issues further with a concrete example.

Example 22.1 *(The cafeteria example; Sunstein and Thaler, 2003): The "cafeteria example" features prominently in the work of Thaler and Sunstein and nicely illustrates the essentials of libertarian paternalism. It is typically the case that the arrangement of food items (a form of framing) has a significant effect on what people choose in a cafeteria or a supermarket. For instance, people might be more inclined to choose items that are placed at "eye level" or items that are placed at the front end. Consider a cafeteria in which the manager must decide on the arrangement of food items in the store. Suppose that the cafeteria manager has the following options:*

1. *She could make choices that she thinks would make the customers best off, all things considered.*
2. *She could make choices at random.*
3. *She could choose those items that she thinks would make the customers as obese as possible.*
4. *She could give customers what she thinks they would choose on their own.*

Anti-paternalism dictates choosing option 4. However, it is difficult to determine what the consumers would choose on their own if the cafeteria manager were not to first choose one of the options. As Sunstein and Thaler (2003, p. 1164) write: "If the arrangement of the alternatives has a significant effect on the selections the customers make, then their true 'preferences' do not formally exist." Ruling out option 4, the only reasonable option is option 1. Hence, a degree of paternalism is inevitable. Furthermore, option 1 illustrates why coercion is not an inevitable part of paternalism in the following sense. The cafeteria manager simply engages in framing the choices without reducing the number of choices. This should hardly have an effect on the neoclassical fully rational consumers, so it imposes no change in their well-being. However, by perhaps placing the healthy options at eye level, the cafeteria manager may improve the well-being of boundedly rational consumers who might be subject to self-control problems and choose unhealthy snacks otherwise.

One may also think about the likely outcomes if the cafeteria manager were purely a profit maximizer (Sugden, 2008, p. 236). However, if the cafeteria manager believes that current preferences are not well formed and she anticipates that in the longer run consumers will like healthier options, then profit maximizing and welfare issues could be closely related.

In actual practice, the job description of a choice architect already exists in many domains. For instance, doctors often advise patients on a choice menu of treatments chosen by them, architects often advise clients about house designs, interior decorators often choose a subset of designs to show to clients, based on their perceived suitability, and so on. However, the strict definition of a choice architect is satisfied when the suggested interventions are designed to make individuals better off, *as judged by themselves* (Thaler and Sunstein, 2009, ch. 5). Furthermore, the choice architect does not unduly limit the freedom of choice of individuals. Hence, *nudges* are not only

admissible, they are the primary means of producing outcomes from the choice architecture. Nudges are believed to be most effective when the decisions of individuals are infrequent and their decisions give rise to noisy feedback; when the choice architect has more experience and expertise; and when individual heterogeneity is low (Thaler and Sunstein, 2009, ch. 4).

In the actual practice of public policy, one observes a wide variety of nudges that could be imagined to originate from a choice architect. For instance, cash machines dispense cash only when the relevant bank card has been taken out first; nozzles for dispensing different kinds of car fuel, e.g., diesel and petrol, are often of a different size or color to prevent incorrect use of fuel; many cars produce a beeping sound if the driver is not wearing a seat belt; automatic electric switches in cars and offices are often used to conserve energy or prevent batteries from running out; default options in savings and pension plans enable people to invest better; the Save More Tomorrow plan takes account of current self-control problems to enable greater savings in the future as income grows over time (see Section 22.3.4 below); color coded labels on food (red for high fat and green for low fat) and government mandated guidelines on the number of fruit portions a day; warnings on cigarette packs; and prohibition of advertising of cigarettes.

Thaler and Sunstein (2009, pp. 102–3) give an example of a regulation called RECAP, short for *record, evaluate, and compare alternative prices*, that allows better across-firms comparability of prices. This regulation calls for a personalized electronic summary of the consumer's use of the good/service sold by the firm in a standardized form, that allows comparison of prices from alternative suppliers. This is similar, in spirit, to the practice in UK supermarkets of separately stating sticker prices of goods of different brands/sizes in standardized units, say per 100 grams, or per 100ml.

In each case, the justification for these nudges in the libertarian-paternalistic framework is that some people are making undesirable choices, as judged by themselves (Thaler and Sunstein, 2009, p. 5). Thaler and Sunstein present a range of evidence to underpin the claim that individuals would wish to be nudged in many cases. Examples include, preference for the median savings portfolios relative to the own portfolio; New Year resolutions for smoking and other addictions that often fail, or activities that involve procrastination; evidence from statistics that show health related risks from temptation activities such as smoking, drinking, and obesity inducing foods (Thaler and Sunstein, 2009, ch. 1, 4). This evidence is certainly impressive but some have doubted if it clinches the argument that people wish to be nudged (Sugden, 2009, p. 371).

22.3.2 Asymmetric paternalism

Concurrently and independently of the development of the idea of libertarian paternalism, Camerer et al. (2003) propose the closely related concept of *asymmetric paternalism* or *paternalism for conservatives*. They define and describe the main features of asymmetric paternalism as follows (p. 1212). “A regulation is asymmetrically paternalistic if it creates large benefits for those who make errors, while imposing little or no harm on those who are fully rational. Such regulations are relatively harmless to those who reliably make decisions in their best interest, while at the same time advantageous to those making suboptimal choices.”

How should alternative paternalistic policies be evaluated? Camerer et al. (2003) propose a simple cost–benefit framework. Suppose that a fraction p of the consumers are boundedly rational, while the rest are fully rational. Consider a proposed paternalistic policy that is implemented by incurring some implementation costs, I . The policy imparts some benefits $B > 0$ to the bound-

edly rational consumers, but also imposes some costs $C > 0$ on the fully rational consumers. Suppose that the incremental effect on the profits of the firms on account of the change in consumer behavior is $\Delta\pi$ which could be positive or negative. The proposed policy is beneficial if

$$pB - (1 - p)C - I + \Delta\pi > 0. \quad (22.25)$$

An asymmetrically paternalistic policy is characterized by a high value of B and low values of C and I . In the limit, the purest case of asymmetric paternalism arises when $C = 0$. There is still no guarantee that I will be low, although this would seem to be an important requirement of asymmetric paternalism. Asymmetric paternalism becomes more attractive as the proportion of boundedly rational consumers, p , increases. The effect on and sign of $\Delta\pi$ depends on the precise context. For instance, if bounded rationality creates errors in the direction of buying too little, then under an asymmetric paternalistic policy, $\Delta\pi > 0$; and in the converse case, $\Delta\pi < 0$. While (22.25) is useful to organize our thoughts, in actual practice, the measurement of costs and benefits in the presence of behavioral biases is a vexed one. For instance, perceived benefits, B , may depend on behavioral factors such as one's psychological state, framing, context, and status-quo bias, which may be inherently hard to measure.

The class of policies that alters "default outcomes" comes closest to satisfying the requirements of asymmetric paternalism. This is particularly the case when a status-quo bias arises from procrastination. Think, for instance, of two alternative states s_1 and s_2 with respective benefits B_1 and B_2 such that $B_1 > B_2$; the states are chosen by the individual. If the status quo is s_1 and individuals exhibit a status-quo bias that prevents them from choosing s_2 then a paternalistic policy might wish to set the default option to s_2 . An example of a default option is automatic enrollment in 401(k) pension plans, which has been shown to improve participation rates (Madrian et al., 2001).

The paternalistic policy of automatic enrollment in 401(k) plans also alerts us to a careful consideration of unexpected aspects of the problem. In this case, in actual practice, the default enrollment meant choosing also the default contribution rates and the default portfolio of assets. The default rates were set quite low (2–3%) and the default portfolio was a conservative one (mostly money market funds). These defaults, however, turned out to be disadvantageous. The low default rates reduced the effectiveness of the default, while the conservative portfolio earned lower returns because stocks have traditionally outperformed bonds. This offset some of the beneficial effects of the increased enrollment rates arising from the new default option (Choi et al., 2004; Madrian et al., 2001).

A second example of unexpected outcomes comes from the detailed text that accompanies mortgage and lease contracts in the US. The detail in the text is a response to the Federal Truth in Lending Act. The text typically gives details of all costs and terms of the contract; multiple, mandatory, signatures certify that the reader has read the separate sections of the contract. Borrowers have to read the following statement and sign that they understand it: "If you obtain this loan, the lender will have a mortgage on your home. You could lose your home, and any money you have put into it, if you do not meet your obligations under the loan." This text exemplifies asymmetric paternalism because the rational types will not need to read the fine print while the boundedly rational types are expected to benefit from reading it. Furthermore, once the text is designed, the cost of producing copies is relatively low. However, a potentially unexpected outcome is that because individuals have limited attention (see Part 7 of the book), the boundedly rational types simply ignore the fine print, in which case the policy might have limited benefits. Indeed, such individuals might have read a more abbreviated version of the text (see the

relevant evidence in Part 7). The condition in (22.25) advocates taking account of all such relevant costs.

Many existing regulations discussed in this section can be interpreted as a form of asymmetric paternalism. Indeed, it is hard to give an interpretation of these regulations on traditional grounds of efficiency when all economic agents are rational. Hence, somewhat paradoxically for neoclassical economics, one might argue that the law reflects asymmetric paternalism (Camerer et al., 2003, p. 1223).

22.3.3 A pragmatic approach

Loewenstein and Haisley (2008) propose a pragmatic approach, light paternalism, in which welfare judgments can be made under two main conditions that they term as *dominance* and *self-officiating*.⁴ The dominance criteria comes into play if one strategy clearly dominates another, in terms of outcome or risk. For instance, full utilization of 401(k) matching programs typically dominates underutilization. Thus, light paternalism would advocate a default option that equals the maximum allowable contribution that is matched by the employers.

Self-officiating means that individuals choose their own goals and then they are provided with the necessary information and incentives to aid self-control. The choice of goals is to be made in a “cold” rather than a “hot” state (see Part 6 for this distinction). Consider the example of an office in which employees determine that they would be better off without a soda machine that dispenses unhealthy drinks. The self-officiating criterion allows for the removal of the machine from the premises provided that the employees make this decision in a cold state. This appears relatively heavy-handed but allows individuals to pursue their stated goals more directly.

The pragmatic approach also differs from the choice-based approach of Bernheim and Rangel (2007) in that it focuses on preferences rather than choice. Loewenstein and Haisley (2008, p. 223) capture the spirit and thrust of their approach in the following quote.

Our own opinion is that the welfare criterion for evaluating paternalistic policies should be based on preference. Much as a psychotherapist would likely take at face value a client's professed desire to become happier, more sociable, or less anxious, even if she engaged in patterns of thinking and behavior that led to the opposite result, we would argue that the economist-as-therapist should treat verbal statements of preference as useful information, even if choice is not in line with professed preference. If people express a desire to lose weight but make choices that cause them to gain weight; if they express a desire to be financially solvent, but make choices that lead to burdensome debt; if they want to stop smoking but continue to smoke; if they want to take prescription medications but fail to do so, these are all situations in which paternalistic interventions could be helpful.

22.3.4 Some applications of soft paternalism

We briefly consider some applications of soft paternalism. Our discussion is abbreviated because these issues have been considered in varying degrees elsewhere in the book.

⁴ A third criterion, *clearly negative outcomes*, is also identified. However, it appears that in a broader definition this criterion can be subsumed within the dominance criteria.

THE ROLE OF DEFAULTS

The status-quo bias is a well-documented finding in behavioral economics.⁵ Suppose that there are two possible actions/policies, x and y , with respective benefits to an individual equal to b_x and b_y . Suppose that the cost of changing from any action to the other is c . Then there is a status-quo bias in favor of x if $b_y - b_x > c$, but the individual prefers to persist with x . Concurrently, there might be a status-quo bias in favor of y too.

Camerer et al. (2003) note several sources of status-quo bias. (1) Loss aversion: If a change in policy is associated with a potential loss, then the status-quo policy is preferred. (2) Asymmetric experienced regret: There is some evidence that people regret bad outcomes that result from their own actions as compared to bad outcomes that result purely from inaction. (3) Hyperbolic discounting: Current-bias might cause individuals to procrastinate in shifting to the economically more beneficial policy.

Status-quo bias opens the possibility that paternalism, in the form of default options, in the form of the policy associated with higher welfare, may be effective. The status-quo bias ensures the individuals are unlikely to override default options. The stickiness of default options has been widely documented: for instance, in decisions to donate one's organs (Johnson and Goldstein, 2003; Abadie and Gay, 2004); default options in car insurance plans (Park et al., 2000); and pension and savings plans (Thaler and Sunstein, 2009). Defaults can be viewed as a form of asymmetric paternalism because they can be set in a manner that is beneficial for a boundedly rational consumer, but impose little or no cost on the rational individuals who can easily override them.

Consider the following natural experiment that illustrates the effect of default rules for insurance policies. In New Jersey, the default was a limited right to sue the other party following an accident, while, in Pennsylvania, the default was a full right to sue. In each state, customers could switch to the other possibility by either paying more (full right to sue in New Jersey) or by receiving a discount (for a limited right to sue in Pennsylvania). Johnson et al. (1993) show that when offered a choice to alter the defaults, only 20% of drivers in New Jersey chose full rights to sue while 75% of drivers in Pennsylvania retained full rights to sue.

RETIREMENT SAVINGS AND DEFAULTS

We have shown in Part 3 that in the presence of present-biased preferences, such as hyperbolic discounting, individuals may save insufficiently for retirement. Other behavioral factors may also reduce savings; we give two examples. (1) By reducing current consumption relative to a reference level of consumption, an individual with a mental account for consumption may suffer loss aversion. Thus, individuals might be reluctant to save more. (2) If the default is zero savings/contributions to a retirement plan, the status-quo bias could further weaken the desire to save.

The idea behind Benartzi and Thaler's (2004) SMarT ('save more tomorrow') plan is to take account of behavioral features in such a manner that encourages more savings. Essentially, under a SMarT plan, individuals agree to save a part of their next pay increase. Since only a fraction of the pay increase in the future is used towards savings, loss aversion from a fall in consumption does not kick in sufficiently. Furthermore, since the saving is made in the future, it is not influenced by the current-bias in preferences arising through, say, hyperbolic discounting. Finally,

⁵ See, for instance, Samuelson and Zeckhauser (1988), Kahneman et al. (1991), and Kahneman and Tversky (2000).

the status-quo bias can be exploited by altering the default option to contribute rather than not-contribute. In the spirit of soft paternalism, individuals can choose to opt-out if they wish to.

Evidence supports the “savings enhancing effects” of changing the default option (Choi et al., 2004; Madrian and Shea, 2001). An alternative to using the default option is to ask employees when they begin their jobs if they would like to enroll or not. Choi et al. (2004) show that this method increases enrollment relative to the case of the opt-in option (default is not-contribute) but less relative to the opt-out option (default is to contribute).

Loewenstein and Haisley (2008) note two limitations of empirical studies that suggest a positive welfare effect of paternalistic savings policies. First, where does an increase in retirement savings come from? If other items of expenditure are compromised and there is an associated increase in credit card debt, then these factors need to be factored in. Second, a range of behavioral factors come into play in many policies. For instance, in the operation of SMarT savings plans, these factors include hyperbolic discounting, loss aversion, and status-quo bias. What is the individual contribution of each of these factors? These are pertinent questions for future research.

OBESITY

Loewenstein and Haisley (2008) report on their research on the effects of deposit contracts in countering obesity. People who would like to lose weight are asked to pay a deposit of \$90/month, or \$3 a day. Individuals are set a target of losing a pound of weight over a week. They are weighed each day against a target. If the weight is below the target, the experimenter doubles the deposit and gives the money, \$6, to the subject. If the weight is greater than the target then the subject loses the deposit.

The scheme works on two behavioral principles; the various components are described in Parts 3 and 7 of the book. First, initial overconfidence contributes to the decision to join the plan. One could also use a hyperbolic discounting argument similar to the optimistic take-up of annual versus daily offers in gym memberships. Second, a loss of the deposit money is viewed as out-of-pocket costs, which are overweighted relative to forgone-gains, on account of loss aversion. Hence, individuals are more reluctant to avoid out-of-pocket costs, which contributes to their efforts to get their weight below the target. The take-up of deposit contracts is a voluntary decision (liberalism). Rational individuals have no need for them. However, boundedly rational individuals can benefit from them (paternalism). Furthermore, the idea of respecting individually chosen goals in a cold state of the mind, and aiding in the achievement of these goals is consistent with the *self-officiating criterion* under light paternalism (see Section 22.3.3).

Wisdom et al. (2010) consider two kinds of interventions to reduce calorie intake in fast food restaurants. In the first, individuals are simply provided with information about the calorie content of their food items. This is consistent with recently proposed laws and the contractarian approach that we outline in Section 22.8. The second proposal is a form of nudge, consistent with soft paternalism, that simply makes choosing high calorie foods more inconvenient, for instance by altering the placement of items on a menu.

Both interventions were effective in reducing calorie intake. However, each intervention had drawbacks. Providing calorie information did not reduce the calorie intake of overweight people who presumably are the main target of these interventions. The nudge worked less well in a weak form as compared to a strong form. An unexpected effect of the nudge was that individuals reduced consumption of high calorie sandwiches but increased their consumption of higher calorie side orders and drinks. These results show that a combination of the two policies may work

better and nudge proposals may have unexpected effects on complementary activities. Hence, careful pilot studies must precede actual policy recommendations.

COOLING-OFF PERIODS

Consumers might sometimes regret a purchase decision on account of several factors. These include: buying in an emotional hot state (see Part 6); or being subject to projection-bias that overstates the future need for the purchase (see Part 6); or present-biased preferences, arising, say, from hyperbolic discounting, that make the benefits of the current purchase salient relative to its future costs (see Part 3). Regulation, in the form of *cooling-off periods*, protects consumers from potential future regret arising from current decisions. Cooling-off periods typically take one of two forms.

(a) Individuals are free to choose but allowed to reverse their decision within some legal time horizon, T . This typically occurs with the purchase of consumer durables and many services such as insurance products.

(b) Having made their initial decision, individuals are not free to exercise their choice until some time, $\tau > 0$, has passed. For instance, once one files for a divorce, the divorce is not granted until some time has elapsed. There might be a mandatory period of separation that precedes the divorce. Often, there is also a mandatory period (usually, a few days) between filing for a marriage request and the granting of a marriage license.

Several examples of cooling-off periods for the US, for the two cases, (a), (b) above, are given in Camerer et al. (2003). Cooling-off periods are an example of asymmetric paternalism. Rational individuals do not need such periods, and when cooling-off periods take the form in (b), a delay of a few days is probably not too costly for such people. Indeed, being rational, they would take the waiting period into account. However, for boundedly rational individuals, cooling-off periods can be welfare improving on account of the various behavioral features noted above.

DEADLINES

One implication of hyperbolic discounting is the tendency of individuals to procrastinate in the completion of some task (see Part 3). In a multiple selves interpretation, the current self typically believes that the future selves will suffer from fewer self-control problems. Hence, the current self, rationally, puts off completion of the task to some future date. If this optimism about future self-control problems is misplaced, the task might never be completed.

Deadlines can often be self-imposed. However sometimes they are externally imposed, and are legally binding. The paternalistic aim of deadlines is to ensure the completion of a task when the planner perceives that the benefits exceed the costs, say, in the sense of (22.25). For instance, in order to contribute to the tax-exempt individual retirement account (IRA), individuals not only must respect the constraint of a given maximum amount each year, but also the deadline of April 15. Thus, if individuals procrastinate beyond this deadline, their entitlement to the maximum allowed tax allowance for the year is lost. This magnifies the effect of missing the deadline. Presumably, this deadline should help individuals who procrastinate, to take advantage of the investment opportunity.

Ariely and Wertenbroch (2002) run the following experiment using MIT students. Students had to write three short term papers for a course and they were sorted into two groups. One of the groups was given three equally spaced deadlines for completion of the term papers (the restricted choice condition). The other group could decide to hand in their term papers any time before the end of term (the free choice condition). An identical late submission penalty applied to both groups. Few students in the free choice condition chose to impose equally spaced deadlines on themselves but most did impose some sort of deadlines on themselves. In general, students who

were subjected (externally or internally) to evenly spaced deadlines performed better in terms of grades. One possible reason is that this allowed for better budgeting of time. In the absence of such deadlines, students might try to do too much, too late, lowering their grades.

22.3.5 A critique of soft paternalism

Recall the *informed desire criterion* in our discussion of libertarian paternalism in Section 22.3.1. Essentially, the criterion suggests that individual choices may be suboptimal, as judged by themselves, in the presence of incomplete information, limited cognitive abilities, and self-control problems. Sunstein and Thaler (2003, p. 1162) treat individuals as not acting in their best interests if they make “decisions that they would change if they had complete information, unlimited cognitive abilities, and no lack of self-control.” Sugden (2008, 2013) questions the empirical usefulness of the informed desire criterion on the following grounds.

The concepts of incomplete information, limited cognitive abilities, and self-control need a formal definition. Even if one could define these concepts in a satisfactory manner, it might still be difficult to pin down human behavior if these factors were taken into account. Should we expect preferences to be consistent when we relax these constraints? For instance, a preference for status quo or anchoring might be hardwired, so the informed desired criterion might not eliminate such preferences. This is clearly an unresolved issue and should be a interesting area of research for the future. As Sugden (2008, p. 234) puts it: “Thus, while we can understand how planners might be motivated by paternalism, we are entitled to be sceptical about how far the planners’ judgements about other people’s well-being are justified, either objectively or inter-subjectively.”

How does the choice architect figure out what is best for the individual? The suggestion in libertarian paternalism is that the planner uses cost–benefit analysis (as in (22.25)) or rules of thumb. These recommendations may appear to many readers as vague (Sugden, 2008, p. 233). Guala and Mittone (2015) propose that if it is the case that individual welfare is impossible or too costly to discover, then one may make a case for a political justification for nudging. They argue that since individual biases are likely to cause externalities to others (and to the individuals themselves), the associated social costs may justify nudging interventions on political grounds.

Soft paternalism is relatively silent about issues of political failure. Regulatory capture, where regulators are unduly influenced by special interest groups, and the bounded rationality of bureaucrats, does not ensure that paternalistic choices will be good ones (Glaeser, 2006). However, it is also true that, in a wide range of cases, government policies are the outcome of more deliberate and detailed information processing. These cases are all counterexamples to the proposition that the bounded rationality of governments will necessarily preclude the possibility of any beneficial paternalistic interventions. For instance, laws on cooling-off periods, various default options, mandatory third party car insurance, compulsory seat belts in cars, and reminders to fill-in tax returns by publicly known deadlines. Behavioral economists take great pains to stress caution, field testing, and pilot studies in the application of paternalistic policies.

Gul and Pesendorfer (2008) argue for a separation of the roles of economists as social scientists, in which they attempt to explain economic phenomena, and their roles as advisors/advocates, e.g., for public policy. Economists who do not respect the distinction between these two roles are termed by them as economist-therapists. We examine this critique in more detail in Section 22.5. Gul and Pesendorfer assign a relatively restricted role for economics and economists, a view that is not shared by all economists. Espousing the opposite view, Loewenstein and Haisley (2008, p. 9) write:

Although Gul and Pesendorfer seem to view “therapist” as a pejorative label, we see no reason to not embrace it. Therapy is, in fact, not a bad metaphor for the new types of policies that behavioral economists have been proposing. . . . Just as the therapist endeavors to correct for cognitive and emotional disturbances that detract from the well-being of the patient, such as anxiety, depression, or psychosis, the economist/therapist endeavors to counteract cognitive and emotional barriers to the pursuit of genuine self-interest.

Klick and Mitchell (2006) argue that soft paternalism ignores the long-run costs. By this, they mean (p. 1625), costs that “offset short-run gains because of the negative learning and motivational effects of paternalistic regulations.” Paternalistic policies, it is argued, interfere with the learning-by-doing process via a moral hazard argument. Individuals, because they are protected via paternalistic policies, may have a lower incentive to protect themselves against bad outcomes.

While this argument may be true in some cases, there are several counterexamples that apply in cases where learning is either very slow, feedback is poor, or where the outcomes of one’s actions are known with substantial delay. For instance, if an individual saves inadequately for his retirement on account of present-biased preferences, then by the time that the outcome (low savings, post-retirement) is revealed, it is too late to learn. This critique is also silent about the forms that learning takes. For instance, suppose that learning takes place through reinforcement learning. Then, presumably, soft paternalism, which nudges the individual in the direction of actions that produce a better outcome, reinforce the use of such actions in future.

22.3.6 *A note on behavioral welfare criteria*

Given the discussion so far, what is the appropriate welfare criterion? This is a difficult question on which there appears to be no consensus. This section briefly outlines some of the issues.

One approach is to use self-reported assessments of happiness as a proxy for welfare. This approach draws on the distinction between *experienced utility* and *decision utility*. These issues have already been discussed in Part 6, so our discussion here will be brief. Decision utility refers to the practice in economics of inferring utility functions from the observed choices of individuals. Experienced utility is the hedonic experience associated with an outcome (Bentham, 1789; Kahneman et al., 1997).

Since experienced utility is one’s subjective recollection of the satisfaction derived from a good or service in the past, its measurement can pose problems. For instance, it is well established that self-reported survey questions that report experienced utility are subjected to biases, such as *duration neglect*, and *peak-end evaluation* (Kahneman et al., 1997). However, evidence shows that subjective measures of well-being are consistent with the actual choices of the vast majority of subjects and these measures are the best predictors of choices (Benjamin et al., 2012). As shown in Part 6, such subjective measures successfully predict a range of phenomena such as job quits, probability of suicides, productivity, and probability of a future divorce among couples.

The phenomenal growth in the literature on happiness economics shows that an increasing number of economists take experienced utility as an objective and useful indicator of an individual’s welfare. Some have forcefully advocated the maximization of happiness as the goal of government policy (Layard, 2005). However, others continue to be skeptical (Loewenstein and Haisley, 2008). First, as noted above, the recollection of a past event could be influenced by current moods or emotions and subject to duration neglect and peak-end bias. Second, individuals might adapt to the current level of conditions in their life. For instance, the happiness levels of patients on dialysis were found to be very similar to a control set of healthy patients (Riis et al.,

2005; Ubel et al., 2005). Many would find the resulting policy recommendation that ill-health is not an undesirable outcome to be troublesome.

If people make mistakes in their choices, then one solution is to provide them with the correct information first. Individuals can then maximize their utility in the usual manner. This is the so-called *informed utility approach*. In actual practice, government regulation tries to impart a variety of information to consumers, such as warnings on cigarette packs, content-labeling on food items, and even, in the case of the UK, daily recommended amounts of fruits and vegetables etc. Armed with the appropriate information, the resulting debiased choice of individuals could then form the basis of welfare criteria.

A concern with this approach is that it does not specify the exact method of imparting information; this is important because there are potential framing effects associated with different methods of information provision. Recall, from Chapter 20, the well-known example from Tversky and Kahneman (1981) in which the outcomes are quite different when the instructions are phrased in terms of “lives lost” rather than “lives saved.” Indeed, deciding on the appropriate frame/context within which the information is to be presented will require some welfare criteria, but this takes us back to the original problem of constructing welfare criteria in the presence of behavioral biases. A second source of concern for the informed utility approach is that while it can potentially counter mistakes that are made on account of imperfect information, it does not address the issue of mistakes (relative to the neoclassical benchmark) that are made when all information is available. For instance, individuals might have self-control problems, or they might rely on a range of judgment heuristics (see Parts 3 and 7).

There is growing popularity of choice-based measures of welfare that correspond to decision utility (Bernheim and Rangel, 2007, 2009); see Sections 22.5 and 22.6 below. These measures assume that despite the enumeration of behavioral biases in behavioral economics, there is a strong systematic component to human behavior that is enriched by psychological considerations. Hence, choice data can be used to disentangle the systematic component. However, when choices are inconsistent as, for instance, in the evaluation of new products (see the discussion on coherent arbitrariness in Part 7) then one may trim the existing choices of these troublesome aspects. Once trimmed of “undesirable” choices, the choice criteria retain their usefulness. However, it is not always clear how choices are to be trimmed. Bernheim and Rangel (2009) suggest using non-choice data, for instance, data from brain scans. However, reservations have also been expressed about the ability of neuroeconomic studies to determine which choices are welfare enhancing (Loewenstein and Haisley, 2008).

22.4 Regulation under imperfect self-information

In neoclassical economics, the standard assumption is that consumers have better information about their types relative to firms. In the leading textbook example of asymmetric information, from the area of health insurance, consumers are better informed than firms about their own state of health. The firm then designs clever contracts to truthfully elicit the state of health of the consumers. However, in many cases, consumers might have limited information on their usage of a product, and the usage data collected by the firm could give an informational edge to firms. This could be the case, for instance, in information collected by cell phone companies or Internet search engines.

It is pertinent to explore the effects on consumer welfare when information is provided to consumers of their own usage. Examples include power companies in Britain offering to put

electricity meters in peoples' houses to enable them to monitor their usage on a continual basis. Once consumers are informed of their usage, they might wish to switch firms. Mullainathan et al. (2012) show this to be the case in a randomized field experiment in the usage of Medicare Part D prescription drug plans in the United States; 28% of the group of subjects that received a letter with personalized cost information switched their providers, as compared to 17% who switched in a comparison group.

The RECAP regulation that we considered above, highlights the welfare effects of information provision in this case. We now consider a model, due to Kamenica et al. (2011), that formalizes the effect of RECAP type regulations.

Suppose that consumers are distributed uniformly over the unit interval, $[0, 1]$. Any consumer uses $x \in \{1, 3\}$ units of service, per month, provided by a private firm. Half the consumers use 1 unit and the other half uses 3 units. However, consumers are imperfectly informed of their own private usage, say, due to limited attention.

Each consumer receives an iid signal, $\theta \in \{1, 3\}$, of her usage. Signal accuracy is captured as follows. Let $p(x | \theta)$ be the conditional probability that the usage is x , conditional on the signal, θ . Then,

$$p(x | \theta) = \begin{cases} \lambda \in [\frac{1}{2}, 1] & \text{if } x = \theta \\ 1 - \lambda & \text{if } x \neq \theta \end{cases},$$

thus, the signal is accurate about the true usage with probability λ . We can now determine the expected usage of each consumer, conditional on the signal received, $E[x | \theta]$.

$$E[x | \theta] = x_i p(x_i | \theta) + x_j p(x_j | \theta); x_i \neq x_j \in \{1, 3\}, \theta \in \{1, 3\}.$$

Thus,

$$E[x | \theta = 1] = 3 - 2\lambda \text{ and } E[x | \theta = 3] = 1 + 2\lambda. \quad (22.26)$$

The utility function of each consumer, on receiving the signal, θ , is given by

$$U(\theta, P) = 1 - E[x | \theta] - M, \quad (22.27)$$

where M is the monetary payment from the consumer to the firm. In this reduced-form model, utility is decreasing in expected usage to reflect the fact that often when consumers are informed about their usage, they are likely to be surprised by the level of usage, and often switch providers of the service in the hope of getting a better deal.

Suppose that there are two providers of the service, A and B . Firm A charges a fixed monthly price of $P_A = 2$, irrespective of usage, thus, $M = 2$ for $x \in \{1, 3\}$. Firm B charges a per unit price equal to $P_B = 1$, so $M = x$ for $x \in \{1, 3\}$.⁶ Thus, if consumers knew their own usage, those with $x = 3$ would prefer to buy from firm A and those with $x = 1$ from firm B . A simple calculation shows that consumers who receive a high usage signal will buy from firm A .

$$U(\theta = 3, P_A) - U(\theta = 3, P_B) = 3P_B - P_A > 0.$$

⁶ It is possible to endogenously derive the optimality of different pricing schemes for firms; see, for instance, Piccione and Spiegler (2012).

Similarly, those who receive a low usage signal purchase from firm B .

$$U(\theta = 1, P_B) - U(\theta = 1, P_A) = P_A - P_B > 0.$$

The expected expenditure of all consumers, given that half get a high signal and half get a low signal, is given by $Y = \frac{1}{2}P_A + \frac{1}{2}E[x | \theta = 1]P_B$. Using (22.26) and $P_A = 2, P_B = 1$, we get

$$Y = 1 + \left(\frac{3}{2} - \lambda\right). \quad (22.28)$$

The first term in (22.28) is the expected expenditure of those who receive a high signal and the second term is the expected expenditure of those who receive a low signal.

Proposition 22.2 *Consumer regulation that provides consumers information about their own usage, decreases expected consumer expenditure and increases their welfare.*

Proof: Regulation that provides consumers information about their usage is equivalent to increasing the accuracy of the signal, λ , in informing consumers of their usage. From (22.28), overall expected expenditure of consumers falls, as λ increases. From (22.27), this increases $U(\theta, P)$, which is a measure of consumer welfare. ■

In actual practice, providing information to consumers under regulations such as RECAP will also alter prices, which are assumed given in the model, above. Perhaps intuition might suggest that the resulting desire of consumers to change providers may also put downward pressure on prices. However, Kamenica et al. (2011) give the following useful analogy to show that this need not be the case.

Suppose that two firms with identical costs compete in prices under Hotelling price competition and consumers are uniformly distributed over the interval $[0, 1]$. When consumers do not know their own usage, $x \in \{1, 3\}$, and receive no signals, both firms appear identical to them. Thus, in a Bertrand equilibrium, the firms set prices equal to their marginal costs. However, now suppose that consumer regulation informs each consumer of their usage. Now firms can employ different pricing schemes (e.g., a fixed monthly price or a price based on usage). Based on their level of usage and travelling costs, consumers no longer find the firms to be symmetric. Firms can take advantage of this “new” market power by extracting some of the surplus. Thus, consumers may pay a higher price.

22.5 Choice and non-choice data: What is the scope of economics?

Should economics restrict itself to choice data alone, or should it also be open to the use of non-choice data? In “The case for mindless economics,” Gul and Pesendorfer (2008) present an extreme case for restricting the boundaries of economics to choice data. Indeed, it would not be surprising if many economists hold the same views, although most behavioral economists are unlikely to subscribe to it. Hence, it is important to examine the opposing arguments in more detail. Gul and Pesendorfer’s views can be summarized as follows under four main points, GPI to GP4, below.

GP1. Non-choice data is inadmissible in economics because the scope of economics is restricted to inferring an individual's preferences from observed choices. Using non-choice data, such as neuroeconomic data, or self-reported survey data, or non-incentivized experimental data, is inconsistent with the scope of welfare economics. It is not entirely clear if Gul and Pesendorfer approve of data from experiments conducted under proper incentive compatibility conditions or whether experimental data may also be classified as choice data.

Welfare and choice are identical. If an individual chooses option A over B then his welfare is higher under option A. The reasons why he chose A, for instance, impulsive buying, limited attention, free will, and biological necessity, are irrelevant to economics. They write (p. 5): "The standard definition of welfare is appropriate because standard economics has no therapeutic ambition; it does not try to improve the decision-maker but tries to evaluate how economic institutions mediate (perhaps psychologically unhealthy) behavior of agents."

Non-choice data cannot be used to invalidate neoclassical economics. In particular, data from neuroeconomics has no bearing on the validity of economic models because economic models make no assumptions about the neuroeconomic foundations of choice. They write (p. 7): "Our conclusion is that the neuroeconomic critique fails to refute any particular (standard) economic model and offers no challenge to standard economic methodology."

GP2. There is a limited role for insights from psychology in economics because the two disciplines ask different questions, and have different conventions and terms. They differentiate the scope of economics and psychology by writing (p. 10): "economics and psychology do not offer competing, all-purpose models of human nature. Nor do they offer all-purpose tools. Rather, each discipline uses specialized abstractions that have proven useful for that discipline." And on p. 22 they write: "Economics and psychology differ in the question they ask. Therefore, abstractions that are useful for one discipline will typically be not very useful for the other."

GP3. Economists are not advocates for any normative criteria, even when the "right normative criterion is unambiguous" (p. 29). They give the example of Bayes' rule, which could be advocated on normative grounds, yet economists typically do not do so. They write (p. 35): "Populating economic models with 'flesh-and-blood human' beings was never the objective of economists."

GP4. Neoclassical economics is sufficiently flexible that the observed departures from the neoclassical model can be explained by modifying the neoclassical model without changing its nature. Mistakes or biases that are highlighted in behavioral economics can be accommodated within neoclassical economics by changing the feasible strategies. They give the example of American tourists in the UK who suffer injuries and fatalities because they look to the left (Bernheim and Rangel, 2004). The solution proposed by Gul and Pesendorfer (p. 23) is as follows. The action "look right" is not in the feasible set of American tourists visiting the UK. However, if they are alerted with a sign "look right" then their feasible set changes and they will indeed "look right" eliminating the supposed mistake.

These views, and others expressed in the Gul–Pesendorfer paper, are the subject of an insightful set of contributions in the edited volume by Caplin and Schotter (2008). We mostly focus on three contributions, each with a fairly instructive title: Camerer (2008) replies under the title "The case

for mindful economics.” Schotter (2008) offers the views of an experimental economist under the title “What’s so informative about choice?” Köszegi and Rabin (2008a) give their views under the title: “Revealed mistakes and revealed preference.”⁷

GP1, GP2, and GP3 are subjectively defined boundaries of economics that are extremely narrow. Indeed, great progress has been made in behavioral economics by looking beyond these boundaries. Camerer (2008, p. 44) speaks for many behavioral economists, when, commenting on GP1 and GP2, he writes: “This argument is simply a definition of economics as inherently mindless, and there is no debating a definition. The definition simply draws a preferred boundary rather than makes an evidentiary ‘case’ for mindless economics.” He continues (p. 59): “Much of GP’s paper is linguistic gerrymandering by defining economics as the revealed preference approach, then constantly reminding the reader that anything else is, by their definition, not economics.”

This leaves only GP4 as the potentially interesting argument to ponder over. In fact, this argument goes to the heart of how economists should practice their craft; for a fuller discussion, see the introductory chapter. Economists have been extremely ingenious in constructing theoretical arguments to explain empirical phenomena. Indeed, given enough freedom to make auxiliary assumptions in standard neoclassical models, one can explain almost any observed behavior. Consider, for instance, the solution suggested by Gul–Pesendorfer to the problem of mistakes made by American tourists in GP4. Their suggestion is to keep the neoclassical framework intact, but explain the observation through changes in the feasible set.

A behavioral or an experimental economist would naturally ask: “Does the data on American tourists in the UK support the proposed explanation?” Indeed, many neoclassical economists would also agree with this suggestion. However, so rapid is the pace of theoretical proposals in the neoclassical tradition to explain departures from the neoclassical model that the supply of models has outstripped the supply of diagnostic data (Camerer, 2008, p. 47). In many cases, the freedom in the use of auxiliary assumptions in neoclassical economics to explain observed departures from the model, often by invoking bogus “as if” arguments, accords well with the definition of “ad-hoc” in science.

In some places, Gul and Pesendorfer (2008) appear to protect neoclassical theories from stringent diagnostic testing by suggesting that they should not be taken literally. For instance, commenting on the Kreps–Porteus explanation of a psychological phenomenon (the details are not relevant here), they write (p. 15): “While the formula is suggestive of a mental process, this suggestiveness is an expositional device not meant to be taken literally. The formula encapsulates the behavioral assumptions of the theory in a user-friendly way and thereby facilitates applications of the theory to (more complicated) economic problems”

Among some other important responses made by Camerer (2008) to the Gul and Pesendorfer framework are the following (the details of the neuroeconomic aspects are left for Part 9).

1. Firms that introduce brand new products (or services), and regulators asked to formulate regulations for these products often cannot depend on existing choice data to predict consumer behavior. Hence, choice data, by itself, has limitations to make predictions in some domains. Yet, in actual practice, firms do introduce billions of dollars worth of new products every year after market research that is largely based on non-choice survey data of potential consumers. Yet, most economists will not argue that these economically significant decisions lie outside the scope of economics.

⁷ The reader may also wish to consult the paper by Caplin (2008) in the same volume. Since we have covered similar ground elsewhere in Parts 4 and 6, we omit a repetition of these issues here.

2. When there are several competing theories that explain given choices, then studying the neural basis of these theories can potentially help us to choose among competing theories. Unlike the assertion of Gul and Pesendorfer, important economic variables of interest such as the coefficient of relative risk aversion and loss aversion have been measured using neuroeconomic data (p. 49). The neural correlates of several economic variables, such as discount rates have been discovered (p. 49). Deception is correlated with pupil dilation (non-choice data), hence, it can be used as the basis for predicting the state of the world from a message (p. 51).
3. Economists can benefit by importing terms from psychology (and vice versa). Consider the following case study. Gul and Pesendorfer (2008) give an example of the term *cue* that they feel is unsatisfactory for economists because it lumps together two distinct phenomena: *complements* and *externalities* (e.g., fries might be the cue to eat hamburgers). Hence, they argue (pp. 11, 12): “The concept of a ‘cue’ offers a good illustration of how abstractions from psychology are inappropriate for economics and, conversely, how the corresponding economic concepts are inappropriate for psychology and neuroscience.” They suggest that economists stick to defining terms in a manner that is underpinned by their own terminology and abstractions.
Camerer (2008) argues that precisely because we do not have words in economics that describe goods/states that are complements and externality causing, importing terms such as cues is useful. He writes (p. 55): “So we could adapt the language of economics to describe ‘cues’ as ‘dynamically adaptive, rapidly reinstatable asymmetric complements to consumption’.” Or we can just learn a new vocabulary word— ‘cue’ which summarizes certain kinds of complements.”
4. Let us revisit the example of the American tourist in GP4, above. Camerer (2008) suggests, sensibly, that one can interpret the American tourist looking for cars in the UK as engaged in a *Stroop task* (see Part 7) and his mistake in looking left and crossing a road is a *Stroop mistake*.⁸ The fact that people get better at a given Stroop task with experience, i.e., commit fewer Stroop mistakes, suggests that it is not very probable that changes in the feasible set account for the mistakes of American tourists in the UK.

Schotter (2008) examines the Gul–Pesendorfer arguments from an experimental economist’s point of view. We summarize some of his arguments below.

In making their case for using choice data, Gul and Pesendorfer (2008, p. 8) argue that: “Standard economics focuses on revealed preference because economic data come in this form.” Schotter (pp. 73–74) does not dispute this claim or the importance of choice data, however, he argues that we should not make a virtue out of a constraint (that most economic data comes in this form). Indeed, all sorts of data that help us to test the predictions of economic models should be used.

Furthermore, it is worth asking if actual choice data materializes for the reasons assumed in neoclassical economics, say, the maximization of a well-defined utility function. What if choice data also arises from other considerations such as emotions, advice from friends, satisficing heuristics, simply following the herd, or any of a range of judgment heuristics described in Part 7. Choice data does not reflect the reasons for choice. So, in order to test alternative theories, say,

⁸ A Stroop task is an automatic, practiced, response by System 1 (such as looking left while crossing a road) that can be incorrect (e.g., in the UK, the American tourist must look right before crossing a road). System 2, the deliberative system, may or may not be able to rein in System 1 while it is making a Stroop mistake.

	Left	Right
Top	$x, 40$	$40, 80$
Bottom	$40, 80$	$80, 40$

Figure 22.1 A 2×2 matching pennies game.

Source: Goeree and Holt (2001) with permission from the American Economic Association.

based on emotions or psychological game theory, it might be unavoidable to use non-choice data, for instance, based on elicited beliefs.

In some cases, it might be important to discover why choices are made in order to test the predictions of the relevant theory. Consider the 2×2 matching pennies strategic form game shown in Figure 22.1 from Goeree and Holt (2001). First, suppose that $x = 80$ in the payoff matrix in Figure 22.1. In this case, there is a unique Nash equilibrium in mixed strategies in which each player randomizes equally between the two pure strategies. In their experiments, Goeree and Holt find that the players played each of their pure strategies with nearly equal probabilities, which is consistent with the Nash outcome.

Now let $x = 320$. This should leave unchanged the behavior of the row player because, in neoclassical theory, his choice of a probability distribution over his pure strategies is designed to make the opponent indifferent between his equilibrium pure strategies. Thus, in the mixed strategy equilibrium, corresponding to $x = 320$, the row player plays Top and Bottom with equal probabilities while the column player plays Left with probability $\frac{1}{8}$ and Right with probability $\frac{7}{8}$. However, in contrast to the predictions of classical game theory, row players now choose Top, 96% of the time.

When $x = 80$, the outcome is close to a Nash equilibrium, yet the choice of a Nash equilibrium is accidental in the symmetric game. A researcher who construed support for a mixed strategy Nash equilibrium based on this choice data would be misled because in the asymmetric game, the outcome differs significantly from a Nash equilibrium. It would appear that some economists view experimental data, even when subjects are properly incentivized, as non-choice data. In this strict view, only market data, say, price–quantity pairs, are admissible. Yet, if the market was like the world with $x = 80$, but almost never like the world with $x = 320$, it would not be able to refute the mixed strategy Nash equilibrium. Furthermore, this example provides a neat illustration of why experiments in economics are critical.

Another example of non-choice data is the use of *mouselab software* to infer whether the pattern of search and look ups in experiment games is consistent with backward induction; see Johnson et al. (2002) and the discussion in Part 4. The observed pattern of lookups appears to be inconsistent with backward induction. For instance, subjects search forward, not backward, and some do not even look at the information in the last stage. These results are valuable, despite the fact that economic theory makes no direct assumptions about the pattern of search and lookups. In a clever experimental design, Johnson et al. (2002) find that when subjects are trained in backward induction, they are more likely to search backwards, giving credence to their proposed method. This highlights why the experimental method must be critical to economics. Indeed, one might never be able to stringently test economic theory otherwise.

Survey data, another form of non-choice data, has provided a major impetus to progress in the rapidly expanding literature on happiness economics. This is discussed in more detail in Part 6

Table 22.1 Fraction of subjects who pay the fee in each treatment.

Situation	Fee (\$)	Probabilities		Fraction of subjects who pay fee	
		$p_l(A)$	$p_h(A)$	Treatment 1	Treatment 2
3	0.50	0.61	1	0.783	0.063
6	2.00	0.51	1	0.565	0.125
12	0.50	0.51	1	0.739	0.125

Source: Schotter (2008).

and other observations on survey data are made in the introductory chapter, so we omit a discussion here.

Actual choice data might be restricted to a subset of the parameter space, which might not allow for stringent tests of theory. We have exploited this observation above, but consider another insightful example by Schotter (2008) based on the work of Eliaz and Schotter (2010).

There are two boxes, A and B. A monetary prize is hidden in one of the boxes. The probability that the prize is hidden in box A in state $i = h, l$ is $p_i(A)$ such that $0.5 < p_l(A) < p_h(A)$. In treatment 1, there are two stages. In the first stage, the subject can pay a fee to learn the state, $i = h, l$; if no fee is paid, then the state is not learnt. In the second stage, subjects choose a box, A or B followed immediately by the payment of the prize if the correct box is chosen. In treatment 2, in the first stage, subjects choose a box, A or B. In the second stage they can pay a fee to learn the state of nature before being paid a prize if they choose the correct box.

Since $0.5 < p_l(A)$, choosing box A first order stochastically dominates choosing box B, irrespective of the state or the individual's attitudes towards risk. Thus, anyone who pays a fee has non-classical preferences. Two alternative explanations for the payment of a fee are considered. The first is the Kreps and Porteus (1978) framework that can accommodate the payment of a fee because subjects have a preference for an early resolution of uncertainty. The second explanation is the *confidence effect*; here people pay a fee just to feel happier from choosing a box that has a higher probability of a prize.

The two frameworks make different predictions. For instance, if as in Kreps–Porteus, individuals have a preference for early resolution of uncertainty, then such a preference should manifest in both treatments. However, if behavior is explained by the confidence effect, then it should have bite only in treatment 1 where subjects learn of the state before they actually make their choice of a box. The results are shown in Table 22.1 for three situations, which correspond to different combinations of the fee and the probabilities (the situation numbers, 3, 6, 12, correspond to those in Eliaz and Schotter, 2010).

Since the fraction of subjects who decide to pay the fee is so starkly different in the two treatments, the results reject the Kreps–Porteus model for the given data. Furthermore, a far higher fraction of subjects are willing to pay the fee in Treatment 1, which is consistent with the confidence effect. Suppose that the real-world choice data came only from Treatment 1, but never from Treatment 2, then choices might suggest support for the Kreps–Porteus framework. However, appropriate experimental design allows for stringent testing of theory, leading to the opposite conclusion.

Yet another reason for looking beyond choice data comes from framing effects. A pervasive finding is that choices are sensitive to framing, hence, without considering the frame, incorrect inferences are likely to be made. Consider, for instance, the well-known problem from Kahneman

and Tversky (1984) in which the same information is framed either in terms of lives saved or lives lost (see Part 7). The results demonstrated that individuals make choices consistent with risk averse behavior in the domain of gains and risk-loving behavior in the domain of losses. However, if the frame is not explicitly considered as a part of the description of the problem, one may make an incorrect inference about the risk preferences of individuals from the observed choices.

Kőszegi and Rabin (2008a) chiefly address the issue raised in GP4. They argue (p. 198) that introducing ad-hoc auxiliary assumptions in neoclassical models can produce *wacky theories* if one is not willing to entertain the possibility that individuals may make mistakes in their choices. This is particularly disturbing when more plausible, psychologically well-founded, alternatives are available. They propose a framework, developed slightly more formally in Kőszegi and Rabin (2008b), that recognizes mistakes in choices, yet attempts to recover the underlying preferences.

The supposition is that despite the propensity of people to make mistakes, there is a underlying degree of coherence to preferences. Thus, choices are not simply random and mistakes are systematic. They propose proceeding along the following lines.

1. Impose as few restrictions as possible on preferences.
2. Compare the subjective beliefs of individuals against objective probabilities. If the two tally then choices reflect welfare.
3. If subjective and objective beliefs do not tally then propose a theory in which individuals have well-defined underlying preferences but they make mistakes.

Let us consider the possibility that individuals suffer from the *gambler's fallacy* (see Part 7 for details). Suppose that an individual observes a series of coin tosses and q is her elicited belief that the next toss will produce tails (T) rather than heads (H).⁹ When no flips are observed, the individual chooses $q = 0.5$; when the sequence of consecutive flips $HHHHHH$ is observed, $q = 0.55$; when the sequence of consecutive flips $TTTTTT$ is observed, $q = 0.45$. One may, as in GP4, propose a neoclassical model in which the individual has a preference to bet on T after observing the sequence $HHHHHH$ and a preference for betting on H after observing the sequence $TTTTTT$; but this is a “wacky” set of preferences that is hard to justify. A more compelling argument is that the individual's subjective beliefs differ from the objective beliefs, and the individual makes a mistake that is consistent with the well-known gambler's fallacy.

Conditional on the assumption that the individual suffers from the gambler's fallacy, we give examples of inferences about preferences that we could draw from these choices.

Example 22.2 *Suppose that following any sequence of coin tosses, the individual is asked which of the two fruits, blueberries or strawberries, would he commit to strictly preferring, following the outcome of the next toss. When the sequence is: (i) HHHHHH, he reveals that he strictly prefers blueberries to strawberries if a T comes up on the next toss (and the converse preference if H comes up). (ii) TTTTTT, he reveals that he strictly prefers strawberries to blueberries if a T comes up on the next toss (and the converse preference if H comes up). (iii) the null sequence (no flips of the coin yet), he reveals that he is indifferent between the fruits, irrespective of the outcome on the next toss. Given that the individual exhibits the gambler's fallacy and has preferences that do not depend on coin flips, therefore, he prefers blueberries to strawberries.*

⁹ For the details of the proposed modification to the Becker–DeGroot–Marshak mechanism that may be used to produce these elicited beliefs, see Kőszegi and Rabin (2008a).

Example 22.3 Consider two individuals, classical (C) and behavioral (B). They choose one of two consumption bundles (x_h, x_t) or (y_h, y_t) . The first and the second elements in each bundle can be consumed if the next toss turns out to be, respectively, H and T. Individual C chooses (x_h, x_t) over (y_h, y_t) if

$$0.5u(x_h) + 0.5u(x_t) > 0.5u(y_h) + 0.5u(y_t). \quad (22.29)$$

Individual B chooses as follows. For the null history (no observed flips) B chooses in accordance with (22.29). If the history is HHHHHH, then (x_h, x_t) is chosen over (y_h, y_t) if

$$0.45u(x_h) + 0.55u(x_t) > 0.45u(y_h) + 0.55u(y_t). \quad (22.30)$$

Finally, if the history is TTTTTT, then B chooses (x_h, x_t) over (y_h, y_t) if

$$0.55u(x_h) + 0.45u(x_t) > 0.55u(y_h) + 0.45u(y_t). \quad (22.31)$$

The behavior of individual C is unproblematic for neoclassical economics. Following GP4, if one stays within the neoclassical framework, then one will need to invoke ad-hoc auxiliary assumptions to explain the choices of individual B, interpreted now as a classical individual. For instance, one could postulate that the individual has a preference for betting on changes in coin flips. However, because the classical individual does know and use objective probabilities (see (22.29)), such wacky preferences will lose money for the individual.

A simpler explanation is that individual B is not a classical individual and suffers from the gambler's fallacy. Indeed, we can identify the preferences of B from the choice data. Such an individual is an expected utility maximizer who suffers from the gambler's fallacy. In an apparent response to the Gul–Pesendorfer framework, Köszegi and Rabin (2008a) write (p. 201):¹⁰ “We admit to not really understanding why we would at all be inclined to ban the study of B from economics departments, but we are especially chagrined at that prospect in light of the fact that we can use the same powerful tools of economics to study B as C in this case. B has well-ordered and coherent preferences. She is making an error in statistical reasoning. The two can be jointly identified. It is useful to do so.”

Example 22.4 There is one caveat to the explanation in Example 22.3. Choice behavior may indicate a mistake. However, there could be more than one competing explanation for the mistake, and the welfare implications are dependent on the chosen explanation. Consider an experiment in which pregnant women are asked, prior to going into labor, if they would ask for an epidural injection during childbirth. Suppose that most women choose natural childbirth, yet when they go into labor, they change their mind and ask for an epidural injection.

One potential explanation of this empirically observed phenomenon is projection-bias (see Part 6) in which pregnant women underappreciate how painful their labor will be. A second explanation is hyperbolic discounting in which the present bias at the time of labor tips the scales in favor of the epidural. There is no role for commitment under projection bias. Indeed, the decision to take an epidural at the time of labor is the welfare maximizing choice for the individual. However, if the source of the mistake is hyperbolic discounting, then the individual might find it optimal to invest in a commitment technology to bind herself to the original, ex-ante, choice. In this case, one will need extra information to determine which of the two explanations is the correct one.

¹⁰ In the quote we have replaced Fiona with individual B and Giles with individual C from the original quote.

22.6 Choice-based behavioral welfare economics

In this section, we consider the choice-based measure of welfare of Bernheim and Rangel (2007, 2009); the basic rationale is discussed in Section 22.3.6.

In neoclassical economics, one may proceed with standard welfare analysis in one of two possible ways. One can either specify preferences and then derive a *choice function* from it. Alternatively, one can directly specify a choice function and then derive the corresponding preferences. The two methods are equivalent; for a lucid account see Chapter 2 in Kreps (1990). Bernheim and Rangel choose the second method, i.e., begin with a choice function and then derive implications for preferences.

Suppose that X is the set of all possible choice objects. Choice objects are not restricted, so these could, for instance, be consumption bundles, or lotteries of dated profiles of intertemporal consumption. Denote the constraint set by $X \subseteq \mathbf{X}$. The general nature of the choice set allows this framework to address static and dynamic problems under risk. An ancillary condition, A , is any set of environment conditions such as default options, status quo, anchors or other framing conditions that are not relevant to the evaluation of an individual's welfare by a social planner.

Define a *generalized choice situation* (GCS) by $g = (X, A)$. Suppose that the individual must choose from n different GCSs: $(X_1, A_1), \dots, (X_n, A_n)$. These can be collapsed into a single GCS, $g' = (X', A')$ where $X' = \bigcup_{j=1}^n X_j$ and A' specifies how the choice is to be made from the n different GCSs. Denote by G^* , the universal set of all GCSs allowed by a theory; thus, G^* is theory-specific. For instance, anchors or defaults might play no role in some theories but play a central role in others. Let $G \subseteq G^*$ be a trimmed-down universal set that contains all the *welfare relevant* GCSs; G is the *welfare relevant domain* for normative analysis in which choices reveal “true objectives” of the decision maker. We discuss the issue of specifying G later.

Next, we define a choice correspondence that maps any GCS into a set of choice objects.

Definition 22.4 (*Choice correspondence*): A choice correspondence is defined as $C : G^* \rightarrow \mathbf{X}$ and $C(g) \subseteq X$ for all $g = (X, A) \in G^*$.

We now define the set of all constraint sets, X , such that for some ancillary condition, A , the resulting GCS, (X, A) is in the welfare relevant domain, i.e.,

$$\chi = \{X : \exists A \text{ such that } (X, A) \in G\}. \quad (22.32)$$

The following two assumptions are needed to make further progress.

A1: Every nonempty finite subset of $\mathbf{X}$ is a member of the set χ , defined in (22.32).

A2: $C(g)$ (see Definition 22.4) is nonempty for all $g \in G$.

The assumptions made so far allow for intransitive choices, and choices that are influenced by irrelevant alternatives or context. The constraint sets are not required to be compact. In conjunction with the general nature of choice objects, this framework allows a very wide range of questions to be asked.

Example 22.5 (*The cafeteria problem; Sunstein and Thaler, 2003; Thaler and Sunstein, 2003*): Consider the cafeteria example from Section 22.3.1. Suppose that there are only two choice objects: healthy snacks, h , and unhealthy snacks, u , so $\mathbf{X} = \{h, u\}$ and $X = \{h, u\}$; from

assumption A2, choices from singleton sets are trivial, hence, we suppress them. The ancillary condition specifies which object is placed at eye level where it is more salient. Let the ancillary conditions A_1 and A_2 denote, respectively, healthy and unhealthy snacks placed at eye level. Let $g_1 = (\{h, u\}, A_1)$, $g_2 = (\{h, u\}, A_2)$, then the universal set of all GCSs is (also recall that singleton choice sets are suppressed)

$$G^* = \{g_1, g_2\}.$$

Suppose that the welfare relevant domain is defined to be the one where the healthy snacks are always placed at eye level, so

$$G = \{g_1\}.$$

Using (22.32),

$$\chi = \{\{h, u\}, \{h\}, \{u\}\}.$$

Suppose that the choice correspondence is defined by

$$C(g_1) = \{h\}, C(g_2) = \{u\}. \quad (22.33)$$

The non-empty subsets of X are $\{h, u\}$, $\{h\}$ and $\{u\}$ and all three sets are in χ so assumption A1 is satisfied. Furthermore, from (22.33), assumption A2 is also satisfied (note that we have suppressed singleton sets in (22.33)).

Example 22.6 (Time inconsistency): Consider an example motivated by time inconsistency of preferences, say, due to hyperbolic discounting. Suppose that $X = \{x, y\}$, where x and y are time t payoffs. Two ancillary conditions, A_1 , A_2 , respectively, denote choices made at time $t - 1$ and time t , respectively. Thus, $g_1 = (\{x, y\}, A_1)$, $g_2 = (\{x, y\}, A_2)$ and $G^* = \{g_1, g_2\}$. The generalized choice problem g_1 specifies the choice set from which the choice is made at time $t - 1$; g_1 plays the corresponding role at time t . Suppose that the choices of the individual are

$$C(g_1) = \{x\}, C(g_2) = \{y\}. \quad (22.34)$$

Thus, the decision maker reveals preference reversals in dated choices. Since choice-based measures can only rely on the choices made, thus, the choices in (22.34) provide conflicting guidance to the planner. One way out for the planner is to treat the welfare relevant domain either as $G = \{g_1\}$ or $G = \{g_2\}$. In this case, inconsistent choices are trimmed away. But which of the two welfare relevant domains should the planner choose?

This aspect of the framework has been criticized by Sugden (2008) and Loewenstein and Haisley (2008). Clearly, one is then forced to rely on non-choice data to discover the welfare relevant domain. But if some welfare criteria has to be employed by the planner to resolve this problem, then we are back to square one. One suggestion by Bernheim and Rangel (2009) is to use neuroeconomic data to see which of the two choices really reflects the individual's welfare. Whether advances in neuroeconomics will allow us to do so is an open question.

Let us now develop the analogs of the neoclassical weak and strong choice relations. We restrict ourselves to the welfare relevant domain, G , ignoring questions about how it might have been discovered, or whether it could ever be discovered.

Definition 22.5 (Weak unambiguous choice): Let $x, y \in X$ and let G be the welfare relevant domain.

(a) We say that x is “weakly unambiguously chosen” over y , written xRy , if

For all $(X, A) \in G$ and $x, y \in X$, if $y \in C(X, A)$ then $x \in C(X, A)$.

(b) Let P denote the asymmetric component of R , so

$$xPy \Leftrightarrow xRy \text{ but not } yRx.$$

(c) Let I denote the symmetric component of R , so

$$xIy \Leftrightarrow xRy \text{ and } yRx.$$

This means that whenever x is chosen, y is also chosen and vice versa.

Example 22.7 Let $X_1 = \{x, y\}$, $X_2 = \{y, z\}$, $X_3 = \{x, z\}$, $X_4 = \{x, y, z\}$, $X_5 = \{x\}$, $X_6 = \{y\}$, $X_7 = \{z\}$ and $\chi = \{X_1, \dots, X_7\}$. We ignore the singleton sets X_5, X_6, X_7 because the choice from these sets is trivial, given assumption A2. Suppose that we observe the following choices for all ancillary conditions, A_i , $i = 1, 2, \dots, n$

$$C(X_1, A_i) = \{x\}, C(X_2, A_i) = \{y\}, C(X_3, A_i) = \{z\}, C(X_4, A_i) = \{x, y\} \quad (22.35)$$

Given (22.35), we can conclude that xRy . But yRx does not hold because in the set X_1 , x is chosen but y is not. Hence, xPy .

Definition 22.6 (Strict unambiguous choice): Let $x, y \in X$ and let G be the welfare relevant domain. We say that x is “strictly unambiguously chosen” over y , written xP^*y , if

For all $(X, A) \in G$ and $x, y \in X$, $y \notin C(X, A)$.

Since the choice set is non-empty (assumption A2), thus, when $X = \{x, y\}$ then x is chosen but y is not. In general, whenever x and y are available in any constraint set then y is never chosen.

Example 22.8 (Example 22.7 continued . . .): From (22.35), we cannot conclude that xP^*y because in the set X_4 , both x, y are available, yet y is also chosen. If instead $C(X_4, A_i)$ was $\{x\}$, or $\{z\}$, or $\{x, z\}$, then we have xP^*y . Suppose that $C(X_4, A_i) = \{x\}$, while the choices in all other cases are given by (22.35). In this case, we have xPy and xP^*y , thus, P and P^* coincide

Definitions 22.5 and 22.6 get around the problem of human choice that is coherent but not well behaved in the sense of being intransitive, or influenced by irrelevant contexts or alternatives. However, they do not get around the problem where decision makers appear to act against their own interest such as not making full use of 401(k) plans, addictions, and low aspirations; see Dalton and Ghosal (2016) for some progress in this direction.

Whenever $C(X, A)$ is unique for each $(X, A) \in G$ then P (Definition 22.5) and P^* (Definition 22.6) coincide (see Example 22.8). Otherwise they can differ. For several binary relations, P^* is the asymmetric part, as, for instance, is the case in the next definition.

Definition 22.7 For all $(X, A) \in G$ and for all $x, y \in X$ the binary preference relation, R^* , on G is defined as follows

$$xR^*y \Leftrightarrow \text{not } yP^*x,$$

i.e., y is never strictly unambiguously chosen over x . So there exists some GCS such that x and y are available and x is chosen. Similarly, one may define the indifference relation I^* as follows

$$xI^*y \Leftrightarrow xR^*y \text{ and } yR^*x,$$

i.e., there is some GCS for which x is chosen when y is available and another for which y is chosen when x is available.

In the absence of ancillary conditions, the binary relations (1) R and R^* are equivalent to the standard weak revealed preference relation in neoclassical economics, and (2) P and P^* are equivalent to the standard strong revealed preference relation in neoclassical economics. Thus, the neoclassical case is subsumed within this framework.

Proposition 22.3 (Bernheim and Rangel, 2009): For all $(X, A) \in G$ and for all $x, y \in X$, we have

$$xP^*y \Rightarrow xPy \Rightarrow xRy \Rightarrow xR^*y,$$

so P^* is the coarsest and R^* the finest of these binary relations.

Example 22.9 (Intransitivity of P^* and incompleteness of R): Let $X_1 = \{a, b\}$, $X_2 = \{b, c\}$, $X_3 = \{a, c\}$, $X_4 = \{a, b, c\}$, $X_5 = \{a\}$, $X_6 = \{b\}$, $X_7 = \{c\}$, and $\chi = \{X_1, \dots, X_7\}$. There are no ancillary conditions, so the welfare relevant domain comprises the sets $X_1, \dots, X_7$. The choices are given by (choices are trivial for singleton sets so these are not listed):

$$C(X_1) = \{a\}, C(X_2) = \{b\}, C(X_3) = \{c\}, C(X_4) = \{a\}. \quad (22.36)$$

In choice sets X_1 and X_4 , the options a and b are available but b is never chosen, while a is chosen from both choice sets, so aP^*b . In choice sets X_2 and X_4 , options b and c are available yet c is never chosen and b is chosen once, so bP^*c . In contrast, consider choice sets X_3 and X_4 where options a and c are present. Option c is chosen from X_3 and option a is chosen from X_4 . Thus, we have both cases $\sim aP^*b$ and $\sim bP^*a$, so using Definition 22.7, we have aI^*b . Thus, P^* is intransitive. From (22.36) we neither have aRc nor cRa so R is incomplete.

Example 22.10 (Incompleteness of R when the ancillary condition varies): Consider two GCSs (X, A) , $(X, A') \in G$ where $X = \{x, y\}$. Suppose that we have the following choices

$$C(\{x, y\}, A) = \{x\} \text{ and } C(\{x, y\}, A') = \{y\}.$$

The second choice prevents xRy and the first choice prevents yRx . Thus, R is not complete in this case. In contrast, R^* is always complete.

While Example 22.9 shows that P^* is intransitive, none of the binary relations defined so far need to be transitive. The next example illustrates intransitivity of P ; an exercise asks you to show that R^* is intransitive.

Example 22.11 (Intransitivity of P and R): Consider the setup in Example 22.9 but replace (22.36) by the following choices

$$C(X_1) = \{a\}, C(X_2) = \{b\}, C(X_3) = \{c\}, C(X_4) = \{a, b, c\}.$$

We have aRb but not bRa , so using Definition 22.5 we have aPb . We can analogously check that bPc . But we do not have aRc so it is not the case that aPc . Thus P is intransitive. R is also intransitive because aRb and bRc but not aRc .

Intransitivity does not prevent a welfare analysis. If a preference relation on a finite set is acyclic then a maximal element exists. The binary relation P^* is acyclic, as shown in the next proposition, however, P need not be acyclic.

Proposition 22.4 (Bernheim and Rangel, 2009):

- (a) Consider outcomes $x_1, \dots, x_N$ such that $x_1Rx_2, \dots, x_{N-1}Rx_N$ and for some $1 \leq k \leq N-1$ we have $x_kP^*x_{k+1}$. Then $\sim x_NRx_1$.
- (b) P^* is acyclic, so for outcomes $x_1, \dots, x_N$ such that $x_1P^*x_2, \dots, x_{N-1}P^*x_N$ we have $\sim x_NP^*x_1$.

Next we give the welfare criteria for improving choices.

Definition 22.8 (Welfare improvement):

- (a) (Strict welfare improvement): Given any choice $x \in X$ we can “strictly improve” upon the choice if there exists some $y \in X$ such that yP^*x . When a strict improvement is impossible then x is a weak individual welfare optimum.
- (b) (Weak welfare improvement):
Given any choice $x \in X$ we can “weakly improve” upon the choice if there exists some $y \in X$ such that yPx . When a weak improvement is impossible then x is a strict individual welfare optimum.

The following proposition is an implication of Definition 22.8.

Proposition 22.5

- (a) If $x \in C(X, A)$ for any $(X, A) \in G$ then it is a “weak individual welfare optimum” in X . If $x \in C(X, A)$ is unique then it is a strict welfare optimum.
- (b) x is a weak welfare optimum in X if and only if for each $y \in X$ such that $y \neq x$ there is some GCS (X', A') in which x is chosen when y is available (although y might be chosen as well).

Proof: (a) Since $x \in C(X, A)$ there does not exist any $y \in X$ such that yP^*x . Hence, it is not possible to strictly improve on x , so x is a “weak individual welfare optimum” in X . If x is unique then there does not exist any element y such that yPx . Thus, a weak improvement is impossible, which implies that x is a strict welfare optimum.

(b) Since there exists some GCS (X', A') in which x is chosen when y is available, we cannot have yP^*x . Thus, we cannot strictly improve welfare starting from x , so x is a weak welfare optimum. ■

Choice sets are non-empty (assumption A2), thus there always exist some $x \in C(X, A)$. Hence, from Proposition 22.5(a), a weak individual welfare optimum always exists within any X such that $(X, A) \in G$. This is the sense in which the welfare criteria respects individual liberty because, by definition, x is voluntarily chosen. Thus, within any welfare relevant domain, G , it is not possible

to improve upon individual decisions by a third party. This is not the case when we choose the universal domain of GCSs, G^* . But this is precisely the role of the choice architect under soft paternalism: to trim down the universal domain of GCSs to the welfare relevant domain. Indeed, no assumptions such as continuity and compactness are needed for this result because of our assumption of well-behaved choice correspondences.

Example 22.12 *Let us revisit Example 22.9. The choices are given in (22.36), which give rise to an intransitive P^* . In the set X_4 , a is the unique choice element, so it is a strict welfare optimum (Proposition 22.5a) in X_4 while b and c are not even weak welfare optima in X_4 . Also using Proposition 22.5a, in set X_1 , b is not a weak welfare optimum and in set X_2 , c is not a weak welfare optimum.*

In set X_3 , where the unique choice is c we cannot conclude that it is the only strict welfare optimum. The reason is that since $C(X_4) = \{a\}$ when c is available, so we cannot compare a and c under P^* . Indeed we cannot affect a weak welfare improvement starting from a or c because neither aPc (c is chosen in X_3 when a is available) nor cPa (a is chosen in X_4 when c is available). Since weak improvements are not possible, thus, both a and c are strict welfare optima in the set X_3 although only c is chosen.

Different ancillary conditions may produce conflicting choices, so much so that it becomes difficult or impossible to identify weak or strong unambiguous welfare improvements. Exercise 7 illustrates this problem. On the other hand, the framework potentially solves some empirical puzzles as the next example shows.

Example 22.13 *(Iyengar and Lepper, 2000; Bernheim and Rangel, 2009) Suppose that the potential options are 29 different brands of jams; the 30th option is “choose nothing.” The welfare relevant domain contains all pairs of 30 options and the entire set of 30 options but no other combinations of options. There are no ancillary conditions. In any pairwise comparison, strawberry jam (say option x) is preferred to any other jam. However, when asked to choose among all 30 alternatives in one go, the individual chooses the 30th option, i.e., choose nothing (say, option y). In the Bernheim–Rangel framework this is not problematic. Let us start with x . We cannot find any other option, say, z such that zP^*x . Thus, a strict improvement over x is impossible and it is a weak individual welfare optimum. x is not chosen when all 30 options are presented but option y is chosen. Yet, it is not the case that yP^*x . Nor is it the case that xP^*y (although xP^*z and yP^*z for all other options z). Thus, both x and y are weak individual welfare optimum.*

In Example 22.13, if one restricted the welfare relevant domain to only all pairwise comparisons of the options, but do not allow all options to be presented together, then we do get a unique choice, x . However truncating in this manner can lead to cyclicity of P^* as shown in the next example.

Example 22.14 *Consider the setup of Example 22.9 but restrict the welfare relevant domain to $X_1 = \{a, b\}$, $X_2 = \{b, c\}$, $X_3 = \{a, c\}$, leaving out $X_4 = \{a, b, c\}$. This violates assumption A1 because $G = \{X_1, X_2, X_3\}$ does not contain all finite subsets of the options a , b , and c . Hence, Proposition 22.4 need not apply. Indeed this is the case because P^* cycles: aP^*b (whenever a and b are present in a set, b is never chosen); bP^*c (whenever b and c are present, c is never chosen); cP^*a (whenever a and c are present, a is never chosen). This intransitivity of P^* makes it difficult to figure out what the choice would be (normative or positive) in the set X_4 ; each of a , b , and c is improvable.*

We can relate the binary choice relations P and P^* to traditional measures of welfare improvement such as Pareto optimality in the following *multi-selves interpretation*. Suppose that the welfare relevant domain G is *rectangular*, i.e., $G = \chi \times \Gamma$ where Γ is the set of all ancillary conditions. Suppose that each ancillary condition corresponds to a different self of the same individual and for each ancillary condition $A \in \Gamma$ the individual maximizes a utility function u_A ; thus, preferences are well behaved for each ancillary condition. We now apply the Pareto criteria across the different selves of the individual.

Definition 22.9 (*Multiself Pareto dominance*): Let $y \in X$ for some X in χ .

- (a) x “weakly multiself Pareto dominates” y , written xMy , if $u_A(x) \geq u_A(y)$ for all $A \in \Gamma$.
- (b) x “strongly multiself Pareto dominates” y , written xM^*y , if $u_A(x) > u_A(y)$ for all $A \in \Gamma$.
- (c) $x \in X$ is a weak multiself Pareto optimum if there is no $y \in X$ such that yM^*x .
- (d) $x \in X$ is a strong multiself Pareto optimum if there is no $y \in X$ such that yMx .

The criteria proposed in Definition 22.9 are quite strong. For instance, a typical behavioral finding is that the choice made by an individual often changes with framing effects, which corresponds here to changes in ancillary conditions. However, the multiself Pareto criteria requires that we do not get reversal of rankings when framing changes.

The next proposition establishes the link between multiself Pareto dominance and the relations P and P^* .

Proposition 22.6 (*Bernheim and Rangel, 2009*): Suppose that G is rectangular, and each $A \in \Gamma$ activates a different multiself of the same individual. Assume that for each $A \in \Gamma$, preferences are well behaved, so we can imagine that the self corresponding to $A \in \Gamma$ maximizes a utility function u_A . Then $M^* = P^*$ and $M = P$.

Proposition 22.6 implies that saying $x \in X$ is a weak (strong) multiself Pareto optimum is equivalent to saying that $x \in X$ is a weak (strong) individual welfare optimum.

The main difficulty in operationalizing this choice-based proposal is to specify the welfare relevant domain and to deal with cases where the welfare criterion is not very helpful in narrowing down choice or choice is cyclic (as in Example 22.14). Indeed, one might need to invoke another normative criterion in this case, or use non-choice data such as data from neuroeconomics. Bernheim and Rangel (2009) make other suggestions too. For instance, one might appeal to other principles such as *preponderance* (p. 89); thus, if a decision maker typically chooses x over y but rarely y over x then the rare exceptions should be ignored.

These comments illustrate two things. First, not everyone will find the choice-based approach to be utterly persuasive. Second, there are severe difficulties in making progress with a choice-based approach even when the choice-based approach is stretched to its limits. Indeed, the choice-based approach has been criticized on these grounds and alternatives have been proposed. For instance, Sugden (2013) proposes a contractarian approach (see Section 22.8 below) and Loewenstein and Haisley (2008) advocate welfare criteria based on preferences rather than choice. Cason and Plott (2014) show that subject misconceptions in choice problems provide poor guidance to their underlying preferences, even in relatively simple problems. Choi et al. (2014) find considerable heterogeneity among individuals in their decision making quality/ability and consistency with utility maximization. Relatively poorer decision making quality associated with lower wealth and lower educational attainment, produces choices that have lower consistency with utility maximization.

For an extension of the Bernheim–Rangel framework to an interpersonally comparable concept of well-being and the development of criteria for policy evaluation, see Fleurbaey and Schokkaert (2013). For a survey of alternative measures of GDP in the presence of interpersonal comparisons, see Fleurbaey and Blanchet (2013). Rubinstein and Salant (2008) have independently developed a choice-based framework in which the ancillary conditions in the Bernheim–Rangel framework are called *frames*. They also develop binary relations similar to P and P^* . Manzini and Mariotti (2014) give a useful discussion of the relative merits of the model-free approach, e.g., the Bernheim–Rangel framework, and the model-based approach, e.g., Masatlioglu et al. (2012) (see Section 22.7 below).

Rubinstein and Salant (2012) depart from the Bernheim–Rangel framework by giving up the choice between alternatives based on the Pareto criterion. Their argument is that as the behavioral data grows, the ability of the Pareto criterion to rank alternatives, shrinks. They propose a welfare ranking of any two alternatives even when one alternative does not Pareto dominate the other alternative. Rather than postulate that inconsistent behavior arises by changes in, say, frames, their model supposes that individuals have more than one preference ordering. Inconsistent choices arise because the individual switches between the preference orderings while making different choices.

22.7 Revealed preference under limited attention

So far, we have assumed that individuals pay attention to all alternatives. In contrast, the evidence supports *limited attention* to alternatives in many domains (see Part 7). In this section, we consider the welfare implications of limited attention using the model of Masatlioglu et al. (2012), who also give the following useful taxonomy of limited attention, to which their framework can be applied.

1. *Choose the top N* : Examples include choosing from among the N cheapest suppliers and ignoring the rest; considering only the N most advertised products; and considering only the products advertised on the first page of results from a search engine.
2. *Top on each criteria*: One may sequentially choose among criteria (Manzini and Mariotti, 2007, 2012). For instance, a potential buyer of a car may narrow down his choices by first choosing the safest car, then choosing the most fuel efficient car from the narrower set, and then the cheapest car from the even narrower set.
3. *Most popular category*: A potential buyer of a bike might first check online the shop that offers the largest variety of bikes and then visit it to buy the bike, ignoring all other bike shops.

The proposed framework has three stages. In the first stage, the decision maker chooses to narrow down the set of alternatives to those that he desires to pay attention to, or he may simply be unaware of the existence of some alternatives. In the second stage, he has a complete and transitive preference ordering over the remaining alternatives. In the third stage, we examine the welfare implications of the first two stages.

The Bernheim–Rangel approach described in Section 22.6 is model-free because the choice function is the primitive and any choice procedure that generates the stated choices is consistent with the model. In contrast, the approach under limited attention specifies a particular choice procedure that relies on limited attention, so it is not model-free. In the Bernheim–Rangel approach, x is unambiguously preferred over y when y is never chosen if x is available (Definition 22.6). Under limited attention, by contrast, the decision maker might sometimes not pay

attention to x , resulting in a choice of y . We shall see an example below where the two criteria are in conflict.

We now develop the formal model. Suppose that X is a finite set of alternatives and χ is the set of all non-empty subsets (or feasible sets) of X . Suppose that the decision maker faces the set $S \in \chi$ of alternatives. The decision maker gives attention to a non-empty subset $\Gamma(S)$ of elements in S , known as the *consideration set*.

Definition 22.10 (*Consideration set mapping*): $\Gamma : \chi \rightarrow \chi$ is a consideration set mapping such that for all $S \in \chi$, $\Gamma(S) \neq \emptyset$ and $\Gamma(S) \subset S$.

If the decision maker is in the market for a new car, then S is the set of all models of all cars sold in the country where he lives. Suppose that the decision maker pays attention to a subset $\Gamma(S)$ of the cars. Consider a model of car $x \notin \Gamma(S)$ but $x \in S$. It is quite possible that the decision maker is unaware of the existence of all models in the set $S \setminus \Gamma(S)$. If now the decision maker were to face the set $S \setminus x$ then it is reasonable to assume that his consideration set would not change, so that $\Gamma(S) = \Gamma(S \setminus x)$. In other words, the decision maker is unaware that he is unaware of all the alternatives. This motivates the next definition

Definition 22.11 (*Attention filter*): A consideration set mapping is an attention filter if for any S , if $x \notin \Gamma(S)$ but $x \in S$ then $\Gamma(S) = \Gamma(S \setminus x)$.

We shall assume that the underlying preferences, $\succ$, on X are complete and transitive and the decision maker chooses a unique best element in each consideration set. We denote $\succ$ by P .

Definition 22.12 (*Choice function*): A choice function $c : \chi \rightarrow \chi$ assigns to each element of the feasible set, $S \in \chi$, a unique element $c(S) \in S$.

Definition 22.13 (*Choice with limited attention, CLA*): A choice function c is a choice with limited attention (CLA) if $\exists$ a complete transitive preference $\succ$ over X and an attention filter Γ such that $c(S)$ is the $\succ$ best element in $\Gamma(S)$.

The preferences of the decision maker $\succ$ and the attention filter used by the decision maker are not observable to a third party. Yet, we show below, that from the observed choices we can infer preferences and the attention filter used.

Example 22.15 Let $X = \{x, y, z\}$, so $\chi = \{\{x, y, z\}, \{x, y\}, \{y, z\}, \{x, z\}\}$; we omit the singleton sets that have a trivial choice. Suppose that the choice function is given by

$$c(x, y, z) = x, c(x, y) = x, c(y, z) = y, c(x, z) = z. \quad (22.37)$$

Two possible preferences, $\succ_1$ and $\succ_2$, along with the corresponding attention filters, Γ_1 and Γ_2 , are given in Figure 22.2. The reader is invited to check that Γ_1 and Γ_2 satisfy Definition

Preference		Attention filter			
		$\{x, y, z\}$	$\{x, y\}$	$\{y, z\}$	$\{x, z\}$
$z \succ_1 x \succ_1 y$	Γ_1	xy	xy	y	xz
$x \succ_2 y \succ_2 z$	Γ_2	xyz	xy	yz	z

Figure 22.2 Two possible preferences, $\succ_1$ and $\succ_2$, along with the corresponding attention filters, Γ_1 and Γ_2 .

Source: Masatlioglu et al. (2012) with permission from the American Economic Association.

22.11. Γ_2 gives full attention to all elements of all subsets of χ except $\{x, z\}$ where it ignores the element x . Γ_1 gives full attention to the sets $\{x, y\}$ and $\{x, z\}$ but limited attention to the sets $\{x, y, z\}$ and $\{y, z\}$.

The reader can easily confirm that given the attention filters, Γ_1 and Γ_2 , the preferences, $z \succ_1 x \succ_1 y$ and $x \succ_2 y \succ_2 z$, are entirely consistent with the choices given in (22.37). Yet we cannot observe either the attention filters or the preferences of the individual. However, suppose that we could discover somehow that $\succ_1$ and $\succ_2$ are the only two pairs of preferences that the decision maker could have. Then, it can be shown that the individual prefers x to y on the following grounds. The decision maker gives attention to both x and y under both attention filters, whenever x and y are available in any set in χ . However, x is always chosen when both x and y are available (see (22.37)).

The observations in Example 22.15 give rise to the next definition, which is based on a fairly conservative criterion of unanimity among all pairs of preferences and attention filters. Requiring unanimity has the advantage that the criterion is not controversial, but the disadvantage is that it limits the domain of applicability of the criterion.

Definition 22.14 Suppose that c is a CLA (Definition 22.13) and there are k different pairs of preferences and attention filters that can represent the stated choices, given by $(\Gamma_1, \succ_1)$, $(\Gamma_2, \succ_2), \dots, (\Gamma_k, \succ_k)$.¹¹ Then

- (a) x is revealed preferred to y if $x \succ_i y$ for $i = 1, \dots, k$.
- (b) x is revealed to attract attention at S if $x \in \Gamma_i(S)$ for $i = 1, \dots, k$.
- (c) x is revealed not to attract attention at S if $x \notin \Gamma_i(S)$ for $i = 1, \dots, k$.

Determining revealed preference and revealed attention/inattention can get cumbersome for large k . However, there is a simpler method that we turn to.

Consider Example 22.15 and compare the two sets $S = \{x, y, z\}$ and $S' = \{x, z\}$. From (22.37), the choices given by $c(x, y, z) = x$, and $c(x, z) = z$ reveal a choice reversal. If the decision maker paid attention to x and z in both sets, then we should not expect a choice reversal. A plausible suggestion is that the removal of y from S alters the attention set; thus, the decision maker must have paid attention to y in the set S (revealed attention). If this reasoning is accepted, then the presence of y induces the decision maker to pay attention to x in the set S and the absence of y from the set S' draws attention away from x , which explains the choice $c(x, z) = z$. Since the decision maker pays attention to x and y in the set S yet chooses x we say that x is revealed preferred to y .

More generally, if the choice changes upon removing an alternative, then the originally chosen alternative is preferred to the alternative that is removed. We can formally state the general idea as follows.

Definition 22.15 Let $x, y \in X$ be two distinct elements. Then xPy (x is revealed preferred to y) if there exists a set $S \in \chi$ such that $c(S) = x$ but $x \notin c(S \setminus y)$.

The preference relation $\succ$ on X is transitive. Thus, we if discover that xPy and yPz , then it follows that xPz although x is not directly revealed preferred from the individual's choices. Denote by P_R , the transitive closure of P . The transitive closure of a binary relation P on a set X is the minimal transitive relation P_R on X that contains P . Thus, $x_1 P_R x_n$ if there exist $x_1, \dots, x_n \in X$ and we have $x_j P x_{j+1}$ for $1 \leq j \leq n - 1$. The next proposition directly follows from these observations.

¹¹ For instance, in Example 22.15, $k = 2$.

Proposition 22.7 (Revealed preference; Masatlioglu et al., 2012): Suppose that c is a CLA. For $x, y \in X$, x is revealed preferred to y if and only if xP_{RY} .

In the next example we see how the predictions of this model differ from the Bernheim–Rangel framework.

Example 22.16 Consider a supermarket that displays four products x, y, z, t , each packed in a box. All four products need not be available at all times. The products can be displayed in two aisles, using the following five rules:

- (1) Each aisle can carry no more than two products because the supermarket manager has determined that it would otherwise overload customers' cognitive abilities. Furthermore, none of the aisles contains only one product unless the supermarket runs out of all remaining products.
- (2) x and y come packed in big boxes, which cannot be stored in the same aisle.
- (3) Since the first aisle is more visible to customers, the supermarket fills it first before any products are placed in the second aisle.
- (4) Due to an exclusive agreement with the manufacturers of products y and z , these products are placed in the first aisle whenever they are available.
- (5) The firm deems product t to be the least popular, so it is placed in the first aisle only if there is space remaining in the aisle after displaying the remaining available products.

Suppose that a customer has privately known preferences

$$t \succ x \succ z \succ y. \quad (22.38)$$

His choice is determined as follows: He visits the first aisle and picks his most preferred item; due to limited attention he never visits the second aisle. Let us separately consider the inferences drawn in the Bernheim–Rangel approach and the CLA approach.

I. Bernheim–Rangel approach: Suppose that products x and y are available. From rules 2 and 4, if y is available, then x is never placed in the first aisle. Suppose that x is placed in the second aisle. Since the customer only explores the first aisle due to limited attention, hence, x never gets chosen by the customer. In the terminology of Bernheim–Rangel, yP^*x , yet the individual prefers x to y , i.e., $x \succ y$. Thus, under limited attention, the Bernheim–Rangel criteria is not suitable; indeed it was never designed for limited attention.

II. The CLA approach: The five rules, given above, suggest that the consideration sets are given by

$$\Gamma(xyzt) = \{y, z\}, \Gamma(xzt) = \{x, z\}, \Gamma(xyt) = \{y, t\}, \Gamma(yzt) = \{y, z\}, \Gamma(xyz) = \{y, z\},$$

and when any two products remain we have $\Gamma(ab) = \{a, b\}$ for all $a, b \in \{x, y, z, t\}$. Suppose all four products are available. Then, by rules 2 and 4, y and z are placed in the first aisle and x and t in the second aisle. The individual picks z from the first aisle because $z \succ y$. Once y is unavailable then x is moved to the first aisle; since $x \succ z$, x is now chosen. Suppose that z is sold out. Then t is placed in the first aisle. Given $t \succ x$, the customer then chooses t . In summary the choices are

$$c(xyzt) = z, c(xzt) = x, c(xt) = t.$$

The choice data, if available, correctly identifies the underlying preferences in (22.38).

If removing an option from a set causes a choice reversal then the decision maker must have been paying attention to the removed object (Definition 22.13). It is also possible to determine attention even if no choice reversals take place, as shown in the following example.

Example 22.17 Suppose that a decision maker chooses y from sets S and T . Removing x from T causes a choice reversal, so the decision maker must have been paying attention to x , i.e., $x \in \Gamma(T)$. Suppose that we now consider elements in the set $A = S \cup T \setminus S \cap T$ and suppose that all of these are revealed preferred to y . Then it must be the case that elements in the set A were not under consideration when y was chosen, i.e., did not belong to $\Gamma(S)$ and $\Gamma(T)$. But then removing these elements from S and T will not change the consideration set, so

$$\Gamma(S) = \Gamma(S \cap T) = \Gamma(T). \quad (22.39)$$

Since $x \in \Gamma(T)$, (22.39) implies that $x \in \Gamma(S)$. Thus, we can infer which elements belong to $\Gamma(S)$ without actually looking for choice reversals from the set S . The next proposition formalizes these ideas about revealed inattention (recall that P_R is the transitive closure of P).

Proposition 22.8 (Revealed inattention; Masatlioglu et al., 2012) Suppose that c is a CLA, then

- (a) x is revealed not to attract attention at a set S if and only if $xP_{RC}(S)$,
- (b) x is revealed to attract attention at a set S if and only if there exists a set T such that $c(T) \neq c(T \setminus x)$, and the following two conditions hold:
 - (bi) $yP_{RC}(S)$ for all $y \in S \setminus T$
 - (bii) $zP_{RC}(T)$ for all $z \in T \setminus S$.

Part (a) is obvious. If x is revealed preferred to all elements actually chosen then it must be the case that the decision maker does not pay attention to x . Part (b) is exemplified in Example 22.17.

Propositions 22.7 and 22.8 require that c is a CLA. We now explore how we can characterize a CLA. We first define the analog of the weak axiom of revealed preference (WARP) under limited attention. Under full attention, WARP implies that every set S has a best alternative, x^* , such that if the choice from some set T lies in S then that choice must be x^* . Under limited attention we must, in addition, ensure that the decision maker pays attention to x^* .

Definition 22.16 The weak axiom of revealed preference with limited attention (WARP(LA)) requires that for a non-empty set S there exists $x^* \in S$ such that if $x^* \in T$ and,

$$\text{if (i) } c(T) \in S \text{ and (ii) } c(T) \neq c(T \setminus x^*) \text{ then } c(T) = x^*.$$

The requirement $c(T) \neq c(T \setminus x^*)$ ensures that the decision maker gives attention to x^* ; this is the only component in WARP(LP) that differentiates it from WARP.

Proposition 22.9 (Characterization of CLA; Masatlioglu et al., 2012): The choice function c satisfies WARP(LA) if and only if c is a CLA.

Proposition 22.9 shows that the characterization of CLA is simple and depends on just one behavioral axiom, WARP(LA).

The basic idea that in choosing an element from a set S , the decision maker only considers a subset $T \subset S$ plays a role in several different models. The reader might wonder how these alternative models relate to the CLA approach. For instance, Cherepanov et al. (2013) propose a model in which decision makers evaluate alternatives based on their ability to *rationalize* them.

There can be several *rationales* and each rationale is a transitive binary relation, R_i . In particular, R_i need not be complete, so R_i may not be able to compare two distinct alternatives, x and y . The set of rationalizable alternatives in S , denoted by $\hat{\Gamma}(S)$, comprises alternatives that dominate all others based on at least one rationale. Thus,

$$\hat{\Gamma}(S) = \{x \in S : \text{For some } R_i, xR_i y \text{ for all } y \in S\}. \quad (22.40)$$

It can be shown that $\hat{\Gamma}$ is not an attention filter. Furthermore, there are rationales that do not imply the CLA model and there are attention filters that do not satisfy $\hat{\Gamma}$. Hence, neither CLA implies rationalizability nor is it implied by it; an exercise invites the reader to explore this for an example.

Manzini and Mariotti (2007) propose the application of two rationales applied in a sequential and fixed manner, irrespective of the set of alternatives. The available choices are pruned by the first rationale (shortlisting) to enable the second rationale to select the final outcome. This is termed as the Rational Shortlist Method.¹² Masatlioglu et al. (2012) state that most models in this class satisfy the *Weak WARP axiom*, which states that if x is chosen from the pair of alternatives $\{x, y\}$ and from a set with a larger number of alternatives, then y cannot be chosen from a set with an intermediate number of alternatives.

Definition 22.17 (*Weak WARP; Masatlioglu et al., 2012*): Suppose $\{x, y\} \subset T \subset S$. If $c(xy) = c(S) = x$, then $y \notin c(T)$.

It is not hard to construct examples where the choice function satisfies WARP(LA) but does not satisfy Weak WARP; an exercise asks you to show this. The converse can also be shown, namely, that choice patterns that satisfy Weak WARP need not satisfy WARP(LA); see Example 3 in Masatlioglu et al. (2012).

22.8 The contractarian approach

In neoclassical economics, and in the soft paternalism approach, one typically advises a choice architect, a social planner, the government, or the policymaker. In the contractarian approach, by contrast, one advises directly the individual who suffers from behavioral biases. Sugden (2013) outlines the *contractarian approach* to normative economics that is based on Buchanan (1968, 1975, 1986); we consider this below.

The contractarian recommendation to any set of individuals is addressed to each individual in the set, not by appeal to the common interest, but by stressing the advantage to each individual. Appeals in the common interest may involve winners and losers, yet the losers may agree to participate voluntarily, perhaps motivated by the common good. However, this is not part of the contractarian framework. If the contractarian recommendations lead to harm to some parties, the proposal may include compensating mechanisms. However, a contractarian recommendation cannot impose net losses on anyone, say, on the grounds that the recommendation imparts high benefits to others.

A contractarian's role is similar to that of a mediator who makes suggestions that are mutually beneficial to different parties in a conflict situation. The contractarian does not advise any par-

¹² For a generalization of the model to accommodate more general menu dependence, see Manzini and Mariotti (2012).

ticular party to strategically out-think the others. Rather he is advising all individuals to reach a mutually beneficial settlement, although some individuals may gain relatively more than others. Since the aim is to achieve a mutually beneficial outcome for all parties, one begins from some baseline position or the status quo. For instance, consider several buyers and a seller of an object. The baseline position is that the seller has the object and the buyer has the money and both parties expect to trade voluntarily without any incurring losses. The contractarian might now recommend a second price auction to all parties as a trading mechanism to realize gains from trade.

Those who formulate contractarian policies are required to have incentives to make the necessary recommendations to individuals. Thus, the underlying framework needs to be broad enough to capture the incentives of the relevant actors, which might include political actors or policymakers who make the recommendations.

Example 22.18 *Consider the recommendation to motorists approaching a roundabout to “give way to traffic from the right.” This is, for instance, enshrined in the traffic laws in the UK. This rule is mutually beneficial to drivers at the roundabout by giving rise to common expectations. In any particular trip, a driver who has to wait for the vehicle from the right loses some time while the other party gains. However, on average, as one makes many car trips, the average gains and losses counterbalance each other, and drivers benefit from a smooth and less stressful flow of traffic, so all drivers benefit.*

Sugden (2013, pp. 530–1) raises an important issue about the relevant *addressee* in normative economics and in the contractarian approach: “The literary convention of addressing normative economics to a public decision-maker seems rather out of place when what are being discussed are (supposed) mistakes in decisions that are made by private individuals and that do not affect anyone else. Advising individuals on how to pursue their own interests in their private lives is a natural counterpart to advising them about how to pursue common interests through agreement. In other words, it is a natural counterpart to the contractarian approach.” The next example illustrates this point further.

Example 22.19 *Suppose that an epidemiologist discovers a link between some human actions, say, eating a particular food and some human disease or illness. The contractarian approach would take the form of direct information dissemination of the findings to the public. Indeed, in the case of smoking and alcohol consumption, this has been the approach in many countries (although belatedly in the case of smoking). So while normative economics is largely addressed to a hypothetical social planner, in actual practice it is not unusual that public policy takes the form of directly informing people about their biases.*

As these examples show, the contractarian approach takes the form of recommendations made directly to individuals, once it is discovered that they exhibit some form of bias. No paternalism is allowed in the contractarian approach. In particular, the contractarian cannot choose to nudge an individual who chooses not to be nudged, no matter how much the contractarian deems the nudge to be in the individual’s interest. This stands in contrast to the libertarian-paternalism approach.

Nudges recommended under soft paternalism are viewed to be paternalistic in the contractarian approach. The reason for disallowing paternalism under the contractarian approach is not because paternalism might unduly intrude into individual liberty. Rather, the reason stems from a practical concern, namely, that *paternalism lacks a valid addressee*. Sugden (2013, p. 537) writes:

“Unless individuals themselves wish to license others to act as their guardians, there is no common interest in paternalism.”

The contractarian approach is not without its limitations. Providing information in the contractarian approach may also be considered by some as a form of nudging. How does one know if individuals wish to be nudged, i.e., be provided with the relevant information? Since issues of the framing of information can be critical, who decides how the relevant information is to be framed? How does one determine if the same information will not be interpreted by different people in a different manner? Perhaps these concerns apply with less force in some examples rather than others, but clearly there are no simple answers to these vexed questions.

22.9 Behavioral public finance and welfare

In Part 7, we noted the effect of limited attention on the demand for a single good under salient and non-salient taxes (Chetty et al., 2009); we assume familiarity with the material on limited attention in Part 7 in this section. When taxes are made more salient, individuals reduce their demand for the good. In this section, we first extend this framework to consider the tax incidence and excess burden under limited attention. Then, we consider some reduced-form frameworks that are suitable to model the welfare implications of alternative policies.

22.9.1 Tax incidence under limited attention

Suppose that a consumer consumes two goods, x and y , and the utility function is additively separable

$$u(x) + v(y). \quad (22.41)$$

The consumer has income $I > 0$. Good y is the numeraire good, whose price is normalized to 1, and it is not taxed. Let p denote the pre-tax unit price of good x that is posted on the price stickers and saliently observable to all buyers. The post-tax price of the good is $q = p + \tau$, where τ is a specific consumption tax. Under full attention, denote the solution to the consumer's optimal demand for x by $x^*(p, \tau, I)$ and for y by $y^*(p, \tau, I)$; we assume that demands are continuously differentiable in the parameters p, τ, I . All feasible choices for the levels of demand must satisfy the consumer's budget constraint

$$(p + \tau) x^*(p, \tau, I) + y^*(p, \tau, I) = I. \quad (22.42)$$

Under classical public finance theory, the consumer takes account of the full post-tax price, q , so the sales tax, whether posted on price stickers or imposed at the sales register, is fully salient, so $x^*(p, \tau, I) = x^*((1 + \tau)p, 0, I)$. Thus, there is no need to separately observe the sales tax if the consumer already takes account of the post-tax price. Under full attention, an increase in the post-tax price, q , on account of an increase in τ , or an increase in p , has identical effects on the demand for x , thus, $\frac{\partial x^*}{\partial q} = \frac{\partial x^*}{\partial p} = \frac{\partial x^*}{\partial \tau}$.

Now introduce the possibility of limited attention towards the sales tax. Denote the optimal consumption bundles under limited attention by $x(p, \tau, I)$ and $y(p, \tau, I)$. In this case, whether the sales tax is displayed on the price stickers, or imposed at the sales register, has an important effect on demand, and empirical evidence indicates that $\frac{\partial x}{\partial p} \neq \frac{\partial x}{\partial \tau}$ (see Part 7 for the evidence).

Following Chetty et al. (2009), let us conduct a partial equilibrium analysis below. Our main focus will be on the demand for good x ; let us write $x(p, \tau, I) \equiv D(p, \tau, I)$ to make the demand connotation salient. Let $\epsilon_{D,p} = -\frac{\partial D}{\partial p} \frac{p}{D}$ and $\epsilon_{D,\tau} = -\frac{\partial D}{\partial \tau} \frac{\tau}{D}$ denote, respectively, the elasticity of demand for $x \equiv D$ with respect to p and τ ; we assume that both p and τ influence demand through their effect on $q = p + \tau$, so we avoid the more cumbersome but accurate notation for these elasticities: $\epsilon_{D,q|p}$ and $\epsilon_{D,q|\tau}$.

Under classical public finance theory, we should, by definition, observe that $\epsilon_{D,p} = \epsilon_{D,\tau}$; but this is no longer guaranteed under limited attention. One way of representing limited attention to taxes that are not displayed on price stickers, but rather added at the point of sale, is to write $D(p, \tau, I)$ in a log-linear form (let us keep I fixed).

$$\log D(p, \tau, I) = \alpha + \beta \log p + \theta \beta \log \tau. \quad (22.43)$$

In (22.43), the coefficients β and $\theta\beta$, respectively, give the two elasticities $\epsilon_{D,p}$ and $\epsilon_{D,\tau}$. Thus,

$$\theta = \left(\frac{\partial \log D}{\partial \log \tau} \middle/ \frac{\partial \log D}{\partial \log p} \right) = \frac{\epsilon_{D,\tau}}{\epsilon_{D,p}} = \left(\frac{\partial D}{\partial \tau} \middle/ \frac{\partial D}{\partial p} \right). \quad (22.44)$$

Under classical public finance, $\theta = 1$ (unlimited attention). However the empirical evidence suggests that $\theta < 1$ because demand is less responsive to less salient taxes, relative to changes in the price of the good, which is salient. Thus, the magnitude of θ may be taken to be a measure of inattention. More generally, $\theta < 1$ may capture other errors in optimization caused by cognitive limitations that are unrelated to inattention. While $\theta < 1$ is in agreement with the empirical evidence on consumption taxes, for other taxes, such as estate taxes, the evidence appears consistent with $\theta > 1$ (tax rates are overestimated).

From (22.44), we get

$$\epsilon_{D,\tau} = \theta \epsilon_{D,p}. \quad (22.45)$$

The problem of *tax incidence* essentially involves figuring out the relative shares of the tax burden borne by producers and consumers of an increase in the tax rate. Let us introduce producers in a simple manner. Suppose that identical competitive firms produce good x . The cost of production faced by a representative firm is $c(S)$, where $S = 1, 2, \dots$ is the number of units of the good produced by a firm, and $c' > 0$, $c'' > 0$. The profit level of the representative firm is given by $\pi = pS - c(S)$. Profit maximizing firms produce at the point where $p = c'(S(p))$; $S(p)$ is the optimal supply of the representative firm as a function of the price of the good. Denote the elasticity of supply by $\epsilon_{S,p} = \frac{\partial S}{\partial p} \frac{p}{S}$.

In a partial equilibrium analysis, we are interested in market clearing in the market for good x , so

$$D(p, \tau, I) = S(p). \quad (22.46)$$

Implicit differentiation of (22.46) gives: $\left(\frac{\partial D}{\partial p} - \frac{\partial S}{\partial p} \right) dp + \frac{\partial D}{\partial \tau} d\tau = 0$. Simplifying, and using (22.45), we get

$$\frac{dp}{d\tau} = \frac{\frac{\partial D}{\partial \tau}}{\frac{\partial S}{\partial p} - \frac{\partial D}{\partial p}} = -\frac{\epsilon_{D,\tau}}{\frac{q}{p}\epsilon_{S,p} + \epsilon_{D,p}} = -\frac{\theta \epsilon_{D,p}}{\frac{q}{p}\epsilon_{S,p} + \epsilon_{D,p}}. \quad (22.47)$$

The calculation $\frac{dp}{d\tau}$ in (22.47), gives the incidence of the tax on producers. Using the relation $q = p + \tau$, we can find the incidence of the tax on consumers as $\frac{dq}{d\tau} = 1 + \frac{dp}{d\tau}$; using (22.47) we can rewrite this as

$$\frac{dq}{d\tau} = \frac{\frac{q}{p}\epsilon_{S,p} + (1 - \theta)\epsilon_{D,p}}{\frac{q}{p}\epsilon_{S,p} + \epsilon_{D,p}}. \quad (22.48)$$

Figure 22.3 shows the effect of tax incidence under limited attention; the pretax price is measured on the vertical axis and the quantities demanded and supplied of good x are measured on the horizontal axis. The original equilibrium at a zero tax rate is shown at point 1. The corresponding pre-tax equilibrium price is p_0 and the equilibrium quantity is x_0 ; since the tax rate is zero, so the pre-tax and the post-tax prices coincide, $p_0 = q_0$. Consider now an increase in the tax rate to $\tau > 0$. The increase in the tax rate shifts the demand curve inwards by $\frac{\partial D}{\partial \tau} \tau$. At the existing equilibrium price, p_0 , there is an excess supply of the good, equal to $\frac{\partial D}{\partial \tau} \tau$; $\frac{\partial D}{\partial \tau}$ is evaluated at $p = p_0$. In order to clear the market, competition among producers ensures that the equilibrium price falls to $p = p_1$ and the equilibrium quantity falls to x_1 ; the new equilibrium point is labeled 3. The post-tax price, $q_1 = p_1 + \tau$ is also shown in the figure and τ is shown as the vertical distance between points 3 and 4. The total tax revenue collected by the government is τx_1 .

Clearly, the change in the price $dp = p_0 - p_1 < \tau$. Thus, producers bear a part of the burden of an increase in the tax burden, shown by the area $(p_0 - p_1) x_1$. The remaining tax burden, shown as the area $dq = (q_1 - p_0) x_1$, is borne by the consumers. Thus, $\tau x_1 = (p_0 - p_1) x_1 + (q_1 - p_0) x_1 \equiv x_1 dp + x_1 dq$ gives the split of the increased tax revenues between that paid by firms and consumers, although the tax is levied only on the consumers. Under full attention, we have $\theta = 1$, so $\frac{\partial D}{\partial \tau} = \frac{\partial D}{\partial p}$. Thus, in Figure 22.3, we could have shown the leftward shift in demand as $\frac{\partial D}{\partial p} \tau$, which is not the case under limited attention ($\theta < 1$); this is the only difference in Figure 22.3 relative to the standard model in public finance.

From (22.47), $\frac{dp}{d\tau}$ is directly proportional to the size of θ . In the extreme case of no attention, $\theta = 0$, we have $\frac{dp}{d\tau} = 0$, so producers do not bear any burden of the tax increase; consumers bear the entire tax burden. As θ increases, the relative tax burden borne by the producer increases, and correspondingly, the tax burden on the consumer falls. It is tempting to conclude that if




Figure 22.3 Tax incidence under limited attention.

consumers pay less than full attention to an increase in the tax rate, it is as if their demand is inelastic, hence, producers can pass a greater amount of the tax rate to final consumers. From (22.45) we know that $\epsilon_{D,\tau} = \theta \epsilon_{D,p}$, so a change in θ or a change in $\epsilon_{D,p}$ have similar effects on $\epsilon_{D,\tau}$, yet their effect on tax incidence is different.

To see this, suppose we have two economies, A and B , with identical supply elasticities $\epsilon_{S,p}^A = \epsilon_{S,p}^B = 0.1$. In economy A , all consumers are neoclassical and display perfect attention ($\theta^A = 1$), but in economy B they display limited attention ($\theta^B = 0.3 < 1$). Demand in economy A is inelastic ($\epsilon_{D,p}^A = 0.3$) while in economy B it is more elastic ($\epsilon_{D,p}^B = 1$). Thus, in each case we have $\epsilon_{D,\tau}^A = \epsilon_{D,\tau}^B = 0.3$. However, using (22.47), we have $\left(\frac{dp}{d\tau}\right)_A = -0.75$ and $\left(\frac{dp}{d\tau}\right)_B = -0.27$. Producers in economy A bear a larger share of the tax despite full attention; in contrast the limited attention in economy B allows its producers to shift a greater share of the tax to its consumers. Following an increase in the tax rate, inattention only shifts down the demand curve (as in Figure 22.3). In contrast, a lower elasticity of demand reduces the inward shift in demand and also requires a larger price cut (size of $p_0 - p_1$ in Figure 22.3) to establish the new equilibrium where demand equals supply.

From (22.47), if we hold fixed the tax elasticity of demand, $\epsilon_{D,\tau}$, a larger value of price elasticity $\epsilon_{D,p}$ increases $\frac{dp}{d\tau}$, or reduces the absolute value of the change $\left|\frac{dp}{d\tau}\right|$ needed to equilibrate markets. However, it is possible that the degree of attention given may also depend on the price elasticity, in which case one would need to model the covariance between θ and $\epsilon_{D,p}$.

22.9.2 Excess burden under limited attention

Let us now use the setup in Section 22.9.1 to calculate the excess burden of taxation, but revert back to $x(p, \tau, I)$ instead of $D(p, \tau, I)$. Assume that initially, the commodity tax rate is zero ($\tau = 0$). Suppose that there are constant returns to scale in production and firms are perfectly competitive, so each firm makes zero profits. Thus, welfare measurements can be equated purely with consumer welfare.

Substituting $x(p, \tau, I)$ and $y(p, \tau, I)$ in the utility function (22.41), we get the indirect utility function

$$V(p, \tau, I) = u(x(p, \tau, I)) + v(y(p, \tau, I)). \quad (22.49)$$

Let the expenditure function be denoted by $e(p, \tau_1, I)$; it is the minimum expenditure required to achieve the level of utility $V(p, \tau_1, I)$ (corresponding to a tax rate of τ_1) when the price is p and the tax rate is τ_2 . We denote government revenue by $R = \tau x(p, \tau, I)$. From (22.42), we have

$$px(p, \tau, I) + y(p, \tau, I) = I - \tau x(p, \tau, I) \equiv I - R(p, \tau, I). \quad (22.50)$$

The LHS of (22.50) is the total expenditure on consumption goods, and the RHS is the total disposable income.

We now use these concepts to define the excess burden, or deadweight loss, of a consumption tax by

$$EB(\tau) = I - R(p, \tau, I) - e(p, 0, V(p, \tau, I)). \quad (22.51)$$

$I - R(p, \tau, I)$ is the disposable income when the distortionary consumption tax rate is τ , and the corresponding indirect utility is $V(p, \tau, I)$. Suppose that we ask what expenditure is needed to achieve a utility of $V(p, \tau, I)$ when we remove the distortionary tax and replace it by a lump-sum tax, L ? The budget constraint with a lump-sum tax is $px(p, 0, I) + y(p, 0, I) = I - L$. Since there are no distortions associated with a lump-sum tax, hence, the same level of utility can be achieved with a lower expenditure, $e(p, 0, V(p, \tau, I))$. Thus, the RHS of (22.51) measures the extra expenditure required under a distortionary tax to achieve the same level of utility as a lump-sum tax; in other words, it is the deadweight loss associated with the tax.

An issue in most behavioral welfare analyses where people do not strictly optimize is to compute the indirect utility $V(p, \tau, I)$. The indirect utility function is obtained by substituting the optimal choices into the utility function. Suppose, as an extreme example, that an individual pays no attention to tax rates ($\theta = 0$). Then, as the tax rate changes, the individual does not choose the optimizing bundles of x, y , thus making it difficult to define an indirect utility function. In order to deal with this problem, consider two assumptions.

- A1: Taxes do not directly affect the utility function. Rather, taxes only affect the utility function indirectly through changes in the demands $x(p, \tau, I)$ and $y(p, \tau, I)$. In other words, the indirect utility function is as given in (22.49).
- A2: Recall that $x^*(p, \tau, I)$ and $x(p, \tau, I)$ are the optimal consumption bundles under, respectively, full attention and limited attention, and that our initial tax rate is $\tau = 0$. We require that

$$x(p, 0, I) = x^*(p, 0, I) = \arg \max u(x(p, 0, I)) + v(I - px(p, 0, I)). \quad (22.52)$$

The implication of (22.52) is that under limited attention, we can discover the underlying preferences of a decision maker who exhibits limited attention, if $\tau = 0$; when $\tau = 0$, there are no issues of inattention.

It is best to examine conditions under which these assumptions are violated in order to understand them better.

Assumption A1 is not likely to hold if taxes directly influence individual utility. This may happen, for instance, (1) if individuals are troubled by the fairness of the fiscal system (e.g., those individuals who feel that the tax system is unfair may evade more taxes), or (2) if individuals must expend costly cognitive effort to give more attention to taxes. In these cases, instead of (22.49), the indirect utility function may have to be written as $\tilde{V}(p, \tau, I) = u(x(p, \tau, I), \tau) + v(y(p, \tau, I), \tau)$. Assumption A2 can be violated even when in the absence of taxes, individuals do not choose the optimal consumption bundles; for instance, under *shrouded attributes* (see Part 7).

Thus, our strategy is to first use (22.52) to compute the individual's underlying preferences, and then use the Marshallian demand functions under limited attention, $x(p, \tau, I)$ and $y(p, \tau, I)$, to compute the indirect utility function in (22.49).

Denote by $h = h(p, \tau, V)$ the Hicksian or compensated demand for good x , under limited attention. Then assumption A2, and the Slutsky equation implies

$$\frac{\partial h}{\partial p} = \frac{\partial x}{\partial p} + x \frac{\partial x}{\partial I}. \quad (22.53)$$

Demands, Hicksian or Marshallian, optimally respond to prices, so we have the usual condition $\frac{\partial h}{\partial p} < 0$. However, due to limited attention, there is no presumption that demands respond optimally to changes in the tax rate, τ . So if we define

presence of income effects we would have to consider Hicksian demands in Figure 22.4. In the absence of income effects, we can also replace θ^c by θ in (22.55).

Two kinds of demand curves are shown in Figure 22.4. $x(p, 0)$ is the demand curve when $\tau = 0$, and it shows how optimal demand varies under limited attention when the price alone varies. The initial commodity price of good x is p_0 , and $\epsilon_{s,p} = \infty$, so the supply curve is horizontal at p_0 . The demand curve $x(p_0, \tau)$ shows the variation in optimal demand when price is fixed at p_0 , but the tax rate, τ , varies. In the figure, we have shown $x(p_0, \tau)$ to be steeper relative to $x(p, 0)$; this is consistent with the empirical evidence that consumers respond less to taxes that are less salient.

The equilibrium when $\tau = 0$ is shown at point J and the consumer demands x_0 units at a price p_0 per unit. Now suppose that we impose a commodity tax $\tau > 0$ on x so that the post-tax price of x is $p_0 + \tau$. In response to an increase in the tax, the consumer reduces demand from x_0 to x_1 . However, at a level of demand x_1 , the marginal willingness to pay is given by the demand curve $x(p, 0)$, i.e., the vertical distance FM , which is lower than the actual price paid, $p_0 + \tau$. Suppose that we use as our measure of welfare, the sum of consumer surplus and government tax revenues. Before the imposition of the tax, tax revenues are zero and consumer surplus is AGJ , so initial total welfare is AGJ . After the imposition of the tax, consumer surplus is $AKMF - BKMD = ABC - CDF$. Total government tax revenues equal $BGIFC + CDF$. Thus, total welfare after the imposition of the tax is $ABC + BGIFC$. The difference in the levels of welfare in the two cases is $AGJ - (ABC + BGIFC) = FIJ$; this is the excess burden of the tax and is shown as the shaded area in Figure 22.4.

As noted above, in the absence of income effects, $\frac{\partial x}{\partial I} = 0$ (e.g., underlying utility function is quasilinear), so we may write $EB(\tau) \simeq \frac{1}{2}(\tau\theta)^2 x(p, \tau, I) \frac{\epsilon_{x,p}}{p+\tau}$. Thus, if $\theta = 0$ (tax changes are ignored, so extreme limited attention), then $EB(\tau) \simeq 0$. The reason is that starting from a zero tax rate, if the imposition of a positive tax rate is ignored by the individual, then there are no substitution effects, only income effects, hence there is no deadweight loss either. When $\theta = 0$, the budget constraint of the consumer must nevertheless be respected on imposition of the tax. Thus, the distortionary tax under extreme limited attention ($\theta = 0$) becomes just a lump-sum tax, that leads to a choice of the first best consumption bundle, x^* . On the other hand, if $0 < \theta < 1$ then the excess burden increases with the square of θ , just as it does for the tax rate, τ . Thus, limited attention reduces the excess burden of taxation, which is welfare improving. In classical public finance theory, it does not matter if taxes are imposed on consumers or producers. However, under limited attention, the two taxes may lead to different effects. The reason is that taxes on producers are often included in the posted price of the good, hence, they are more salient (higher θ); thus the excess burden arising from these taxes is relatively greater.

Going back to Figure 22.4, where we assumed $\frac{\partial x}{\partial I} = 0$, suppose that the curve $x(p, 0)$ becomes steeper (price elasticity reduces), then the area FIJ expands, so the excess burden increases.

If there are income effects $\frac{\partial x}{\partial I} > 0$, then limited attention may increase the excess burden of taxation; this result arises even if $\theta = 0$, so the consumer pays no attention to tax, $\frac{\partial x}{\partial \tau} = 0$ (see (22.44)). When $\frac{\partial x}{\partial \tau} = 0$, we get from (22.54) that $\frac{\partial h}{\partial \tau} = x \frac{\partial x}{\partial I}$. Substitute this, and $\theta^c = \left(\frac{\partial h}{\partial \tau} / \frac{\partial h}{\partial p} \right)$, into the first line of (22.55), to get

$$EB(\tau) \simeq \frac{-1}{2} (\tau x)^2 \left(\frac{\partial x}{\partial I} / \frac{\partial h}{\partial p} \right) \frac{\partial x}{\partial I} > 0. \quad (22.56)$$

The sign in (22.56) follows because $\frac{\partial h}{\partial p} < 0$. Thus, under complete inattention, we can get a positive excess burden of taxation when there are income effects.

To see this, consider an example in which x is the number of visits to the theater and y is food. Suppose that the government levies tax, $\tau > 0$, on theater visits and the consumer is completely inattentive to these taxes ($\theta = 0$). The budget constraint of the consumer is given by $(p + \tau)x(p, \tau, I) + y(p, \tau, I) = I$. Differentiate both sides with respect to τ to get

$$\frac{\partial x}{\partial \tau} + \frac{\partial y}{\partial \tau} = -x.$$

Thus, one method by which the budget may be balanced is that the individual ignores the effect of tax on x , i.e., sets $\frac{\partial x}{\partial \tau} = 0$. In this case, as the consumer's income goes down after paying the tax, the full effect is felt as a decrease in the consumption of food, y . Thus, the new income level, net of taxes, is not allocated efficiently between theater visits and food, which creates a deadweight loss from taxation, hence, the excess burden is positive.

Finally, note from Proposition 22.10, when tax increases are imposed on pre existing positive tax rates, then the tax increases can have first order effects (as in the term $\tau_0 \Delta \tau$), although these effects are weakened by inattention (low values of θ^c).

22.9.3 A note on reduced-form approaches to behavioral public finance

A departure in behavior from the neoclassical benchmark can sometimes be explained by several competing behavioral explanations. Consider, for instance, the finding that individuals might buy energy inefficient lightbulbs and cause carbon externalities on others (Hausman, 1979; Allcott et al., 2014). Do individuals engage in this decision because they are present-biased (lightbulbs impose current costs, but future benefits; see Part 3 of the book), or simply because they possess limited attention (insufficient attention to future benefits; see Part 7 of the book)? Competing behavioral explanations of the same phenomenon are viewed by some as a drawback of behavioral economics. This is a mistake. There are almost always competing explanations of a phenomenon in any science; even in neoclassical economics one can often explain the same phenomenon by simply invoking some form of asymmetric information or other auxiliary assumptions. It is the role of empirical evidence to help us sift through these alternative explanations and pin down the empirically correct explanation.

Reduced-form approaches to behavioral public finance seek to provide a framework that is robust to alternative behavioral explanations, within some subset of explanations. Mullainathan et al. (2012) propose a reduced-form approach that nests competing behavioral explanations that draw on the distinction between experienced utility and decision utility (see Parts 6 and 7 of the book). We now consider the basic model.

Suppose that there are n individuals indexed by $i = 1, \dots, n$. Individual i can take a binary action $a_i \in \{0, 1\}$, for instance, buy or not buy an energy efficient product; 0 denotes *inaction*, while 1 denotes taking the *action*. The action costs p (price and other costs associated with the action). The benefit to individual i from the action, in income equivalent units, is b_i ; benefits have a distribution function F in the population. By taking the action, an individual imposes an externality e on all individuals (e.g., higher carbon emissions for all when an individual buys a more polluting car). So if $m \leq n$ individuals choose to take the action, then the externality on any individual is me .

The government levies a tax t on the action taken by an individual. Hence, the post-tax price paid by an individual who chooses to take the action is $q(t)$; for instance, in the case of a specific

consumption tax, $q(t) = p + t$. The government also transfers an amount T to each individual, so the government budget constraint can be written as $G(t, T) = 0$.

The initial income of each individual is Y , and each individual has an increasing and concave utility function, U , that is defined as follows.

$$U = \begin{cases} U(Y + T + (b - q(t)) - me) & \text{if } a_i = 1 \\ U(Y + T - me) & \text{if } a_i = 0 \end{cases} . \quad (22.57)$$

It follows that $a_i = 1$ is strictly preferred to $a_i = 0$ by a neoclassical consumer, if

$$b > q(t). \quad (22.58)$$

We can also find that $m = 1 - F(q(t))$. In contrast to (22.58), behavioral consumers, in this reduced-form approach, chose $a_i = 1$ over $a_i = 0$, if

$$b + d(b, t) > q(t), \quad (22.59)$$

where $d(b, t) \geq 0$ is a non-stochastic distortion term that reflects behavioral factors, and is increasing in b . We may describe the decision rule (22.59) as reflecting *decision utility* and the decision rule (22.58) as reflecting *experienced utility*. This framework applies to any situation where there is a wedge between experienced and decision utility. Clearly this does not apply to many kinds of behavioral factors such as social preferences or to the presence of reference points, where no such meaningful distinction between the two notions of utility is required.

Example 22.20 (Laibson, 1997; O'Donoghue and Rabin, 1999): Suppose that the individual takes some action that involves current benefits and future costs (e.g., buying consumer durables on credit; see Part 3 for several examples). Let $b = v - c$, where v is the current benefit and c is the future cost. Suppose that $t = T = 0$ and the price $p = p_0 + p_1$, where p_0 is the price paid now and p_1 is the price paid in the future. Suppose that $\beta \in (0, 1)$, then present-biased individuals choose to take the action if $v - \beta c - (p_0 + \beta p_1) > 0$, or

$$v - c + (1 - \beta)(c + p_1) > p_0 + p_1.$$

Setting $d(b, t) = (1 - \beta)(c + p_1)$ and $q = p_0 + p_1$, we get the decision rule in (22.59).

Example 22.21 (Nudges, Thaler, and Sunstein, 2009): Suppose that there are $j = 1, 2, \dots, k$ possible nudges. Then, we may index $d(b, t)$ in (22.59) by j , so that the decision rule under nudges is $b + d_j(b, t) > q(t)$.

The main insight offered by this reduced-form approach is to show that the identity of the marginal individuals on which policy applies in the presence and in the absence of behavioral factors is different. Suppose, for instance, that the government increases the tax rate from t to $t + \Delta t$. Prior to the tax increase, the individuals who took the action satisfied the condition

$$b > p + t - d,$$

while after the tax increase, the individuals who take the action satisfy the condition

$$b > p + t + \Delta t - d.$$

The marginal behavioral individual for small changes, $\Delta t \rightarrow 0$, is one whose benefit from the action is $b \approx p + t - d$. In neoclassical economics, the marginal individual has benefits $b \approx p + t$, which reflect the true marginal cost of taking the action. Suppose that $d > 0$ (e.g., sin goods; see Part 3 for a framework of analysis), then for a behavioral individual, the marginal benefit, b , of the good is below the marginal cost, $p + t$. The extra tax could then dissuade this individual from undertaking the activity, which is welfare improving. In Part 3, we report evidence, based on Canadian data, that cigarette smokers welcome an increase in the tax rate. As Mullainathan et al., (2014, p. 17.9) put it: “For many policy problems, the feature of the behavioral bias that matters most is how it affects who is marginal. This change is at the heart of our entire approach.”

Farhi and Gabaix (2015) propose an ambitious research program in which they model limited attention by using the sparse max operator (see Part 7 for the details) and develop the theory of optimal taxation with behavioral agents. They derive optimal tax formulae under these conditions for commodity taxes (the analogue of Ramsey taxes), income taxes (the analog of the Mirrlees approach to non-linear income taxation), and corrective taxes in the presence of externalities (the analog of the Pigouvian framework in classical public finance). The optimal taxes in classical public finance are often too complex relative to real-world taxes that are relatively simple. The authors exploit insights from limited attention to show that taxpayers may give even lower salience to complex tax systems, driving an even greater wedge between the private marginal cost and the private marginal benefit of a decision, making it welfare improving to design simpler tax systems.

Exercises for Part 8

1. (Camerer et al., 2003) (a) Evidence shows that people are ambiguity averse (see Part 1) and have social preferences (see Part 2). Do these findings create a need for paternalism?
 (b) Use asymmetric paternalism in order to evaluate the following policy proposal. A policy that requires patients to get a second opinion before undergoing surgery.
 (c) Can one use asymmetric paternalism to rationalize the existing licensing regulations in many countries for physicians (and car drivers) that bar entry into the profession if one does not have the appropriate license?
 (d) How can *cooling-off periods* for consumers, advocated under soft paternalism, influence the behavior of sellers?
2. (Camerer et al., 2003) Suicide is currently illegal, which is a form of *heavy-handed paternalism*. How might you deal with the problem of suicide using soft paternalism?
3. (Camerer et al., 2003) How might one use asymmetric paternalism in order to address the problems indicated in the following empirical findings?
 (a) Enron's 401(k) plan was heavily invested in the company's own stock with the result that when Enron collapsed, so did the pension plans of the employees.
 (b) Evidence indicates that people vastly overestimate the chances of winning lotteries, which might account for the large participation rates in lotteries.
 (c) "Rent to own" establishments lease consumer durables and furniture, typically to low-income consumers. Most US states treat these contracts as rental agreements even though 70% of the consumers buy the products they rent. This frees up firms from the regulations that cap interest rates. Consumers end up paying high interest rates.
 (d) People often run up high credit card debts because they fail to fully appreciate the difference between simple and compound interest rates.
4. Derive the analog of Proposition 22.1 under Type 2 asymmetric information (Definition 22.3b).
5. (Dalton and Ghosal, 2011) Consider the framework developed in Section 22.2. Suppose that the action set is $A = \{a_1, a_2, a_3\}$ and the set of psychological states is given by $\Theta = \{\theta_1, \theta_2, \theta_3\}$. The decision maker faces payoff relevant uncertainty. There are two equally likely states of the world in which the payoffs corresponding to each action–psychological state pair are given in Tables 8.E1 (for state 1) and 8.E2 (for state 2).
 (a) Suppose that the decision maker must choose his action before uncertainty about the state of the world resolves and that he is an expected utility maximizer. What is the unique solution to the behavioral decision problem (BDP) and the corresponding payoff of the decision maker? Would the decision maker be better off if he had known the state of the world with certainty?
 (b) Consider the payoff matrix in Table 8.E3.

Table 8.E1 Payoffs in state 1.

	θ_1	θ_2	θ_3
a_1	−1	0	0
a_2	0	3	0.5
a_3	1	4	1

Source: Dalton and Ghosal, 2011.

Table 8.E2 Payoffs in state 2.

	θ_1	θ_2	θ_3
a_1	1	4	1
a_2	0.5	3	0
a_3	0	0	−1

Source: Dalton and Ghosal, 2011.

Table 8.E3 Payoffs for each pair of actions and psychological states.

	θ_1	θ_2	θ_3
a_1	−1	0	0
a_2	0	3	0.5
a_3	1	4	1

Source: Dalton and Ghosal, 2011.

Suppose first that $A = \{a_1, a_2\}$ and $\Theta = \{\theta_1, \theta_2\}$, so the relevant payoff matrix is the 2×2 payoff submatrix in Table 8.E3. Now expand the set of actions and psychological states such that $A = \{a_1, a_2, a_3\}$ and $\Theta = \{\theta_1, \theta_2, \theta_3\}$; the payoff matrix is given by the entire 3×3 payoff matrix in Table 8.E3. Find the unique BDP solution in each case. Is the decision maker better off or worse off following an increase in the set of actions?

6. Consider the Bernheim and Rangel (2009) framework developed in Section 22.6. Consider the setup in Example 22.9 but replace (22.25) by the following choices

$$C(X_1) = \{a\}, C(X_2) = \{b\}, C(X_3) = \{c\}, C(X_4) = \{c\}.$$

Show that the binary relation R^* is intransitive.

7. (Bernheim and Rangel, 2009) Consider the Bernheim–Rangel framework in Section 22.6. Let $X_1 = \{a, b\}$, $X_2 = \{b, c\}$, $X_3 = \{a, c\}$, $X_4 = \{a, b, c\}$, and $\chi = \{X_1, X_2, X_3, X_4\}$. Suppose that we have two possible ancillary conditions A and A' . The welfare relevant GCSs are given by $G = \chi \times \{A, A'\}$. Under the ancillary condition A , b is never chosen when a is available and c is never chosen. Under the ancillary condition A' , b is never chosen when c is available and a is never chosen.
- (a) Are any of the alternatives, a, b, c , comparable under P or P^* ?
- (b) What conclusions can you draw about the set of weak and strict individual optima (see Definition 22.8).
8. (Masatlioglu et al., 2012) Consider the CLA method in Section 22.7. Show that the complete and transitive relation P in this model is acyclic if and only if the choice function c satisfies WARP(LA) given in Definition 22.16.
9. (Masatlioglu et al., 2012; Cherepanov et al., 2013) Consider the Rationalization model in Section 22.7 and the associated set of rationalizable alternatives, $\hat{\Gamma}$, given in (22.40). Construct examples to show the following.
- (a) Show using an example with 3 alternatives and 2 rationales that $\hat{\Gamma}$ is not an attention filter (Definition 22.11).

(b) Show that for any rationalization, R_i , and some set S of alternatives

$$x \in \hat{\Gamma}(S) \Rightarrow x \in \hat{\Gamma}(T) \text{ for all } x \in T \subset S.$$

Does this property hold in the CLA model?

(c) What relationship can you conclude between the CLA model and the Rationalization model from your answers to (a) and (b)?

(d) Show that the Rationalization model is a special case of the CLA model if the rationales, R_i , were transitive and complete for all i (Note: in the Rationalization model R_i is incomplete).

10. (Masatlioglu et al., 2012) Consider the model of CLA in Section 22.7. Suppose that there are four alternatives x, y, a, b . The alternatives a and b are never chosen (unless there is no other alternative) but they alter the attention of the decision maker. The decision maker's preference is given by $y \succ x \succ a \succ b$ and she picks the best alternative using these preferences. She considers y only when either a or b is feasible, but not both, and always considers all other alternatives. Show that her consideration set mapping satisfies WARP(LA) (Definition 22.16) but it does not satisfy Weak WARP (Definition 22.17).
11. (Sugden, 2013) Thaler and Sunstein (2009), in their libertarian paternalism approach, give the example of the operation of a cash machine. The machine returns the card first, which if taken out first, triggers the delivery of cash. This prevents individuals from forgetting their card in the cash machine. Consider now the contractarian approach outlined in Section 22.8. Suppose that we are back in time when machines returned cash before the card. How would a contractarian design the appropriate intervention if it were found that people sometimes forget cards in the machine? In your answer, consider two different cases. (i) The cash machine is owned by a private firm. (ii) The cash machine is owned by a public agency that finances its operations through general taxation.
12. (Mullainathan et al., 2012) Consider the reduced-form model in Section 22.9.3. In each of the two examples below, show that the decision rule for the individual can be stated in terms of the decision utility interpretation in (22.59) and find an explicit formula for $d(b, t)$.
 - (a) Inattention to salient components of price (Gabaix and Laibson, 2006; Chetty et al., 2009; DellaVigna, 2009) Suppose that individuals decide to take some action for which benefits are salient. However, the individual pays limited attention to sales taxes. Inattentive individuals discount non-salient costs by the factor $\theta \in (0, 1)$. Thus, inattentive consumers think that they are paying a price $p + \theta t$, where p is the retail price per unit and t is a sales tax.
 - (b) False beliefs and overconfidence (Sandroni and Squintani 2007, Spinnewijn 2015) Suppose that agents misestimate or have false beliefs about the benefits from an action. For instance, the unemployed could be too overconfident of their ability to find a job, hence, searching suboptimally for a new job. Conversely, people may overweigh the benefits to taking an action, for example, when buying an inefficacious herbal medicine. Instead of taking the action when $b > q(t)$, such people take the action when $\hat{b} > q(t)$.

REFERENCES FOR PART 8

- Abadie, A. and Gay, S. (2004). The impact of presumed consent legislation on cadaveric organ donation: a cross country study. *Journal of Health Economics* 25: 599–620.
- Afriat, S. (1967). The construction of a utility function from expenditure data. *International Economic Review* 8: 67–71.
- Allcott H., Mullainathan S., and Taubinsky D. (2014). Externalities, internalities, and the targeting of energy policy. *Journal of Public Economics* 112: 72–88.
- Ariely, D. and Wertenbroch, K. (2002). Procrastination, deadlines, and performance: self-control by precommitment. *Psychological Science* 13: 219–24.
- Atkinson, A. and Stiglitz, J. (1980). *Lectures on Public Economics*. McGraw-Hill.
- Bandura, A. (1977). *Social Learning Theory*. New York: General Learning Press.
- Benartzi, S. and Thaler, R. H. (2004). Save more tomorrow: using behavioral economics to increase employee savings. *Journal of Political Economy* 112: S164–S187.
- Benjamin, D. J., Heffetz, O., Kimball, M. S., and Rees-Jones, A. (2012). What do you think would make you happier? What do you think you would choose? *American Economic Review* 102(5): 2083–110.
- Bentham, J. (1789). *An Introduction to the Principles of Morals and Legislation*. Oxford: Clarendon Press.
- Bernheim, B. D. and Rangel, A. (2004). Addiction and cue-conditioned cognitive processes. *American Economic Review* 94(5): 1558–90.
- Bernheim, B. D. and Rangel, A. (2007). Toward choice-theoretic foundations for behavioral welfare economics. *American Economic Review Papers and Proceedings* 97: 464–70.
- Bernheim, B. D. and Rangel, A. (2009). Beyond revealed preference: choice-theoretic foundations for behavioral welfare economics. *Quarterly Journal of Economics* 124(1): 51–104.
- Beshears, J., Choi, J. J., Laibson, D., and Madrian, B. C. (2008). How are preferences revealed? *Journal of Public Economics* 92(8–9): 1787–94.
- Buchanan, J. M. (1968). *The Demand and Supply of Public Goods*. Chicago: Rand McNally.
- Buchanan, J. M. (1975). *The Limits of Liberty*. Chicago University of Chicago Press.
- Buchanan, J. M. (1986). *Liberty, Market and State*. Brighton: Wheatsheaf.
- Camerer, C. F. (2008). The case for mindful economics. In A. Caplin and A. Schotter (eds.), *The Foundations of Positive and Normative Economics*. Oxford: Oxford University Press.
- Camerer, C. F., Issacharoff, S., Loewenstein, G., O'Donoghue, T., and Rabin, M. (2003). Regulation for conservatives: behavioral economics and the case for asymmetric paternalism. *University of Pennsylvania Law Review* 151: 1211–54.
- Caplin, A. (2008). Economic theory and psychological data: bridging the divide. In A. Caplin, and A. Schotter (eds.), *The Foundations of Positive and Normative Economics*. Oxford: Oxford University Press.
- Caplin, A. and Schotter, A. (eds.) (2008). *The Foundations of Positive and Normative Economics*. Oxford: Oxford University Press.
- Cason, T. N. and Plott, C. R. (2014). Misconceptions and game form recognition: challenges to theories of revealed preference. *Journal of Political Economy* 122(6): 1235–70.
- Cherepanov, V., Feddersen, T., and Sandroni, A. (2013). Rationalization. *Theoretical Economics* 8: 775–800.
- Chetty, R., Looney, A., and Kroft, K. (2009). Salience and taxation: theory and evidence. *American Economic Review* 99(4): 1145–77.
- Choi, J. J., Laibson, D., and Madrian, B. C. (2004). Plan design and 401(k) savings outcomes. *National Tax Journal* 57: 275–98.
- Choi, S., Kariv S., Müller, W., and Silverman, D. (2014). Who is (more) rational? *American Economic Review* 104(6): 1518–50.
- Dalton, P. S. and Ghosal, S. (2011). Behavioral decisions and policy. *CESifo Economic Studies* 57(4): 560–80.
- Dalton, P. S. and Ghosal, S. (2016). Self fulfilling mistakes: characterization and welfare. forthcoming in *Economic Journal*.
- Dalton, P. S., Ghosal, S., and Mani, A. (2016). Poverty and aspirations failure. *Economic Journal*. 126(590): 165–188.
- DellaVigna S. (2009). Psychology and economics: evidence from the field. *Journal of Economic Literature* 47(2): 315–72.
- Eliaz, K. and Schotter, A. (2010). Paying for confidence: an experimental study of the demand for

- non-instrumental information. *Games and Economic Behavior* 70(2): 304–24.
- Farhi, E. and Gabaix, X. (2015). Optimal taxation with behavioral agents. Manuscript under preparation.
- Fleurbaey, M. and Blanchet, D. (2013). *Beyond GDP: Measuring Welfare and Assessing Sustainability*. Oxford: Oxford University Press.
- Fleurbaey, M. and Schokkaert, E. (2013). Behavioral welfare economics and redistribution. *American Economic Journal: Microeconomics* 5(3): 180–205.
- Gabaix, X. and Laibson, D. (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quarterly Journal of Economics* 121: 505–40.
- Glaeser, E. L. (2006). Paternalism and psychology. *University of Chicago Law Review* 73: 133–156.
- Goeree, J. K. and Holt, C. A. (2001). Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review* 91(5): 1402–22.
- Guala, F. and Mittone, L. (2015). A political justification of nudging. *Review of Philosophy and Psychology*. Online March 5, 2015.
- Gul, F. and Pesendorfer, W. (2008). The case for mindless economics. In A. Caplin and A. Schotter (eds.), *The Foundations of Positive and Normative Economics*. Oxford: Oxford University Press.
- Hausman, J. A. (1979). Individual discount rates and the purchase and utilization of energy-using durables. *Bell Journal of Economics* 10(1): 33–54.
- Houthakker, H. S. (1950). Revealed preference and the utility function. *Economica* 17(66): 159–74.
- Iyengar, S. S. and Lepper, M. R. (2000). Why choice is demotivating: can one desire too much of a good thing? *Journal of Personality and Social Psychology* 79: 995–1006.
- Johnson, E. J., Camerer, C. F., Sen, S., and Rymon, T. (2002). Detecting failures of backward induction: monitoring information search in sequential bargaining. *Journal of Economic Theory* 104(1): 16–47.
- Johnson, E. J. and Goldstein, D. G. (2003). do defaults save lives? *Science* 302: 1338–9.
- Johnson, E. J., Hershey J., Meszaros, J., and Kunreuther, H. (1993). Framing, probability distortions, and insurance decisions. *Journal of Risk and Uncertainty* 7: 35–53.
- Jolls, C., Sunstein, C. R., and Thaler, R. H. (1998). A behavioral approach to law and economics. *Stanford Law Review* 50: 1471–550.
- Kahneman, D., Knetsch, J., and Thaler, R. (1991). The endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives* 5: 193–206.
- Kahneman, D. and Tversky, A. (1984). Choices, Values, and Frames. *The American Psychologist* 39: 341–50.
- Kahneman, D. and Tversky, A. (eds.). (2000). *Choices, Values and Frames*. New York: Cambridge University Press.
- Kahneman, D., Wakker, P., and Sarin, R. (1997). Back to Bentham? Explorations of experienced utility. *Quarterly Journal of Economics* 112: 375–406.
- Kamenica, E., Mullainathan, S., and Thaler, R. (2011). Helping consumers know themselves. *American Economic Review* 101(3): 417–22.
- Klick, J. and Mitchell, G. (2006). Government regulation of irrationality: Moral and cognitive hazards. *Minnesota Law Review* 90: 1620–1663.
- Koszegi, B. and Rabin, M. (2008a). Revealed mistakes and revealed preference. In A. Caplin and A. Schotter (eds.), *The Foundations of Positive and Normative Economics*. Oxford: Oxford University Press.
- Koszegi, B. and Rabin, M. (2008b). Choices, situations, and happiness. *Journal of Public Economics* 92: 1821–32.
- Kreps, D. M. (1990). *A Course in Microeconomic Theory*. Princeton, NJ: Princeton University Press.
- Kreps, D. M. and Porteus, E. L. (1978). Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* 46: 185–200.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 112: 443–77.
- Layard, R. (2005). *Happiness. Lessons from a New Science*. London: Allen Lane.
- Loewenstein, G. and Haisley, E. (2008). The economist as therapist: methodological ramifications of “Light” Paternalism. In A. Caplin and A. Schotter (eds.), *The Foundations of Positive and Normative Economics*. Oxford: Oxford University Press.
- Loewenstein, G. and Ubel, P. A. (2008). Hedonic adaptation and the role of decision and expe-

- rience utility in public policy. *Journal of Public Economics* 92: 1795–810.
- Madrian, B. C. and Shea, D. F. (2001). The power of suggestion: inertia in 401(k) participation and savings behavior. *Quarterly Journal of Economics* 116: 1149–525.
- Manzini, P. and Mariotti, M. (2007). Sequentially rationalizable choice. *American Economic Review* 97(5): 1824–39.
- Manzini, P. and Mariotti, M. (2012). Categorize then choose: boundedly rational choice and welfare. *Journal of the European Economic Association* 10(5): 1141–65.
- Manzini, P. and Mariotti, M. (2014). Welfare economics and bounded rationality: the case for model-based approaches. *Journal of Economic Methodology* (21)4: 343–60.
- Masatlioglu, Y., Nakajima, D., and Ozbay, E. Y. (2012). Revealed attention. *American Economic Review* 102(5): 2183–205.
- Mas-Colell, A., Whinston, M. D., and Green, J. R. (1995). *Microeconomic Theory*. New York: Oxford University Press.
- Mullainathan, S., Kling J. R., Shafir, E., Vermeulen, L., and Wrobel, M. V. (2012). Comparison friction: experimental evidence from medicare drug plans. *Quarterly Journal of Economics* 127(1): 199–235.
- Mullainathan, S., Schwartzstein, J., and Congdon, W. (2012). A reduced-form approach to behavioral public finance. *Annual Review of Economics* 4: 511–40.
- O'Donoghue, T. and Rabin, M. (1999). Doing it now or later. *American Economic Review* 89: 103–24.
- Park, C. W., Jun, S. Y., and MacInnis, D. J. (2000). Choosing what I want versus rejecting what I do not want: an application of decision framing to product option choice decisions. *Journal of Marketing Research* 37: 187–202.
- Piccione, M. and Spiegler, R. (2012). Price competition under limited comparability. *Quarterly Journal of Economics* 127: 97–135.
- Posner, R. A. (1986). *Economic Analysis of Law*. 3rd edition. Boston: Little, Brown.
- Riis, J., Loewenstein, G., Baron, J., Jepson, C., Fagerlin A., and Ubel, P. A. (2005). Ignorance of hedonic adaptation to hemo-dialysis: a study using ecological momentary assessment. *Journal Experimental Psychology: General* 134: 3–9.
- Rubinstein, A. and Salant, Y. (2012). Eliciting welfare preferences from behavioural data sets. *Review of Economic Studies* 79: 375–387.
- Salant, Y. and Rubinstein, A. (2008). (A, f): choice with frames. *Review of Economic Studies* 75(4): 1287–9.
- Samuelson, P. A. (1948). Consumption theory in terms of revealed preference. *Economica* 15(60): 243–53.
- Samuelson, W. and Zeckhauser, R. J. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty* 1: 7–59.
- Sandroni, A. and Squintani, F. (2007). Overconfidence, insurance, and paternalism. *American Economic Review* 97: 1994–2004.
- Schotter, A. (2008). What's so informative about choice? In A. Caplin and A. Schotter (eds.), *The Foundations of Positive and Normative Economics*. Oxford: Oxford University Press.
- Spinnewijn, J. (2015). Unemployed but optimistic: optimal insurance design with biased beliefs. *Journal of the European Economic Association* 13(1): 130–67.
- Stiglitz, J. E. (2000). *Economics of the Public Sector*. 3rd edition. New York: W.W. Norton & Company.
- Sugden, R. (2008). Why incoherent preferences do not justify paternalism. *Constitutional Political Economy* 19: 226–48.
- Sugden, R. (2013). The behavioural economist and the social planner: to whom should behavioural welfare economics be addressed? *Inquiry: An Interdisciplinary Journal of Philosophy* 56(5): 519–38.
- Sunstein, C. R. and Thaler, R. H. (2003). Libertarian paternalism is not an oxymoron. *University of Chicago Law Review* 70: 1159–202.
- Thaler, R. H. and Sunstein, C. R. (2003). Libertarian paternalism. *American Economic Review: Papers and Proceedings* 93: 175–9.
- Thaler, R. H. and Sunstein, C. R. (2009). *Nudge: Improving Decisions about Health, Wealth, and Happiness*. New Haven, CT: Penguin Books.
- Tversky, A. and Kahneman, D. (1981). The framing of decisions and psychology of choice. *Science* 211: 453–8.
- Ubel, P. A., Loewenstein, G., Schwarz, N., and Smith, D. (2005). Misimagining the unimaginable: the disability paradox and healthcare decision making. *Health Psychology* 24: S57–S62.
- Wisdom, J., Downs, J. S., and Loewenstein, G. (2010). Promoting healthy choices: information versus convenience. *American Economic Journal: Applied Economics* 2(2): 164–78.

PART 9

Neuroeconomics

INTRODUCTION TO PART 9

Part 9 offers a brief, selective, and incomplete introduction to the rapidly growing field of *neuroeconomics*. We make no attempt to be comprehensive; the treatment here is best thought of as a brief taster course in neuroeconomics. The aim of neuroeconomics is not only to identify the brain regions that are activated during economic choices, but also to uncover the neural circuits that may suggest new economic theories, and enable new predictions. The hope is that, in due course, economic models can successfully make joint predictions about choice behavior and the relevant neural circuitry. Most of the behavioral theories discussed in this book have not required neural foundations. Whether neural foundations will successfully underpin behavioral theories in the near future cannot be answered with certainty at the moment.

Neuroeconomics is a relatively new field in economics that is also one of the fastest growing. While our understanding of the neural underpinnings of behavior continues to improve, there is still much disagreement about the neural foundations of behavior. Indeed some readers may even object to the inclusion of such a chapter in a book that has the word “Foundations” in its title. However, it would be a mistake to ignore this field, yet at the same time, it must be realized that many of the results are tentative and speculative at the moment. Given the rapid progress in the field, this state of affairs can change in the future.

We introduce the reader to the major brain areas that are of interest to neuroeconomists and give a brief outline of the two main kinds of methods in neuroeconomics: *Measurement techniques* and *manipulation techniques*. Measurement techniques measure some aspect of brain activity; examples are single-neuron recordings, positron emission tomography, and functional magnetic resonance imaging. Manipulation techniques can alter activity in specific areas of the brain, allowing for an observation of the associated behavioral change; examples include transcranial magnetic stimulation, transcranial direct current stimulation, and brain lesion studies. Each of the techniques has advantages and disadvantages, hence, data from several complementary techniques helps us to build a better picture of the neural correlates of behavior. Some of the most interesting results come from attempts to predict the direction of causality.

We only focus on the neural underpinnings of some aspects of human behavior outlined in the first four parts of the book: behavioral decision theory with a focus on prospect theory; social preferences with a focus on inequity aversion, social trust, and punishment; time preferences with a focus on two alternative models, the β - δ model, and the planner-doer model; and strategic interaction with a focus on the neural underpinnings of equilibrium and non-equilibrium reasoning, as embodied in level- k and CH models. Finally, we briefly comment on pharmacoeconomics, which studies the change in human behavior when one interferes with the neuromodulator system, with a focus on the effects of oxytocin.

CHAPTER 23

Neuroeconomics

23.1 Introduction

It is useful to begin with a definition of neuroeconomics. Camerer et al. (2015) define neuroeconomics as follows.

Neuroeconomics is the study of the biological microfoundations of economic cognition and economic behavior. Biological microfoundations are neurochemical mechanisms and pathways, like brain regions, neurons, genes, and neurotransmitters.¹ Economic cognition includes memory, preferences, emotions, mental representations, expectations, anticipation, learning, perception, information processing, inference, simulation, valuation, and the subjective experience of reward. In general, neuroeconomic research seeks to identify and test biologically microfounded models that link cognitive building blocks to economic behavior.

Neuroeconomic insights have, so far, not been critical to the development of behavioral economics, but they have already begun to enrich it, and may do so even more, in due course. For instance, the work of Kahneman and Tversky that laid the seeds of behavioral economics, and much subsequent work, did not need neural foundations. The System 1 and System 2 distinction (see Part 7) that Kahneman (2011) uses to motivate judgment heuristics is just a useful metaphor to organize behavioral insights; there is no presumption that distinct brain regions correspond to these systems.

The advocates of neuroeconomics suggest several potential benefits of their research, many of which are necessarily speculative at this early stage (Glimcher et al., 2009; Glimcher and Fehr, 2014; Camerer et al., 2005; Camerer et al., 2015). The most ambitious hope is that, in due course, economic models will make joint predictions about the choices made (based on observables), and the neural mechanisms that give rise to these choices. Such models can then be tested using either traditional methods in economics or neuroeconomic methods, or both. Furthermore, such theories may make novel predictions that theories based on choice behavior alone might not be able to make.

Other ambitious hopes from neuroeconomic research include direct measurements of human welfare, and the use of neuroeconomic data to spur the development of newer models in economics. The common perception is that neuroeconomics tries to identify particular brain areas for each task. However, as Camerer et al. (2015) note, the main thrust of neuroeconomics is the

¹ Neurotransmitters are molecules that carry neurochemical signals from one neuron to another.

hope that “common patterns of [brain] circuitry will emerge which will inform debates about the computations that are performed, and suggest new theories of behavior and new predictions.”

The literature on neuroeconomics, and the neuroscience that it draws on, is already quite vast. This chapter gives a brief account of a few selected topics from this literature. We make no pretense at being complete, comprehensive, or balanced. One justification for this brevity is that several excellent accounts written by the leaders in the field are already available (Glimcher et al., 2009; Glimcher, 2011; Glimcher and Fehr, 2014; Camerer, 2013; Camerer et al., 2015). We choose to focus only on a few studies that explore the neural foundations of some of the human behavior outlined in the first four parts of the book: behavioral decision theory, social preferences, time preferences, and strategic behavior. The hope is to give the readers a flavor of the methods and results, so as to stimulate their interest in a fuller exploration of the issues.

We omit a range of standard topics in neuroeconomics that can be more fully pursued in the references. These include estimation techniques and computational methods in neuroeconomics, and the neuroeconomics of learning and emotions. We also omit a discussion of the debate between those who argue that economics should only consider choice data and others who argue for a broader view of economics; the thrust of these issues has already been considered in Part 8. In the most restrictive view of economics, all non-choice data is inadmissible, particularly neuroeconomic data; this position is most starkly outlined in Gul and Pesendorfer (2008). For a view offering strong support for a broader view of economics, particularly encompassing neuroeconomic data, see Camerer (2008).² Several key arguments offered by Camerer such as the ability of neuroeconomic data to predict human choice and estimate behavioral parameters of interest are illustrated in this chapter.

The first two sections give an introduction to the basics. Section 23.2 gives a brief introduction to brain areas and regions of interest to neuroeconomists. Section 23.3 gives an introduction to alternative methods of measuring brain activity and identifying neural circuitry responsible for economic decisions. Each method has its own advantages and disadvantages, hence, a more complete picture typically emerges by comparing the results from different complementary techniques. We separate the techniques into two classes: *Measurement techniques* and *manipulation techniques*, and evaluate them on the following three criteria. *Temporal resolution* (time frequency of recording of brain activity), *spatial resolution* (precision with which one can narrow down brain activity to a particular area), and the *degree of invasiveness*.

Measurement techniques typically establish correlations between economic behavior and brain activity. This can take the form of single unit recording of neuronal activity, or forming 2D or 3D images of brain activity. We mainly discuss *single-neuron recording*, *electroencephalography* and *magnetoencephalography*, *positron emission tomography* (PET), and *functional magnetic resonance imaging* (fMRI). Manipulation techniques alter (either increase or inhibit) activity in a part of the brain and then study the resulting change in human behavior. The three main manipulation techniques that we describe are: *repeated transcranial magnetic stimulation* (rTMS), *transcranial direct current stimulation* (tDCS), and studies of *brain lesion patients*.

Section 23.4 considers the neuroeconomics of risky decisions with a focus on prospect theory. We begin in Section 23.4.1, with a brief discussion of the neural correlates of value and probability (Glimcher, 2011; Fehr and Rangel, 2011). We ask if the brain computes something like a utility from simple outcomes that is decisive in choosing among the outcomes? The bulk of the studies that answer this question have been performed using simple food items or simple tasks such as judging the quality of an item. It is not always clear if these results will carry over to more

² See also Bernheim (2008) who offers a constructive view on the role of neuroeconomics.

complex economic decisions such as portfolio choice (Camerer, 2013). The distinction that we drew in Part 7 between *decision utility* and *experienced utility* turns out to be of critical importance in these experiments and is shown to have a good neural basis. Decision utility appears to be associated with the ventromedial prefrontal cortex while experienced utility is associated with the orbitofrontal cortex and the nucleus accumbens (these and other brain regions are described in more detail in Section 23.2). We also comment very briefly on the neural basis of mean variance utility, and the axiomatic approach to modeling brain activity.

Section 23.4.2 considers reference dependence. The amygdala has been implicated in fMRI studies of framing effects that are consistent with reference dependence (De Martino et al., 2006; Roiser et al., 2009). However, contrary to these claims, amygdala lesioned patients also exhibit framing effects (Talmi et al., 2010). In another implication of reference dependence—the WTA/WTP disparity—the neural correlates of WTA and WTP lie, respectively, in the ventromedial prefrontal cortex and the lateral orbitofrontal region (Knutson et al., 2008; De Martino et al., 2009). Section 23.4.3 considers loss aversion. Tom et al., (2007) find that in mixed lottery experiments, there is increased activity in the ventromedial prefrontal cortex, medial orbitofrontal cortex, and the anterior cingulate cortex when subjects face increased gains. However, when they face increased losses there is simply reduced activity in the same brain areas, that is greater in absolute value compared to gains; this is a form of *neural loss aversion*.

In particular, risky losses do not induce activity in the expected brain areas such as the amygdala and the insula, which are associated with fear, anxiety, and emotions. However, contrary to these results, patients with brain lesions in the amygdala show no loss aversion (De Martino et al., 2010). Several reasons may account for this difference, which include the range of losses used in the experiments, and the role of the amygdala as an intermediate output area (Jenison et al., 2011). Finally, Section 23.4.4 considers non-linear probability weighting. Prominent brain regions that appear to be associated with non-linear probabilities are the striatum and the cingulate gyrus (Paulus and Frank, 2006; Berns et al., 2008; Hsu et al., 2009).

Section 23.5, the longest section in this chapter, considers the neural foundations of social preferences. Section 23.5.1 considers the neural circuitry associated with social rewards. There appears to be a common neural activation for private and social rewards (Izuma et al., 2008; Smith et al., 2010). Based on identical underlying tasks, in private and social rewards treatments, areas of ventromedial prefrontal cortex and the ventral striatum, encode signals for both kinds of rewards (Lin et al., 2012). Section 23.5.2 shows that inequity averse preferences have strong neural foundations (Tricomi et al., 2010). Indeed the *neural strength* of these preferences appears to be greater than the *behaviorally observed* inequity averse preferences. The ventral striatum and the ventromedial prefrontal cortex are important brain areas that appear to encode these preferences. Neural data can also differentiate between inequity averse preferences that arise from cases where the income distribution is generated from endogenous or exogenous effort (Cappelen et al., 2014).

Section 23.5.3 considers the neural basis of social punishment and trust. In a trust game with punishment that varied intentions and punishment costs, brain activity in the caudate nucleus under free punishment predicted the level of punishment when it was costly (De Quervain et al., 2004). Furthermore, the punishment decision involved a trade-off between reduction in personal rewards and the satisfaction arising from punishing. Hence, it activated the median orbitofrontal cortex, which has been implicated in similar trade-offs. These results speak to an important puzzle in classical game theory that arises from the empirical results from public goods games with punishment: Why do people punish when bygones are bygones? The neural answer is simply that people derive satisfaction from the act of punishing norm offenders; such punishment activates the brain's reward circuits, hence, punishment confers hedonic benefits. Players who are treated

unfairly in ultimatum games show increased activation of the emotional regions of the brain such as the anterior insula (Sanfey et al., 2003). Thus, unfair offers appear to elicit an emotional or affective response.

The differing neural responses in punishing ingroup and outgroup members provide a neural basis for *parochial altruism* (Baumgartner et al., 2012). In punishing outgroup members there is activation in the right orbitofrontal gyrus, right lateral prefrontal cortex, and the right dorsal caudatus. However, in punishing ingroup members, subjects must generate rationales for mitigating punishments. The neural circuitry that is activated in this case is the temporoparietal junction that is often implicated in neural research on the *theory of the mind*, i.e., mentalizing about the motives, strategies, and beliefs about others. Brain manipulation techniques show that disrupting the right lateral prefrontal cortex reduces rates of rejection in ultimatum games (Knoch et al., 2009). Individuals with lesions in the ventromedial prefrontal cortex exhibit different social preferences relative to non-lesioned subjects (Koenigs and Tranel, 2007; Krajbich et al., 2009). The amygdala has also been implicated in the evaluation of fairness in ultimatum games (Koscik and Tranel, 2011; Gispic et al., 2011). The size of the amygdala is also correlated with the size of one's social network (Bickart et al., 2011).

Section 23.5.4 considers the neural basis of selected human virtues. Neural activity in deceptive and honest players is markedly different when they play the same game. Activity in the anterior cingulate cortex and the right interior frontal gyrus in an anticipation task can successfully predict the extent of deception that an individual will engage in (Baumgartner et al., 2009). Section 23.5.5 shows that there is a neural basis to some of the constructs used in psychological game theory, such as the feeling of *guilt* in trust games when receivers return low amounts to trustors (Chang et al., 2011).

Section 23.6 considers the neuroeconomics of time preferences. Support is found for two different models of time discounting considered in Part 3, (β, δ) quasi-hyperbolic discounting (Laibson, 1997), and the planner–doer models (Thaler and Shefrin, 1981; Fudenberg and Levine, 2006). Separate areas of the brain are activated when subjects express a preference for immediate rewards (β -areas), relative to delayed rewards (δ -areas). β -areas include nucleus accumbens, subgenual cingulate cortex, and medial orbitofrontal cortex, while δ -areas include posterior cingulate cortex, posterior parietal cortex, and the insula (McClure et al., 2004, 2007). However, these results have been questioned by contrary empirical evidence (Glimcher et al., 2007; Kable and Glimcher, 2007). In the planner–doer models, fMRI evidence indicates that the dorsolateral prefrontal cortex plays the role of a patient long-run planner (Hare et al., 2009). Temporary disruption of the dorsolateral prefrontal cortex using brain manipulation techniques leads subjects to choose more immediate rewards, which is consistent with the planner–doer approach (Figner et al., 2010).

Section 23.7 considers the neuroeconomics of strategic interaction. A major thrust of this literature is to uncover neural activity that illuminates how players develop and use a theory of the mind. In other words, how do players infer the likely strategies, beliefs, and mental models of others? One method of revealing a theory of the mind is to check for the contrast when humans play against human opponents and against computers. Different neural circuits appear to be engaged in the human–human and human–computer pairings (Gallagher et al., 2002; McCabe et al., 2001; Coricelli and Nagel, 2009).

Another issue is the differing neural activity of “equilibrium types” who play in accordance with classical equilibrium strategies, and the behavioral types who play non-equilibrium strategies (e.g., level- k players). Since beliefs and choices are aligned for equilibrium subjects, as hypothesized their brain activity occurred in largely similar areas, but for non-equilibrium subjects, brain activity was much more dispersed (Bhatt and Camerer, 2005). Within level- k subjects,

differing neural activity has been found among high and low level reasoners in a p-beauty game; decisions of high level reasoners were associated with relatively greater activity in the medial prefrontal cortex (Coricelli and Nagel, 2009). Two brain areas were particularly highlighted for high level reasoners; one is associated with taking a third person perspective while the other is involved in thinking about the strategies of like-minded other people. However, results from other studies have not always supported these conclusions (Kuo et al., 2009). Finally, we consider the neural basis of deception that is of critical importance in game theory. Emerging evidence supports level- k reasoning in simple, two-player, buyer–seller bargaining experiments where deception is possible (Bhatt et al., 2010, Bhatt et al., 2012).

Finally, in Section 23.8, we briefly touch upon the expanding literature on the pharmacology of economic and social decision making. The main aim of this research is to interfere with the body's *neuromodulator systems* by altering the levels of neuromodulators (dopamine, serotonin, norepinephrine), or hormones such as testosterone and oxytocin. We focus exclusive attention on the role of oxytocin as an illustration of this line of research. Readers interested in an authoritative and accessible survey can consult Crockett and Fehr (2014). Oxytocin has been shown to increase the level of trust that trustors place in trustees in trust games. However, the effect of oxytocin is very selective in influencing only the trust level of trustors; it does not alter the general level of optimism, risk perception, moods, or social trust (Kosfeld et al., 2005). Subsequent studies have enriched our understanding of oxytocin by examining the effects of interpersonal relations, generosity, cooperation, perceptions of trustworthiness, attractiveness, approachability, attachment, and socio-emotional responses such as envy and social aversion (Bartz et al., 2011). In the trust game, it is payoff improving for both parties if the trustor trusts more, hence, oxytocin plays a positive role. However, in other contexts, oxytocin may reduce the ability to recognize deceptive behavior and reduce one's payoffs (Israel et al., 2014).

23.2 A brief introduction to the brain

In this section, we briefly review some features of the human brain that are useful in the study of neuroeconomics; for an authoritative source for economists, see Glimcher and Fehr (2014). *Modularity* refers to the idea that the brain is composed of discrete parts that perform their own specific computations and then pass on the results to other brain areas for further processing and analysis (Gazzaniga, 2009). *Plasticity* refers to the ability of the brain to evolve and learn as the environment changes. *Automaticity* refers to the idea that we engage in certain behaviors/deliberations without being consciously aware that we are doing so.

We might often be unaware of the cause of our own behavior on account of automaticity. As Camerer et al. (2004) put it: “the human brain is like a monkey brain with a press secretary, which puts a coherent spin, and sometimes a wrong one, on what the monkey brain did.” The brain is sometimes “too good” for its own good at *sense making*. Experimental evidence, reviewed elsewhere in the book, indicates that people claim to spot regularity in completely random patterns such as stock market prices. Elsewhere, hot hands in basketball, the gambler's fallacy, the hindsight bias, and regression to the mean may also, perhaps, be guided by a similar desire for sense making even when the underlying data does not support these inferences based on classical statistics (see Part 7).

The flip side of these features is that the brain allocates attention selectively, which protects capacity limits in perception and cognition by biasing/prioritizing salient or relevant representations over others; recall, for instance, the phenomenon of confirmation bias discussed in Part 7.

In an interesting paper by Simons and Chabris (1999) with the title “Gorillas in our midst: Sustained inattention blindness for dynamic events”, subjects watch a video of six individuals playing basketball. The experimental task is to count the number of passes made by each team that wears a different jersey color. Within a minute, a gorilla walks into the center of the game and stays there for 9 seconds. At the end of the experiment, half the subjects in the experiment have no recollection of seeing a gorilla in the video clip.

When humans departed from their nearest primate relatives (about 6 million years ago), the human brain was not redesigned from scratch. Evolution, because it is inherently conservative and cautious, added neural regions imparting significant cognitive ability to humans, such as the prefrontal cortex, on top of the existing brain structures. Going back even further into evolutionary history, at the core of the human brain is a primitive mammalian brain, the limbic system, that controls emotions such as fear, anger, disgust, etc. Thus, experiments conducted on other mammals are of potential value for an understanding of the human brain, particularly those conducted on our nearest primate relatives. For instance, rats become addicted to almost all the substances that humans become addicted to (Camerer, 2007). Wrapped around this primitive mammalian brain is the neocortex (also called the cerebral cortex, which corresponds to people’s perception of “gray matter”) that accounts for about 5/6 of the human brain. The neocortex is responsible for human cognition, language, and speech. For instance, it allows humans, unlike other animals, to think about the longer-term consequences of their actions.

With this brief introductory background, let us consider the various brain regions that are important in neuroeconomics.³ The four main regions of the brain are shown in Figure 23.1. Each of the four lobes is responsible for functions that are either independent or overlapping. The *frontal lobe* is associated with cognition and logic. The *parietal lobe* is associated with recognition and movement; the *occipital lobe* with the processing of visual information; and the *temporal*




Figure 23.1 A classification of the brain into four lobes: frontal, parietal, occipital, and temporal.

³ For textbook treatments see Breedlove et al. (2014) and Bear et al. (2007). For a brief and clear introduction, see Glimcher and Fehr (2014). © Shutterstock

lobe with memory and auditory stimuli. Neuropsychological studies of aphasia (a term used for a collection of speech disorders) suggest that speech production and comprehension might be attributed to distinct frontal and temporal substrates. The lobes come in pairs, one on each side of the brain; language is associated with a left hemisphere function. The lobes are connected by neurons that enable the brain to simultaneously engage its various parts to generate a composite response to external stimuli.

Example 23.1 (Camerer et al., 2004): “When an automated insurance broker calls and says, ‘Don’t you want earthquake insurance? Press 1 for more information’ the occipital lobe ‘pictures’ your house collapsing; the temporal lobe feels a negative emotion; and the frontal lobe receives the emotional signal and weighs it against the likely cost of insurance. If the frontal lobe ‘decides’ you should find out more, the parietal lobe sends signals to your finger to press 1 on your phone.”

The following terminology is useful in identifying various brain regions; see Figure 23.2. We quote the following description from Zak (2004, p. 1739): “Terms for locations of brain regions include: dorsal (top, from the Latin for back); ventral (bottom facing the central axis, from the Latin for belly) or basal; rostral (front, from the Latin for beak) or anterior; caudal (back, from the Latin for tail) or posterior; superior (towards the top); inferior (towards the bottom);




Figure 23.2 A convenient classification of brain regions. © Shutterstock

medial/mesial (middle); lateral (away from midline); and orbital (above the eyes, from the Latin orbital meaning eye sockets).” In evolutionary terms, ventral brain regions are typically phylogenetically older in comparison to dorsal and rostral regions.

It is sometimes helpful to separate the brain into three divisions. These are the forebrain or *telencephalon*, midbrain or *mesencephalon*, and the hindbrain or *brainstem*; see Figure 23.3. The main focus in neuroeconomic research has been on the telencephalon. In turn, the telencephalon is typically divided into three main regions. These are the *cerebral cortex*, *basal ganglia*, and the *thalamus*; see Figures 23.3, 23.4. Our main focus will be on the cerebral cortex and the basal ganglia. Of these, in evolutionary terms, the cerebral cortex is the most recently evolved structure, while the basal ganglia, shown in Figure 23.4, is the oldest (recall that the more ventral structures are typically phylogenetically older).

The two regions in the basal ganglia, *caudate* and *putamen*, together comprise the *striatum*, an important brain region in neuroeconomics. The striatum, an input area of the basal ganglia, receives inputs from the frontal cortex and mainly sends its inputs to the output areas in the basal ganglia, the *globus pallidus* and the *substantia nigra pars reticula*. Another area in the basal ganglia that plays an important role in neuroeconomics is the dopaminergic system that is comprised of neurons that release dopamine in the *ventral tegmental area* and the *substantia nigra pars compacta*.

The brain is covered with an outer layer of neural tissue called the *cerebral cortex*; in common usage this is also referred to as *gray matter*. Dense axons under the gray matter are sometimes known as *white matter*. The most recent evolutionary part of the cerebral cortex is the neocortex, which is a six layered structure shown in Figure 23.5. Each layer can have different kinds of neurons, and has its own functional specialization, but overall the neocortex is fairly uniform. For instance, layer 4 (the internal granular layer) is typically an input region in the cortex and layer 5 comprises mainly the pyramidal neurons that connect, through their axons, with other regions in the cortex. The phylogenetically more ancient part of the cerebral cortex is the hippocampus, which is a three layered structure. In order to increase surface area within the confines of the




Figure 23.3 A classification of the brain into three regions: forebrain, midbrain, and hindbrain.




Figure 23.4 Structure of the basal ganglia and the approximate position of its various parts. © Shutterstock




Figure 23.5 Different layers of the cerebral cortex. © Anna Hine, Shutterstock

THE LIMBIC SYSTEM




Figure 23.6 The limbic system of the brain. © Shutterstock

skull, the cerebral cortex is folded; ridges are known as gyri (singular gyrus) and grooves are known as sulci (singular sulcus).

Figure 23.6 shows a sketch of the limbic system, which is a part of the forebrain; some of the important limbic structures for neuroeconomic research are the hippocampus, amygdala, fornix, mammillary body, and cingulate gyrus. The amygdala is a part of the telencephalon, yet it is not considered to be a part of the basal ganglia or of the cerebral cortex. The amygdala is activated when individuals experience fear or a perceived threat; it receives inputs from several sensory areas and feeds into the hypothalamus. The hippocampus, a three layered structure, is a part of the ancient cortex; it is associated with long-term memory and plays a central role in neuroeconomic studies of learning and memory.

Korbinian Brodmann classified the brain into distinct, numbered, regions based on the differences in neuronal architecture between these regions. This classification, published originally in 1909, is shown in Figure 23.7 (Brodmann, 1909). The Brodmann areas typically perform well-defined and discrete tasks. Some Brodmann regions receive inputs from the eyes and ears, so they are referred to as *sensory areas*, while others are associated with the control of movement, so they are known as *motor areas*. We can use the Brodmann areas to identify an important region of interest for neuroeconomists, the *prefrontal cortex*, abbreviated as PFC; this is neither a sensory, nor a motor, area and is sometimes referred to as an *association area*. The PFC comprises of the following three regions. *Dorsolateral prefrontal cortex* abbreviated as dlPFC (areas 9, 46); *ventromedial prefrontal cortex*, abbreviated as vmPFC (area 45); *orbitofrontal cortex* OFC (area 10).

Different parts of the brain communicate with each other through special cells called *neurons* or *nerve cells*. There are many kinds of neurons in the human body. For instance, the motor neurons receive information from the spinal cord and are responsible for muscle contractions. The sensory neurons respond to the senses, e.g., touch, and transmit signals to the spinal cord. Figure 23.8 shows two neurons and the basic structure of a neuron; multiple arrows represent the direction of the flow of information.




Figure 23.7 Brodmann areas of the human brain.




Figure 23.8 The structure of a neuron. © Shutterstock

The bulbous cell body, or *soma*, encloses the nucleus; different structures inside the cell body process sugars and oxygen to provide the energy that powers the neuron. Several thin structures that can branch several times, called *dendrites*, emanate from the soma; these are the input structures of the neuron and receive information from other neurons. A single, thin, output structure, the *axon*, originates from each cell body of a neuron. It transmits electrochemical information to a neighboring neuron through special structures called *synapses* that come into contact with the dendrites of other neurons. Axons can be quite long, and in rare cases, up to a meter long. Unlike dendrites, each neuron has a single axon, which nevertheless can branch several times.

Signals between neurons are sometimes known as *action potentials* or simply *spikes*. Neurons are surrounded, like all cells, by a thin *membrane* impermeable to water-based chemicals that is perforated by tiny tubes or *channels*. These tubes, also known as ion channels, allow for the exchange of sodium ions (relatively abundant outside the membrane) and potassium ions (relatively abundant inside the soma). *Neurotransmitters* that move from the synapse of one axon onto the dendrite of another axon regulate the opening and closing of the ion channels. Neurons have an electric resting potential of about -70 microvolts. This means that when the neuron is not sending any signals, the inside of the neuron is 70 microvolts less than the outside and there are relatively more sodium ions outside the neuron but relatively more potassium ions inside. Action potentials occur if the voltage in a neuron exceeds a certain threshold, typically around -50 microvolts caused by inputs from other neurons. This allows the ion channels in the membranes of neurons to be activated. Depending on the type of the neuron, the baseline firing rates vary from a few spikes per second to about a hundred spikes per second.

We shall sometimes report activity in an area of the brain by its (x, y, z) coordinates; most software packages use the Montreal Neurological Institute (MNI) coordinates. Camerer et al. (2015) succinctly explain these coordinates: “A higher positive value of y is further forward, or more anterior, in the brain; more negative values are more posterior toward the back of the brain. Similarly, x values range from the left side (most negative) to the right side (most positive), and z -values range from most negative (the inferior part or bottom of the brain) to the most positive (the superior part or top of the brain).”

23.3 An introduction to neuroeconomic techniques

Neuroeconomics uses a range of methods to gather evidence and form inferences.⁴ Many methods of neuroeconomic measurement aim to measure the flow of electrochemical information between neurons. Different techniques measure different aspects of neural activity with varying degrees of accuracy and precision. It might help to distinguish between two categories of techniques.

Measurement techniques measure some form of brain activity using, say, neuronal firing rates, or metabolic activity, when subjects in experiments engage in some relevant economic choice. Such techniques typically establish correlation but not causation between brain activity and human decisions. *Manipulation techniques*, on the other hand, alter some aspect of brain function temporarily or permanently and observe the causal effect on human decisions. It is possible to alter the neuronal firing rates in a selected region of the brain, or to alter the neurotransmitter levels in a transient manner. In animal studies, it is possible to permanently damage brain tissue and then observe which aspects of animal behavior change. Human subjects may have *brain*

⁴ For an excellent and concise introduction, see Ruff and Huettel (2014) that we draw upon below.

lesions caused by infections or by accidents that are also used in manipulation techniques. There is a degree of complementarity between the two techniques because measurement techniques may identify brain areas that can be manipulated later to establish causality.

23.3.1 *Criteria for evaluating neuroeconomic measurement techniques*

Alternative methods of measuring brain activity have their pros and cons. Hence, alternative techniques play a complementary role in building a richer picture of the relation between brain activity and economic decisions. Neuroeconomic methods of measurement may be evaluated on the following three main criteria.

(1) *Temporal resolution* is the frequency with which brain activity is recorded. The best techniques that measure neuronal activity directly, by measuring the flow of ions in brain tissue (electrophysiological measurements), achieve millisecond precision. Indirect measures of neuronal activity, such as measuring changes in blood oxygen levels in the brain, have intermediate temporal resolution that ranges from seconds to minutes. Brain manipulation techniques, say, drug induced temporary damage to brain tissue, have the lowest precision that ranges from minutes to days (the exception is event-related transcranial magnetic stimulation).

(2) *Spatial resolution* is the precision with which one can narrow down brain activity to a particular region of the brain. Naturally, techniques that place electrode sensors directly in the brain, typically in animals, achieve the highest spatial resolution and may measure activity directly for a particular neuron. Functional neuroimaging techniques that are extremely popular (see below) achieve spatial resolution of millimeters to centimeters. Finally, techniques that measure electrical signals of brain activity non-invasively, say, by measurements taken on the skull, achieve the poorest spatial resolution that ranges from centimeters to the entire brain.

(3) *Degree of invasiveness* refers to the extent to which one may have to interfere with brain processes in order to measure brain activity. Examples include knocking-out areas of the brain temporarily or permanently; the latter is typically possible only with animals. In humans, invasive techniques may be used prior to neurosurgery.

Our focus is not on the estimation and computation techniques of neuroeconomics. The interested reader can begin with Busemeyer and Diederich (2014) and follow the references therein.

The older methods measure *psychophysical* activity, i.e., the involuntary reaction to stimuli. These methods include, for instance, pupil dilation, skin conductance, and the heart rate. We now focus on some of the modern methods of measurement.

23.3.2 *Single unit recording*

Using fairly invasive methods, one can measure the electrical activity in a single targeted neuron. An electrode, typically a very thin insulated wire, is inserted in the neural tissue surrounding a single neuron; this is able to measure the neuron's electrical activity. Such techniques are quite invasive because they involve the opening of the skull, hence, they are typically used only on animals. However, they are also sometimes used on humans, who for some medical reasons, e.g., epilepsy, need to have an electrode inserted in their brain. Since animals share the affective rather than the deliberative faculties of the brain with humans, this evidence speaks more to the emotional rather than cognitive aspects of behavior. The advantage of these techniques is that they can target their location very precisely.

Most neurons are extremely small, so it might not always be possible to distinguish the firing rate of a neuron from other neurons in close proximity. However, improvements in techniques

have reduced this problem. Single unit recording has very high temporal and spatial resolution, hence, it has enabled the identification of some of the core brain functions. Since the method is highly targeted at specific neurons, and neurons differ in their firing rates and their responses to different inputs, it helps to build a picture of brain functions that is unmatched in comparison to the other methods. The main drawback of the technique is the degree of invasiveness, and the rate of data collection is very slow because data is collected for a single neuron at a time. Animal colonies for experimental purposes are costly to maintain, so the technique is relatively expensive. Finally, the technique is also not very suitable when analyzing complex brain functions that involve interactions among several parts of the brain or neural coding at the population level.

23.3.3 *Electroencephalography and magnetoencephalography*

The *electroencephalogram* (EEG) measures brain activity, externally, say, by electrodes positioned on the scalp, or positioned internally, while subjects make some decisions. EEG at the scalp measures changes in the polarity of synchronously firing populations of neurons. The advantages of this technique are as follows. (i) It measures brain activity directly, rather than using some proxy such as blood flow or the level of metabolism. (ii) It has a high temporal resolution that allows one to study dynamic processes in the brain. (iii) It is relatively inexpensive in terms of equipment, say, relative to the cost of an MRI scanner (although the set-up time to position the electrodes correctly on the scalp is high). The disadvantage is that it has low spatial resolution, i.e., it measures aggregate brain activity in some region of the brain and cannot always localize the source of that activity.

An alternative to measuring electrical activity is to measure the magnetic fields associated with electrical activity in the neurons; this technique is known as *magnetoencephalography*. It has similarly high temporal resolution as electroencephalography but does not suffer from the problem of weakening of the electrical signal on passing through the skull, which is the case with electroencephalography. However, the magnetic fields are relatively weak and the equipment is relatively expensive, hence, not many studies in neuroeconomics have used it.

23.3.4 *Positron Emission Tomography*

Positron Emission Tomography (PET) relies on early work done in the 1950s and falls within the ambit of *functional neuroimaging*, i.e., producing 2D or 3D maps of brain activity. The method is to inject a weak, short-lived, radioactive isotope into the body of a subject in the experiment. Once enough time has been given for the solution to permeate the tissue of interest, the radioactive solution decays and emits positrons. The emitted positrons collide with electrons within a short distance and emit gamma rays. A *PET scanner* records the gamma rays and assigns to them a likely region of origin in the brain; this gives rise to a *PET image*.

The disadvantages of this technique are that it is invasive, relatively expensive, and its temporal resolution is relatively low (it shows average brain activity over an entire experiment), while its spatial resolution is intermediate. However, there are some advantages over other functional neuroimaging techniques such as fMRI, which relies on blood flow to correlate with neural activity. For instance, PET provides a more direct indicator of neural activity and, by creating customized molecules that are attached to radioactive isotopes, it can provide fairly precise chemical information.

23.3.5 *Functional magnetic resonance imaging*

Magnetic resonance imaging (MRI) is a body imaging technique that is about 35 years old. It can provide detailed information about any part of the human body in any plane. The idea is to expose the body to a powerful magnetic field so that the protons in human cells line up with the direction of the magnetic field. A second electromagnetic field is then switched on which enables the protons to absorb the energy of this field. When this secondary electromagnetic field is switched off, the protons emit the energy that they have absorbed. MRI contrasts are then created based on the speed with which protons realign with the magnetic field once the radio frequency pulse offsets. Protons from different tissues in the body emit energy in different ways, allowing the researcher to form detailed internal images of the body.

A newer version of the MRI is the *functional magnetic resonance imaging* technique (fMRI) that has been developed since the early 1990s and is very popular.⁵ Electrical activity in the brain requires the firing of neurons. However, neurons do not have their own internal source of energy. Hence, when neurons need to fire, they absorb oxygen from the surrounding blood stream. This increases the blood flow towards those neurons (a process called as *hemodynamics*) after a delay of between 1–5 seconds. The response of oxygenated blood to a magnetic field is different from that of deoxygenated blood; this is known as the blood-oxygenation-level-dependent (BOLD) contrast. The BOLD contrast allows a magnetic field to pick out the neurons that are firing (say, in response to some stimuli). The fMRI interface modifies the MRI technique by installing an MR-compatible monitor or a head mounted display, which subjects use to register their choices and decisions by using a joystick or pressing buttons. Arguably the biggest distinction between fMRI and MRI is that fMRI measures changes in the BOLD signals while MRI produces high resolution structural images of the brain's morphology.

The advantage of the fMRI technique is that it is non-invasive (like the EEG) and it has good spatial resolution (3–6 millimeters). The disadvantage, relative to the EEG, is that it has poorer temporal resolution (several seconds), although it is still relatively good. This is because, while the EEG measures electrical activity, fMRI measures the differential response of oxygenated and deoxygenated blood. Other shortcomings are that (1) it is relatively expensive, (2) subjects face many repeated trials to uncover the underlying brain processes, even for essentially static games, and (3) the MRI scanner is loud and it presents a confined environment. The unit of measurement in fMRI trials is typically a *voxel* that localizes brain activity using three-dimensional coordinates. The typical fMRI voxels are 27 mm^3 , i.e., a cube whose sides are of 3 mm length. Each voxel can have more than a million brain cells.

The fMRI technique allows one to correlate an external stimulus with hemodynamic activity in a certain part of the brain. It does not directly establish causality, although it can do so indirectly by running more than one treatment. One can use observations on neural activity from some treatment, and then make predictions of neural activity that will arise in another treatment, also known as “out of treatment forecasting.” Fehr and Camerer (2007) give the following examples (p. 423): “For example, individual differences in caudate nucleus activation when punishment is costless for the punisher can be used to predict how much individuals actually pay for punishment when it is not costless. Likewise, individual differences in striatal activity in the condition in which donations are forced can be used to predict subjects’ willingness to donate money to charities in the condition in which donations are voluntary.”

⁵ For a textbook treatment of the technique, see Huettel et al. (2009).

23.3.6 Manipulation techniques

The techniques discussed above, establish correlation between human decisions and activity in specific regions of the brain. This does not establish causality (unless one uses out of treatment forecasting). We now move into the category of *manipulation techniques* that either temporarily or permanently disable a brain area or heighten activity in that area, and then study the resulting behavioral implications. Brain lesions and neuropharmacological interventions provide an opportunity to apply these techniques in humans. The basic principles have been understood for a long time, initially by using invasive studies for animals (Fritsch and Hitzig, 1870) and humans (Penfield and Rasmussen, 1950).

Most studies nowadays use non-invasive methods; the two most popular methods are *transcranial magnetic stimulation* and *transcranial direct current stimulation*. These studies take a step towards establishing causality but none of the methods, on their own, can conclusively establish causality. A great difficulty in establishing causality in neuroeconomic studies is that the same brain region is typically activated in a range of economic situations. Given the current state of knowledge in neuroeconomics, it is very difficult to establish causality from an activation of a brain region (Poldrack, 2006). One typically uses other supporting evidence to increase faith in the degree of causality; this includes moods, skin conductance measures, and response times; see Hutzler (2014) for a view that is more sympathetic to the ability of neuroscience data to establish causality.

TRANSCRANIAL MAGNETIC STIMULATION

Transcranial magnetic stimulation (TMS) is a non-invasive manipulation method.⁶ As noted above, neurons have an electric resting potential of -70 microvolts. TMS takes the form of stimulating areas in the brain by using the principle of electromagnetic induction. A weak electrical current is induced in the neurons belonging to a selected part of the brain by rapidly changing magnetic fields. This is achieved by placing a looped copper coil on the part of the scalp directly over the brain region that is intended to be stimulated; not all locations in the brain are equally easy to stimulate.

Modern studies typically use repetitive TMS (rTMS) pulses separated by only a few milliseconds to achieve high temporal resolution. One may apply rTMS pulses in one of two ways: in an *online fashion* or in an *offline fashion*. When applied in the offline fashion, rTMS is applied at low temporal frequency for several minutes just prior to subjects' participation in the actual experimental task. This induces lowered cortical excitability in a particular brain region that lasts for 10–30 minutes. When applied in the online fashion, rTMS pulses serve to heighten cortical activity. If part X of the brain is thought to perform some task T, then TMS can be used to momentarily enhance or knock out part X. If this affects performance on task T, then it is suggestive of the causality of brain region X for task T.

A major advantage of the TMS technique is that it is non-invasive and has high spatial and temporal resolution. A disadvantage is that it is more difficult to be precise in targeting a specific brain area, particularly the deeper brain areas. In the offline studies, it is not always clear how long the rTMS pulse should be applied, prior to the experiment. Furthermore, some subjects might find the procedure noisy and uncomfortable.

⁶ See Wassermann et al. (2008) for a book length treatment.

TRANSCRANIAL DIRECT CURRENT STIMULATION

Under *transcranial direct current stimulation* (tDCS), two electrodes are placed on the scalp and an electric current is passed through them. This affects the voltage in the membranes of the neurons along the path of the current and, thus, their firing rates, which can be measured. One of the electrodes is positively charged (anodal) and the other is negatively charged (cathodal).

There are two main types of tDCS simulations. Under *anodal* tDCS, the neurons are upregulated, or stimulated, while under *cathodal* tDCS, they are downregulated. A third kind of simulation, *sham simulation*, applies only a momentary current for the same length of time as the other two simulations; subjects are not told that they are in a sham simulation treatment, hence, this acts as a placebo treatment.

The advantages of tDCS are that it is relatively inexpensive, not noisy, and unlike TMS, it does not have unpleasant tactile sensations. One can run tDCS experiments simultaneously with a group of subjects (not possible under many other methods, such as fMRI). This feature is very useful in running a wider range of experiments, particularly those that involve social preferences.

BRAIN LESION STUDIES

One may study individuals who exhibit behavioral deficits following damage to some area of the brain (brain lesions); this is part of the subject matter of *neuropsychology*. The classic study is based on a grievous accident suffered by a Vermont Railway worker, Phineas Gage, in 1848.⁷ In an injury caused by a dynamite explosion, a tamping iron rod went through the side of his face, through the back of his left eye and broke the top of his skull; see Figure 23.9.




Figure 23.9 A drawing showing the skull of Phineas Gage, after a tamping iron went through it. © Everett Collection Historical/Alamy Stock Photo

⁷ An excellent account of the incident can be found on the wikipedia page: http://en.wikipedia.org/wiki/Phineas_Gage. © Everett Collection Historical / Alamy Stock Photo

Phineas Gage survived the accident but the destruction to his left prefrontal cortex significantly altered his behavior. Post-accident, he exhibited lack of foresight, fickle changes of mind, bad temper, impulsiveness, and several kinds of emotion deficits that were previously absent in his behavior (Harlow, 1868).⁸ A case study of this accident contributed significantly to an understanding of the causal role of the prefrontal cortex.

Brain lesions can occur in many ways, for instance, surgical removal of tumors, brain infections such as meningitis or encephalitis, and exposure to neurotoxins such as excessive alcohol. The challenge is to pinpoint, as accurately as possible, the exact location of the brain lesion. In *double dissociation studies*, one establishes that two groups of patients with different brain lesions perform differently in two distinct tasks. This is an important component in establishing causality. Brain lesion studies are subject to two disadvantages. First, one typically does not have data on patients prior to their brain lesions. Second, as the brain reorganizes in the period following the brain lesion, as it appears to have done in the later years of Phineas Gage's life, one needs to separate two different effects: Those arising from the original brain lesion and those arising from brain reorganization. In animals, brain lesion studies involve more direct and invasive intervention in which the targeted brain tissue is either permanently or temporarily disabled. In a control group, an identical surgery is performed but the relevant brain areas are not disabled; clearly this is not possible with human subjects.

23.4 The neuroeconomics of risky decisions

Part 1 considered alternative theories of risk, uncertainty, and ambiguity. The behavioral evidence indicates that among the available behavioral theories, prospect theory is currently the most satisfactory alternative. For this reason, this section focuses mainly on the neural evidence for prospect theory.

23.4.1 Neural basis of value and risk

There is a large literature on the neural encoding of value for simple binary tasks that differ in the quantity or quality of a reward. For instance, choosing between a higher and lower quantity of apples, or choosing between a more or less tasty variety of apple. The aim is to show that the brain computes something like the utility of an outcome. While there is substantial merit in understanding the correlates of neural activity for simple choice tasks, it is not clear if the insights gained from this analysis carry over to more complex tasks such as deciding among complex portfolios of stocks, or among alternative pension plans (Camerer, 2013). For this

⁸ Here is the full quote from Harlow (1868, pp. 13–14), who was the doctor who treated him using pioneering methods: "The equilibrium or balance, so to speak, between his intellectual faculties and animal propensities, seems to have been destroyed. He is fitful, irreverent, indulging at times in the grossest profanity (which was not previously his custom), manifesting but little deference for his fellows, impatient of restraint or advice when it conflicts with his desires, at times pertinaciously obstinate, yet capricious and vacillating, devising many plans of future operations, which are no sooner arranged than they are abandoned in turn for others appearing more feasible. A child in his intellectual capacity and manifestations, he has the animal passions of a strong man. Previous to his injury, although untrained in the schools, he possessed a well-balanced mind, and was looked upon by those who knew him as a shrewd, smart business man, very energetic and persistent in executing all his plans of operation. In this regard his mind was radically changed, so decidedly that his friends and acquaintances said he was 'no longer Gage'."

reason, our discussion will be brief. The interested reader can consult several references for more details.⁹

A typical experiment in this literature proceeds as follows. The subjective value assigned by subjects to an object is elicited, using standard behavioral economics experiments. For instance, subjects might be asked to bid for an item under institutional rules such as second price auctions, which predict that bids and valuations for the item will be identical. While placing the bids, the brain activity of subjects is measured, say, using fMRI techniques, or single neuron recordings. One could then ask, which brain regions are particularly active, or what brain circuitry is implicated when individuals choose high value items? The items used in the analysis are typically simple food items such as snacks (Plassmann et al., 2007). The main result from these experiments indicates that value is encoded in the vmPFC. Experiments using gambles, charitable contributions, and delayed monetary rewards also support this result (Levy et al., 2010; Kable and Glimcher, 2007; Hare et al., 2009).

An important insight in this literature is to address the observed choice between options by individuals; a choice that appears apparently stochastic. For instance when individuals are repeatedly asked to make a choice between two close options, A and B, or we use a between-subjects design to study the choice between these two options, there is an element of stochasticity to the choices. One way of modeling stochasticity in choices is to use random utility models (see the introductory chapter and Part I). An alternative, and intriguing, possibility is that this stochasticity reflects noise in underlying cognitive processing (Glimcher, 2011; Woodford, 2014).¹⁰ Thinking of options A and B as two external stimuli, the choice made by an individual depends on the number of spikes produced by different populations of neurons, each tuned to respond differentially to the two stimuli. The number of spikes, in turn, is drawn from an underlying probability distribution that could be common to the two options. For binary choice data, Clithero and Rangel (2013) show that when response times are taken into account, the out of sample predictions based on such a model (the drift-diffusion model) outperform the predictions from the standard logistic model.

Recall the distinction between *decision utility* and *experienced utility* that we made in Part 7; the evidence indicates that this distinction has a neural basis. A quick summary of the literature is as follows.

(1) *Decision utility*, or *subjective expected value*, arises at the time of making choices. Different choice options give rise to different *stimulus values* (SV), which are compared to determine the option with the highest SV (Wallis, 2007; Kable and Glimcher, 2009; Rangel et al., 2008; Rangel and Hare, 2010; Rushworth et al., 2009). Furthermore, fMRI evidence shows that SV signals have been found in the vmPFC (Chib et al., 2009; Hare et al., 2008, 2009; Kable and Glimcher, 2007; Padoa-Schioppa, 2009; Plassmann et al., 2007; Plassmann et al., 2010; Tom et al., 2007). The vmPFC also appears to encode value from *appetitive goals* (e.g., tempting food when hungry) and *aversive goals* (e.g., an electric shock) (Litt et al., 2010; Plassmann et al., 2010). Single neuron studies in monkeys show that decision value signals exhibit *range adaptation*, i.e., the neuronal activity range does not depend on the range of values given (Padoa-Schioppa, 2009; Padoa-Schioppa and Assad, 2006). If these findings apply to humans, they could violate the standard assumption of independence of irrelevant alternatives.

⁹ For a book length treatment, see Glimcher (2011); for a survey for economists, see Fehr and Rangel (2011); and for a critical and insightful commentary on this literature, see Camerer (2013).

¹⁰ For a recent theoretical exploration of the relation between the two approaches, see Webb (2015).

(2) *Experienced utility* is the *experienced subjective value* (R) at the time of consumption. While the vmPFC is activated in the case of decision utility, as described above, experienced utility appears to be associated instead with the orbitofrontal cortex and the nucleus accumbens (Blood and Zatorre, 2001; de Araujo et al., 2003; McClure et al., 2004). The initial price of a product may signal the quality of a product, hence, it could also influence the experienced utility from the product. Plassmann et al. (2008) varied the original price for an identical wine after which subjects received sprints of the wine inside a fMRI scanner. They found that experienced utilities were positively influenced by the initial price and higher experienced utility was associated with greater activity in the orbitofrontal cortex.

(3) *Prediction error* (PE) is the difference between decision utility and experienced utility. PE is used to update stimulus values in models of reinforcement learning. The main brain region associated with PE in human fMRI studies is the ventral striatum (Delgado et al., 2000; Berns et al., 2001; Yacubian et al., 2006; Hare et al., 2008).

Many models in finance use mean-variance utility analysis. Evidence suggests that different brain areas encode the mean and variance of rewards. Mean rewards are correlated with activity in the striatum (Preusschoff et al., 2006) while variance of rewards, since it reflects risk, activates the insula (Mohr et al., 2010). The greater activation of the insula when presented with higher risks suggests that there is an *affective response* to risk that is consistent with the *risk as feelings* hypothesis of Loewenstein et al. (2001) (see Part 6). tDCS studies show that stimulation of the right dlPFC increases risk aversion (Fecteau et al., 2007) and disrupting the dlPFC increases risk seeking (Knoch et al., 2006).

An alternative approach to studying neural activity is to propose axioms related to the neural encoding of rewards and prediction errors (differences between decision and experience utilities). This approach has been used, for instance, to test the *dopaminergic reward prediction error hypothesis* (DRPE) (Caplin and Dean, 2008; Caplin et al., 2010). The DRPE hypothesis relates to the responsiveness of dopamine release (a neurotransmitter) from dopamine neurons as various external stimuli, such as rewards or probabilities are manipulated. An advantage of the axiomatic approach is that it makes very precise and testable predictions. The axiomatic approach also allows for a more direct and stringent test of the DRPE hypothesis because one can dispense with auxiliary assumptions based on reinforcement learning that alternative approaches have relied upon. Camerer (2013, Section 4.1) critically evaluates the axiomatic approach relative to the alternative approach of fitting ever improving functional forms as the understanding of a phenomenon improves. Camerer's (2013, p. 1167) own assessment of the axiomatic approach, at this relatively early stage, is that: "I agree that there is an added dimension of 'rigour' . . . but there is not a big leap in knowledge."

23.4.2 Reference dependence

One implication of reference dependence is the identification of differences in behavioral patterns in the domains of gains and losses. De Martino et al. (2006) exploit this implication to study the neurobiological basis of reference dependence. Twenty subjects initially received an endowment of £50 and were then asked to choose between risky options of the form $(x, p; -y, 1 - p)$ and a sure option $(z, 1)$, where x, y, z are monetary amounts and $p > 0$ is a probability. The sure outcome, z , was either framed as a gain or as a loss. For instance, in the *gain frame*, the sure option could take the form "keep £20 of the £50." In the *loss frame*, the equivalent sure option is described as "lose £30 of the £50." Subjects chose the risky option in 42.9% of the cases in the gain frame and 61.6% of the cases in the loss frame; these proportions are different from 50% at the 5% significance




Figure 23.10 fMRI results for the gain and loss frames.

Source: From De Martino, B., Kumaran, D., Seymour, B., and Dolan, B. J. (2006). "Frames, biases, and rational decision-making in the human brain. *Science* 313: 684–7. Reprinted with permission from AAAS.

level. Thus, as expected, under prospect theory, decision makers are risk averse in the gain frame and risk seeking in the loss frame (the authors do not check for the *fourfold classification of risk* under prospect theory; see Part 1).

Decisions between lotteries were made inside an fMRI scanner. Thus, brain activity was monitored when in the gains frame subjects chose the sure option (G_{sure}) or the risky option (G_{gamble}). Likewise in the loss domain they were monitored when they chose either the sure option (L_{sure}) or the risky option (L_{gamble}). The main contrast was between activation in brain areas when subjects chose in accordance with their predominant behavioral tendency ($G_{sure} + L_{gamble}$) or against it ($G_{gamble} + L_{sure}$).

The main results are shown in Figure 23.10. Panel A shows activation in the amygdala when subjects choose according to their predominant behavioral tendencies (sure outcome in gains and risky outcome in losses). Panel B shows the percentage signal change for peaks in the right amygdala; a positive change is observed when choosing according to their behavioral tendencies and a negative change when they do not choose according to it. When subjects choose contrary to their behavioral tendencies (risky outcome in gains and sure outcomes in losses) the fMRI scan in Panel C shows activation mainly in the ACC and to a lesser extent in the dlPFC. In Panel D, one can observe relatively increased activation in the ACC when subjects make choices contrary to their behavioral tendencies.

Subjects who resisted framing effects (gains and losses are the two frames) showed relatively greater activation in the orbitomedial prefrontal cortex (omPFC). The authors interpret

activation in the amygdala when subjects are susceptible to framing effects as driven by an affect heuristic (p. 686). Furthermore, they argue (p. 687) that “As such, our findings support a model in which the omPFC evaluates and integrates emotional and cognitive information, thus underpinning more rational (i.e., description-invariant) behavior.”

These results were replicated in the study by Roiser et al. (2009) who showed that subjects who had a specific genetic neurotransmitter variant were more influenced by framing effects. But when this genetic variant was absent, subjects appeared to show no framing effects. However, Talmi et al. (2010) found that amygdala lesioned subjects continued to exhibit framing effects of the sort studied in De Martino et al. (2006). This makes it difficult to implicate the amygdala conclusively in framing effects. In related neuroeconomic studies of the disparity between WTP and WTA (for which reference dependence is a necessary condition) a correlation has been found between (i) WTP and activation in the vmPFC, and (ii) WTA and activation in the lateral orbitofrontal region (Knutson et al., 2008; De Martino et al., 2009).¹¹

23.4.3 *Loss aversion*

Tom et al. (2007) examine the neural correlates of loss aversion. They presented subjects with a range of choices between accepting or rejecting lotteries of the form $(x, p; -y, 1 - p)$, where $x > 0$ and $y > 0$. In order to ensure that subjects would never be out of pocket, they were given an endowment of money larger than the maximum loss that they could incur. Subjects were asked to rate the strength of their acceptance or rejection on a four point scale.

The median loss aversion was found to be 1.93. Other things fixed, as the size of gains increased, there was increased activation in dorsal and ventral striatum, vmPFC, vlPFC, ACC, OFC, and dopaminergic midbrain regions. There are two competing hypotheses about brain activation when subjects are presented with increased potential losses. One hypothesis is that loss aversion may be driven by a fear response that is caused by activation of brain regions such as the amygdala and anterior insula. An alternative hypothesis is that gains and losses are mediated by a single system. In this study, support is found for the second hypothesis. No brain region is found to be additionally activated when the size of potential loss in gambles increased. The same brain areas involved in the computation of expected rewards for gains showed decreased activity when potential losses are increased; these areas are the striatum, vmPFC, ventral ACC, and medial OFC. In particular, there was hardly any activity in the amygdala and the insula. Thus, it appears that there is a single brain system involved in the computation of gains and losses. Increased activity in the striatum that is associated with gains and decreased activity in the striatum associated with losses is consistent with the findings from some other studies (Breiter et al., 2001; Delgado et al., 2003). Replications of these findings also showed up decreased activity in the vmPFC when losses are increased (Cunningham et al., 2009; Plassmann et al., 2010).

So how do these results compare with other studies that find activation in the amygdala and the insula in the evaluation of lotteries with losses (Kahn et al., 2002; Kuhnen and Knutson, 2005)? Tom et al. (2007) conjecture that their experiments focused on decision utility, while the other studies might have been picking up the effects of prediction errors. De Martino et al. (2006) find a correlation in activity in the amygdala and an increase in risk. However, risk and loss aversion are different concepts. Furthermore, Tom et al. (2007) identify a form of *neural loss aversion*. The positive slope which reflects the increase in activity in the striatum and the vmPFC for gains is found to be smaller than the absolute value of the negative slope that

¹¹ For more details, the interested reader can pursue Louie and De Martino (2014).

measures a decrease in activity for losses in the same brain areas. Thus, these results provide support for a core component of prospect theory—loss aversion—a feature that humans share with close primate relatives such as capuchin monkeys (see Part 1 for the details on the empirical evidence).

Yacubian et al. (2006) consider the neural correlates of positive and negative expected values of lotteries using fMRI data. Let $x < 0 < y$, and consider choices over lotteries of the form $(x, p; y, 1 - p)$. Then the expected value EV can be written as $EV = EV^+ + EV^-$, where $EV^+ = y(1 - p)$ and $EV^- = xp$. The prediction error is given by $PE = R - EV$, where R is the actual outcome, x or y . EV^+ was found to be correlated with activity in the ventral striatum; in particular, the bilateral ventral striatum was activated as the reward magnitude or its probability increased. EV^- was correlated with activity in the bilateral amygdala.

An alternative method of discovering brain areas associated with loss aversion is to study patients with brain lesions. De Martino et al. (2010) study the behavior of two subjects with lesions in the amygdala. Subject SM, a 43 year old woman, had lesions in the entire amygdala region, while subject AM's lesions were confined to 50% of the amygdala. Both subjects were endowed with \$50 cash and then asked to accept or reject a series of mixed gambles in which gains ranged from \$20 to \$50 and losses ranged from $-\$20$ to $-\$50$. The decisions of each subject were compared to a set of six subjects, comparable in terms of education, age, income, and gender; these subjects played the role of control subjects.

Both lesioned subjects exhibit no loss aversion relative to their respective control groups; this finding echoes the spirit of studies with other primates in which amygdala lesioned monkeys approach stimuli that are rejected by non-lesioned monkeys (Mason et al., 2006). In contrast, the loss aversion parameter found in the control group in each case is close to the previous findings. However, both lesioned patients satisfy a basic consistency requirement by preferring larger gains and smaller losses; these results are consistent with evidence from other primate studies (Murray et al., 1996). The more lesioned patient, SM, exhibited a slight loss seeking behavior while the less lesioned patient, AM, was loss-neutral.

Clearly, these results on the role of the amygdala differ from those in the fMRI study by Tom et al. (2007). One possible reason is that the De Martino et al. (2010) study uses a larger range of losses. For smaller losses, the BOLD signal in the amygdala might not be detectable using existing fMRI techniques. A second possibility is that in the evaluation of mixed lotteries, the amygdala might contribute vital inputs to other brain regions (e.g., the striatum) involved with the processing of synaptic processing that is picked up by the BOLD signal. This possibility is confirmed in Jenison et al. (2011) who recorded 16 amygdala neurons whose spike is particularly associated with a WTP task for three subjects. Of these 16 neurons, 9 showed a negative correlation with the WTP and 7 showed a positive correlation. Thus, it is quite likely that in fMRI studies, the neuronal signals from the amygdala cancel out, yet the amygdala may play an important role in providing inputs during mixed lottery decision tasks.

De Martino et al. (2010) were able to distinguish between risk aversion and loss aversion. Loss aversion implies an aversion to mixed gambles because losses are weighted more heavily than gains. In contrast, in gambles with all gains or all losses, loss aversion does not have bite. For such gambles, the spread of outcomes can be varied to test for risk aversion. It turns out that the risk aversion of amygdala lesioned patients is no different from the respective control groups. Thus, amygdala damage seems to lower loss aversion while keeping fixed the level of risk aversion; indeed separate neural systems could be responsible for loss aversion and risk aversion.

The behavioral evidence shows that loss aversion can be reduced in some cases by asking subjects to think like traders and combine gains and losses (Sokol-Hessner, Delgado, et al., 2009).

The neural correlates of such manipulation have been shown to be a reduction in activity in the amygdala and increased activity in the dlPFC (Sokol-Hessner, Hsu, et al., 2009).

23.4.4 Non-linear probability weighting

Under expected utility theory, probabilities enter into the objective function in a linear manner. In contrast, in the main behavioral alternatives such as rank dependent utility and prospect theory, individuals weight objective or subjective probabilities in a non-linear manner. In particular, inverse S-shaped probability weighting is the most common empirical finding, at least for probabilities bounded away from the endpoints of the probability interval $[0, 1]$; see Part I for the details on these assertions. There are many alternative probability weighting functions, but the most satisfactory is the Prelec (1998) function; it is axiomatically founded, and consistent with the empirical evidence.¹² The one parameter form of the Prelec function is $w(p) = e^{-(\ln p)^\alpha}$, where $\alpha > 0$ is the single parameter, $p \in [0, 1]$ is the probability, and w is the probability weight.

Most probability weighting functions are relatively flat in the middle range of probabilities. Hence, in order to detect non-linear probability weighting, one needs to use lotteries with probabilities that are dispersed over the range $[0, 1]$. Neural responses to probabilities have been found in the striatum and the cingulate gyrus (Paulus and Frank, 2006; Berns et al., 2008). A PET study by Takahashi et al. (2010) found a negative correlation between dopamine D_1 receptor binding in the striatum and the degree of non-linear probability weighting (relative to the benchmark of linear weighting).

Hsu et al. (2009) allow for a suitably large range of probabilities; we consider their work, next. In their experiments, on the first screen, subjects observed a lottery $L_1 = (x_1, p_1)$ (x_1 is an outcome and p_1 is a probability) without the confounding influence of a comparison between two lotteries. On the second screen, subjects were asked a question to test their understanding of the lottery (x_1, p_1) (e.g., is the probability lower than or greater than p ?). On the third, and final, screen, subjects were asked to choose between two lotteries, $L_1 = (x_1, p_1)$ and $L_2 = (x_2, p_2)$, where (x_2, p_2) is another outcome–probability pair. The first experiment (first two screens) revealed the brain areas involved in *reward anticipation*, while the second experiment (third screen) revealed brain areas involved in *choice*. Probabilities were chosen from the set $\{0.01, 0.1, 0.3, 0.5, 0.8, 0.95\}$ and outcomes, all positive, were chosen from the set $\{10, 20, 50, 100\}$.

In order to separate the objective function into a linear and a non-linear component, the objective function under prospect theory for lottery $L = (x, p)$ can be written as $PT(L) = \pi(p, \alpha)v(x)$, where $\pi(p, \alpha)$ is the decision weight under prospect theory (α is the parameter of the one-parameter Prelec function) and $v(x)$ is the utility function under prospect theory (see Part I for the details). We can rewrite the objective function as

$$PT(L) = \pi(p, \alpha)v(x) = pv(x) + \Delta(p, \alpha)v(x), \text{ where } \Delta(p, \alpha) = \pi(p, \alpha) - p. \quad (23.1)$$

In the parametric estimation, v is assumed to take the power form $v(x) = x^\gamma$, $\gamma > 0$.¹³

¹² For three different axiomatic foundations, see Prelec (1998), Luce (2001), and al-Nowaihi and Dhami (2006). For empirical support for the Prelec function, see Stott (2006).

¹³ The power form of utility is supported by the evidence (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992; Stott, 2006). For an axiomatization of the power form of utility under prospect theory, see al-Nowaihi et al. (2008).

In the first experiment that involved reward anticipation, the BOLD signal in the brain is regressed on the two terms on the RHS of (23.1), $p\nu(x)$ and $\Delta(p, \alpha)\nu(x)$. If humans only take account of linear probability weighting, as under expected utility, then the BOLD signal should not be correlated with the non-linear term $\Delta(p, \alpha)\nu(x)$; in constructing $\Delta(p, \alpha)$ the group mean value of $\alpha = 0.77$ was used.

A non-parametric method was also employed to correlate brain activity with the behavioral shape of the Prelec function.¹⁴ Each of the probabilities is assigned a dummy variable, $d(p)$ (e.g., $d(0.5) = 1$ and $d(p') = 0$ for any $p' \neq 0.5$). Letting z be the BOLD response in the brain, captured by the fMRI evidence, the relation

$$z = a + \beta d(p)\nu(x) + \varepsilon \quad (23.2)$$

was estimated; ε is a noise term and a, b are the parameters to be estimated.

The results are as follows. In the parametric model, the brain regions that correlated with the linear term ($p\nu(x)$) in the first experiment (anticipated reward case) are the striatum, cerebellum, motor areas, and the median prefrontal cortex. These results are consistent with the identification of similar areas in related tasks. Activation in the following regions was correlated with the non-linear term ($\Delta(p, \alpha)\nu(x)$): striatum, cingulate gyrus, motor cortex, and the cerebellum. The areas of the brain that were equally activated by the linear and non-linear terms are striatum, motor cortex, and the cerebellum. The activation of the striatum for the reward anticipation task is consistent with the results from other studies on reward anticipation/reward processing (Knutson et al., 2000; O'Doherty et al., 2004). This region has also been found responsive to a multiplicative combination of probabilities and reward magnitudes (Tobler et al., 2007).

The coefficient β in the non-parametric model (23.2) is shown for various probabilities as the dots in Figure 23.11 for the left and right striatum, separately. The BOLD signal activation




Figure 23.11 The coefficient β in (23.2) (shown as the dots) for the left striatum (left panel) and the right striatum (right panel) for different probabilities. The behavioral Prelec function for the group average value of $\alpha = 0.77$ is shown as the continuous curve.

Source: Hsu et al. (2009).

¹⁴ Note that the inverse S-shape of the Prelec function arises for $\alpha < 1$ but not for $\alpha > 1$, which is also admissible under the axioms that lead to the Prelec function; see al-Nowaihi and Dhami (2006).

is non-linear in probabilities; it is overweighted relative to the behavioral Prelec function for low probabilities and underweighted for high probabilities.

23.5 The neuroeconomics of social preferences

The long-overdue incorporation of social preferences into economics has led to a huge increase in our understanding of a wide range of phenomena. Part 2 of this book, one of the longest, is devoted to a study of social preferences. In this section, we study the neural basis of some aspects of human sociality.

23.5.1 *Computation of social rewards in the brain*

Most of the studies mentioned in Section 23.4.1 on the neural correlates of private rewards, use non-social rewards such as juice, food, or money. In this section, we ask the question: Which are the neural correlates of social rewards? Are the same regions and neural processes that are involved in private rewards, also implicated in the processing of social rewards? Izuma et al. (2008) found that acquiring a good reputation (a social reward) activates the same brain areas as in the case of monetary rewards. They conclude that there is a common neural currency for private and social rewards, which is supported by other neuroeconomic evidence (Zink et al., 2008; Smith et al., 2010).

We focus now on Lin et al. (2012) who consider identical tasks that can lead to either private or social rewards. The task is shown in Figure 23.12. In the private/monetary and social rewards conditions, individuals chose between two slot machines, each of which produced three possible outcomes, positive (\$1), neutral (\$0), and negative (−\$1); subjects had to learn the probability distribution over outcomes produced by each of the slot machines. In each stage, the choice was between a neutral slot machine (represented by six circular dots) which had a uniform distribution over the three outcomes, and one of two other slot machines (see top two machines in Panel B in Figure 23.12).

If the outcome is positive, then (i) under social rewards, individuals see a smiley face and hear emotionally matched words such as excellent, bravo, and fantastic, and (ii) under monetary rewards subjects simply received \$1. If a negative outcome arises, then (i) under social rewards, subjects see an angry face and hear words such as stupid, moron, and wrong, and (ii) under monetary rewards, \$1 was subtracted from their endowments. Finally, in the case of a neutral outcome, (i) under social rewards, subjects see a neutral face and hear neutral words such as desk and paper, and (ii) under monetary rewards they receive \$0.

As one would expect in such a simple task, participants learnt to choose the positively valenced slot machine after a few trials (the machine in the middle row in Panel B); see Panels C and D in Figure 23.12. The fMRI data for subjective value (SV), reward magnitude (R), and prediction error (PE) for the social, monetary, and a combined condition are shown for various regions of the brain in Figure 23.13. In the monetary and social conditions, the SV on the slot machines was correlated with the vmPFC and additionally with the mid-cingulum, superior frontal gyrus, and angular gyrus. Thus, the neural circuitry for computing SV is similar in the monetary and social rewards tasks. In the monetary condition, activation in the caudate and putamen was correlated with PE, while there were no statistically significant correlations, at the chosen threshold, in the social task. Finally, reward magnitudes (R) are correlated with activation in the vmPFC for social and monetary conditions; in the monetary condition there was also activation in the insula, occipital




Figure 23.12 Panel A: Each choice involved choosing from a neutral or valenced slot machine. The decision and rewards time intervals are shown. Panel B: Probability distributions over rewards of the three slot machines, each represented by a different pattern. Panels C and D: Choices of subjects by disaggregated and aggregated treatment effects.

Source: Lin et al. (2012).

cortex, cingulate gyrus, and the superior frontal gyrus. The authors conclude (p. 280): “Together with other recent findings . . . our results provide increasing support that overlapping areas of vmPFC and ventral striatum encode value signals for both types of rewards.”

23.5.2 The neural basis of inequity averse preferences

We have discussed *models of inequity aversion* in Part 2 (Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000). However, several explanations that were traditionally based on inequity aversion were sought to be supplanted later by concerns for social image (Andreoni and Bernheim, 2009) as well as reciprocity considerations (Falk et al., 2008; Rabin, 1993). In this section, we report neural evidence, based on fMRI tests, that establishes neural networks in the brain, which encode inequity averse preferences in a simple model.




Figure 23.13 Activation of various brain regions correlated with SV (first row), PE (second row), and R (third row).

Source: Lin et al. (2012).




Figure 23.14 The behavioral evidence. Vertical bars denote standard errors of the mean.

Source: Reprinted by permission from Macmillan Publishers Ltd. *Nature* 463: 1089–91 (25 February 2010). "Neural evidence for inequality-averse social preferences," Elizabeth Tricomi, Antonio Ranger, Colin F. Camerer and John P. O'Doherty.

In Tricomi et al. (2010), the data is based on 20 pairs of unacquainted males who were given a baseline pay of \$30. Half were then randomly selected and given \$50 (high-pay group), while the other half was given nothing extra (low-pay group). Individuals in both groups then had to rate an incremental transfer between \$0 and \$50 to themselves and to their partner in the group by assigning a number from a ratings set $\{-5, -4, \dots, 0, \dots, 5\}$.

Figure 23.14 shows the regression coefficient obtained by regressing ratings on the amount transferred to themselves and to others. Both groups rated own-transfers positively. Now consider transfers to the partners. Members of the high-pay group rated transfers to the low-pay

partner positively (advantageous inequity), because such transfers narrowed the income gap. By contrast, members of the low-pay group rated transfers to the high-pay partner negatively (disadvantageous inequity), because such transfers widened the income gap. Thus, the behavioral data is consistent with the Fehr–Schmidt (1999) model of inequity aversion.

fMRI data shows that in the case of own-transfers, there was a statistically significant correlation in the activation of the vmPFC and ventral striatum with the magnitude of transfers. This is not surprising. The same areas are known to be activated in the case of private rewards; see Section 23.4.1.

The authors next considered a statistic, which is the contrast of the fMRI response for own-transfers relative to partner-transfers for both groups. Figure 23.15 shows the resulting effects on the ventral striatum while Figure 23.16 shows the corresponding effects on the vmPFC. For each of the two groups, there is activation in both the regions and the contrast statistic is significantly higher for the low-pay group. For the low-pay group, the activation of both regions is greater for own-transfers as compared to partner-transfers. For the high-pay group, we get the opposite




Figure 23.15 The fMRI evidence on activation of the ventral striatum for the contrast between own-transfers relative to partner-transfers for both groups. (a) Activation in the ventral striatum (circled); contrast was significantly higher for low-pay group. (b) Contrast values for both groups. Vertical bars denote standard errors of the mean.

Source: Reprinted by permission from Macmillan Publishers Ltd. *Nature* 463: 1089–91 (25 February 2010). “Neural evidence for inequality-averse social preferences,” Elizabeth Tricomi, Antonio Ranger, Colin F. Camerer, and John P. O’Doherty.




Figure 23.16 The fMRI evidence on activation of the vmPFC for the contrast between own-transfers relative to partner-transfers. (a) Activation in the vmPFC (circled); contrast was significantly higher for low-pay group. (b) Contrast values for both groups. Vertical bars denote standard errors of the mean.

Source: Reprinted by permission from Macmillan Publishers Ltd. *Nature* 463: 1089–91 (25 February 2010). “Neural evidence for inequality-averse social preferences,” Elizabeth Tricomi, Antonio Ranger, Colin F. Camerer, and John P. O’Doherty.

result, i.e., relatively more activation for partner-transfers relative to own-transfers. The results hold even when wealth effects and prediction errors are controlled for.

Comparing Figure 23.14 with Figures 23.15 and 23.16 for the high-pay group, one notices a discrepancy between the behavioral and the neural results. In the behavioral results, one rates an own-transfer relatively higher than a partner-transfer while the reverse is true for the neural data. The authors note (p. 1090): “basic reward structures in the brain may reflect even stronger equity considerations than is necessarily expressed or acted on at the behavioral level. These findings raise the possibility that even when basic reward responses reflect strong equity considerations, in some cases additional factors may intercede to moderate the influence of such equity judgements on behavior, such as strategizing . . . or the engagement of self-serving biases such as judgements of deservingness or need.” This is an example where neural data enriches understanding of human behavior relative to a consideration of behavioral data alone.

Cappelen et al. (2014) study the neural correlates of inequity aversion when individuals evaluate different income distributions that are generated by an input of endogenous effort. The behavioral data suggests a preference for income distributions in which incomes are proportional to effort. Thus, individuals are less inequity-averse when inequity arises from the input of endogenous effort. There was a significant activation in the striatum whenever the income distribution deviated from the one where incomes are proportional to effort. In competitive environments, such as tournaments between competing agents, workers, or athletes, winning is considered more personally rewarding. In such environments, winning is correlated with increased activity in the vmPFC and the striatum, which are also activated in the case of private rewards (Bault et al., 2011).

23.5.3 Neural basis of social punishment and trust

Classical game theory predicts that bygones are bygones, so we should not observe players imposing costly punishments on others. Yet, the empirical findings of costly punishment in a range of games, precisely when bygones should be bygones, is puzzling. A prominent example is the data from public goods games with punishment (Fehr and Gächter, 2002; see Part 2 for details). In other contexts, one observes costly punishment meted out to others, including third parties, for violating social norms; this is sometimes known as *altruistic punishment* or *indirect reciprocity*. Indeed, such reciprocal behavior appears to have been critical to the survival of traditional societies, prior to the establishment of a formal state, and it has evolutionary foundations.¹⁵ In this section, we consider the neurobiological foundations of such behavior.

De Quervain et al. (2004) consider a *trust game* between player A (the trustor) and player B (the trustee) that allows for punishment of the trustee for betrayal of trust. Players A and B are endowed with 10 MUs (experimental unit of currency). Player A sends an amount, x MUs, to B that the experimenter quadruples and passes on to B. Player B then has an opportunity to repay this trust by choosing to send back one of two amounts: 0 MUs or $\frac{1}{2}(10 + 4x)$ MUs. On observing the amount sent back by B, player A is given 1 minute to deliberate and then decide on the level of punishment to impose on B. Different treatments vary the cost of punishment on both players. During the 1 minute that player A deliberates, his brain is scanned using positron emission tomography (PET). In order to reduce the exposure of radioactive substances, the authors scan in only those cases where player B sends back nothing; we report these results below.

¹⁵ See Part 2 for the details. Several excellent references are available. These include Boyd et al. (2003), Bowles and Gintis (2004, 2011), Fehr and Gächter (2002), Fehr et al. (2002), and Fehr and Fischbacher (2003).

Four different treatment conditions were considered, depending on the cost of punishment; these are IC, IF, IS, and NC. The first letter in any name of treatment captures the intentionality (I) or non-intentionality (N) of the decision made by B. Under N, a random device makes the decision on behalf of player B, while under I, player B directly makes the decision. The second letter of the treatment name determines if the punishment is costly (C), free (F), or purely symbolic (S). Under C, player A gives up 1 MU to reduce the payoff of player B by 2 MUs, while under F, there is no cost to player A. Under S, there is no cost to either of the players.

The four treatments provide the necessary contrasts to evaluate the relevant hypotheses. Player A should have an incentive to punish player B whenever B intentionally chooses the action of sending back nothing (the only case considered below). When punishment is symbolic, player A has the intention to punish but not the means to do so, hence, he might not get any satisfaction in this case. The contrast IF–IS captures the satisfying aspect of punishment. If punishment is satisfactory in the treatment IF, then provided punishment is not too costly, we expect players to also punish in the IC treatment, although the punishment might be reduced. Indeed, brain areas that are involved in the calculation of such trade-offs may come into play.

One might expect that subjects with the greatest activation in the reward areas in the brain in the IF treatment might engage in the greatest punishment in the IC treatment. If so, then this would contribute to the ability of neuroeconomics to make successful predictions. When the decision by player B is non-intentional, then player A should not derive any pleasure from punishing player B. The combined contrast (IC + IF) – (IS + NC) is also useful. In the treatments IC and IF, there is a desire and opportunity to punish while in treatment IS there is no opportunity to punish and in treatment NC there is no desire to punish.

Player A is asked to rank the perceived unfairness of player B and the desire to punish player B on a seven-point *Likert scale* from –3 to +3. A Likert scale is a bipolar scale in which subjects indicate a positive and a negative response, with 0 denoting indifference. The behavioral results are shown in Figure 23.17 for each of the four treatments. Maximal perceived fairness is indicated by –3 and maximum perceived unfairness by +3. The maximal desire to punish is indicated by +3 and the maximal desire to reward by –3.

Mean perceptions $\pm$ standard errors are shown in Panels A, B in Figure 23.17. When the choices of B are unintentional (treatment NC) there is no perceived unfairness (panel A), desire to punish




Figure 23.17 Mean responses by player A when player B returns 0 MUs, and the actual payoff reduction of player B in various treatments.

Source: "The neural basis of altruistic punishment." Dominique J.-F. de Quervain, Urs Fischbacher, Valerie Treyer, Melanie Schellhammer, Ulrich Schnyder, Alfred Buck, and Ernst Fehr, *Science* (27 August 2004) 305 (5688): 1254–8. Reprinted with permission from AAAS.

is very small (panel B), and the payoff of player B is reduced very little through punishment, possibly on account of inequity-aversion (panel C). In the other three treatments, the intentional choices of B are perceived as unfair, there is desire to punish, and actual punishment takes place. The highest punishment takes place when it is free (treatment IF).

Consider the contrasts $(IC + IF) - (IS + NC)$, $IF - IS$, $IC - IS$, $IF - NC$, $IC - NC$. In these contrasts, the PET results show that the highest activation is achieved in the caudate nucleus and this activation is particularly significant in the contrasts $(IC + IF) - (IS + NC)$ and $IF - IS$. The MNI coordinates of this brain area are [6, 22, 4]. A secondary area that was activated in these contrasts is the thalamus.

Figure 23.18 shows the peak blood-flow increase in the caudate nucleus in each treatment, relative to the average level of activation. The greatest activation was in the IC and IF treatments where there is a desire and an opportunity to punish, relative to the other two cases. The caudate nucleus has been found to be associated with the brain's reward circuits in brain lesion studies in rats, single neuron studies in non-human primates (Hollerman et al., 1998) and neuroimaging studies in humans (Knutson et al., 2001; Delgado et al., 2000). The thalamus too has been implicated in the processing of rewards in neuroimaging studies in humans (Delgado et al., 2000). Thus, it appears that the act of punishing an unfair act by player B is hedonically rewarding to player A.

An important step is taken in establishing causality by first identifying the 16 subjects in their roles as player A who exhibit the highest caudate activation in the IF condition when they punish player B. These players are also the ones who punish maximally in the IC condition when punishment is costly. Thus, the PET study appears to clearly identify those players who derive the greatest pleasure from punishing norm violation by player B; these players behave exactly as predicted in the IC condition. Costly punishments involve a trade-off. On the one hand, punishment is satisfying, but on the other hand, punishment imposes a personal cost. In this case, the PET data shows activation also in the mOFC, which is typically engaged in making difficult choices.

Empathy allows individuals to share the emotions, pain, and suffering of others. Neural evidence shows that similar brain regions are highlighted when one experiences painful events or suffering, or observes others experiencing these events with whom one has empathy. These brain areas are typically the anterior insula (AI), front-insular cortex (FI), and the anterior cingulate cortex (ACC) (Wicker et al., 2003; Keysers et al., 2004; Singer et al., 2004; Morrison et al., 2004).




Figure 23.18 Peak activation in the caudate nucleus relative to the mean level in all four treatments.

Source: de Quervain et al. (2004).

One might wonder how empathy towards others varies, depending on how favorably one views them. This question is addressed by Singer et al. (2006) and is the one that we consider next.

They play a modified trust game that they call a sequential prisoner's dilemma. Subjects in their roles as first movers could send an amount of money to a second player who then chooses to either send back a low, seemingly unfair, amount or a high, seemingly fair, amount. The second player is a confederate of the experimenter in the sense that he/she chooses a fair or an unfair amount in accordance with the experimenter's instructions. First movers were then asked to rate the fairness, agreeableness, or likeability of the second player; they rated the second player positively on these traits if they sent a fair amount, and negatively otherwise.

Once first movers had rated the second players, fMRI was used to monitor their empathetic responses when a painful stimulus was applied to them (self-condition) and to the other players who were fair or unfair (others-condition). Consistent with the earlier literature on empathy, in the self-condition, the activated brain areas were AI and ACC, extending into the FI and brain-stem. There was also a high correlation between empathy scores in behavioral tasks and activation of the relevant brain areas associated with empathy. The others-condition gives rise to two main results. First, less brain empathy is generated, measured by the activation of the relevant brain areas in fMRI data, when the partner is unfair, relative to being fair. Second, there is a significant gender difference. Measured by peak brain activity, there was little or no empathetic response in the fMRI data of men when the partner was unfair. In contrast, women still showed a significant empathetic brain response to unfair partners, although lower as compared to fair partners.

In an early study, Sanfey et al. (2003) showed that players who were treated unfairly in ultimatum games (offered a low amount by the proposer) exhibited activation in the emotional regions of the brain, such as the anterior insula. Thus, rejection of unfair offers may be associated with a negative emotional response. Tabibnia et al. (2008) argue that it is not clear if the brain responds to a lower monetary offer or a lower perceived fairness of the offer. They disentangle these two effects in which all players, in their roles as responders, play two ultimatum games. There were no proposers in the game, although the responders were told that proposers had already made offers that they were to respond to and actual payments would be made to the hypothetical proposers, whose pictures were shown to them.¹⁶

Consider, for instance, the offers in ultimatum games shown in Figure 23.19. An offer of \$7 out of \$15 may be deemed as fair but so is an offer of \$0.50 out of \$1. Hence, one may vary the monetary amount but maintain the same degree of fairness. Similarly, the last row of Figure 23.19 shows that depending on the base amounts, two different monetary payoffs, one higher than the other, may be deemed to be unfair.

In the behavioral results, self-reported happiness of proposers is significantly higher for fairer offers and self-reported contempt was significantly higher for unfair offers. The main areas of brain activity revealed by fMRI data that were activated in the fair versus unfair offers contrast were vmPFC, ventral striatum, and the amygdala. The activation of the amygdala is interesting because one typically associates it with a fear response. However, it has also been implicated in reward processing in other studies also.¹⁷ One potential limitation of this study is that the authors only included matched pairs of data in which the same amount was rejected in one case and accepted in the other (see p. 342). Hence, as Fehr and Krajbich (2014, p. 206) note, the "high > low fairness contrast is essentially the same as contrasting accepted > rejected offers."

¹⁶ Experimenters in economics often argue strongly for eliminating deception from experiments (Houser and McCabe, 2014, p. 27). The experiments reported here could possibly be construed as borderline deception.

¹⁷ The interested reader can look up the references in Tabibnia et al. (2008).

		Material payoffs	
		High	Low
Fairness	High	\$7 out of \$15	\$0.50 out of \$1
	Low	\$7 out of \$23	\$0.50 out of \$10

Figure 23.19 Different offers in an ultimatum game that signify varying levels of fairness.

Source: Tabibnia et al. (2008).

Baumgartner et al. (2012) explore the neural basis of *parochial altruism*, which is a form of altruism that is particularly directed at ingroup members only; see Chapter 7 for issues of social identity. The authors give a very clear definition (p. 1453): “Parochial altruism constitutes a persuasive psychological phenomenon which is qualified by a preference for altruistic behavior towards the members of one’s ethnic, racial, or any other social group, combined with a tendency for indifference, mistrust, or even hostility toward outgroup members.” Parochial altruism has played a critical role in norm compliance not only over human evolution but also in modern times; see Parts 2 and 5 of this book for details. Such norm compliance often takes the form of third party punishment of norm violators. Third parties are typically not harmed by observed opportunistic behavior between other parties, yet they are willing to engage in costly punishments that do not give them any direct benefits. The authors conjecture that such behavior, which is the product of human evolution, might also have given rise to specialized neural structures that support it.

The setting for the social identity experiments used by the authors has already been outlined in Chapter 7 (Goette et al., 2006). Swiss army recruits are randomly assigned to platoons over a four week training course. The behavioral results are consistent with the social identity view that own-platoon members are considered as ingroups relative to other-platoon members. A simultaneous move prisoner’s dilemma is conducted with two players A and B in the third or fourth week of training. A and B are given 20 points; the points can be exchanged for money with the experimenter at a known exchange rate. Under cooperation, C, the players can simultaneously transfer 20 points to the other player, which the experimenter doubles before transferring to the other player. So if both players transferred their entire endowments to the other player, each gets 40 points. The players can also choose not to transfer (defection, D) and simply exchange their endowment for money with the experimenter at the end of the game. Since decisions are made simultaneously, the four outcomes or *conditions* are CC, CD, DC, and DD. Thus, for instance, in condition DC, player A plays the strategy defect (D) and player B plays the strategy cooperate (C).

A third player, player C (not to be confused with the strategy C) observes the results of these experiments within five days and must decide on the level of punishment. He can only punish one of the two players. The identities of the players are so arranged that player C can only punish player A, who could be a ingroup or an outgroup member. Punishment costs 1 point but allows for a deduction of 3 points from the target. While engaging in punishment, fMRI data was collected from player C, whose behavior will be our main focus of attention.

There are three main groups/constellations, each with three players. In group ABC, all three players are from the same platoon; in group AC, players A and C are from the same platoon, but B belongs to another platoon; in group BC, players B and C are from the same platoon, but A belongs to another platoon. To summarize:

{ 2 player Conditions/Strategy profiles chosen by players A, B:	{CC, CD, DC, DD}
{ 3 player Groups/Constellations listing ingroup members only:	{ABC, AC, BC}

Consider the condition DC in which player A defects and player B cooperates. In the constellation BC, player A is an outgroup member, while in the constellation AC, A is an ingroup member. Parochial altruism arises if player C punishes A more in the constellation BC (where A is an outgroup member) as compared to the constellation AC (where A is an ingroup member).

In the case of two-player punishment games, the brain regions involved in the punishment of norm violators are the right lateral prefrontal cortex (rLPFC) (Knoch et al. 2006, 2008; Sanfey et al. 2003), and the dorsal caudatus (de Quervain et al. 2004; Seymour et al. 2007). The calculations in third party punishment games are similar to the case of two-player punishment games. One has to trade off the loss in personal material payoff with the desire to punish norm violators; indeed such punishment offers personal satisfaction (de Quervain et al., 2004). Hence, we might expect similar brain regions to be activated in third party punishment.

However, the introduction of third party punishment in a social identity experiment introduces other issues. Third parties will need to engage *mentalizing processes* that justify why punishment meted out to ingroup members should be lower as compared to outgroup members for the same norm violation? Such mentalizing might take the form of mitigating factors to justify norm violations of ingroup members. Neural evidence indicates that the brain regions associated with this type of mentalizing are the dorsomedial prefrontal cortex (dmPFC) and the bilateral temporoparietal junction (TPJ). A range of studies have implicated the TPJ in suggesting *a theory of the mind*, i.e., the ability of an individual to take on the perspective of another person (Decety and Lamm, 2007; Frith and Frith, 2007; Van Overwalle, 2009).

The main condition of interest that we now focus on is DC (because in the conditions CC, CD, and DD, player A behaves no worse than player B so there would appear to be no desire to punish). Two main brain contrasts were used in the condition DC. (1) Constellation BC (player A belongs to the outgroup) minus the constellations (AC + ABC) (player A belongs to the ingroup). (2) The constellation (AC + ABC) minus the constellation BC. The first contrast allows us to examine if there is increased brain activity when one punishes an outgroup norm violator. The second contrast allows one to test if there is increased activity in the mentalizing network when one reduces punishment of ingroup norm violators.

A special software (MarsBaR) is used to identify the *regions of interest* (ROI analysis) during fMRI imaging.¹⁸ This is based on maximum observed voxel activity in these regions during third party punishment. Furthermore, Physio-Physiological Interaction (PPI) analysis was used to determine interaction between brain regions that goes towards establishing causality of one brain region in modulating activity in other brain regions.

The behavioral results confirm the classical results in social identity theory. There is punishment of norm violators, which arises most starkly in the condition DC (relative to the conditions CC, CD, and DD). Furthermore, there is a statistically strong effect of group membership on

¹⁸ Other programs can also be used to define ROI on the basis of anatomical or functional criteria, e.g., SPM and MRICro.

punishment; outgroup members belonging to other platoons are punished much more strongly than ingroup members for the same norm violation.

The fMRI results are as follows. Consider first the outgroup effects as captured by the contrast between the constellation BC and the constellations (AC + ABC); see Figure 23.20 for the results. Here the third party punishes outgroup members who violate norms against ingroup members who follow norms. The expected brain areas involved in punishment decisions showed increased activity in this case. These included the right orbitofrontal gyrus (rOFG), rLPFC and the right dorsal caudatus (rdCau). In order to establish the connectivity across these brain regions, a PPI analysis reveals that greater activity in rOFC is associated with greater connectivity of neural activity between rLPFC and the right caudatus. This suggests that the three regions are part of a functionally connected neural network that determines punishment behavior.

The results for the second brain contrast, the constellation (AC + ABC) minus the constellation BC, that captures ingroup effects is shown in Figure 23.21. The main brain activity in the contrast took place in those regions of the brain that have already been implicated in the mentalizing network; these regions are dmPFC and the temporoparietal junction (TPJ) (Gallagher and Frith, 2003; Rilling et al., 2004a). In Figure 23.21, the ROI analysis revealed the left and right parts of the TPJ, respectively lTPJ and rTPJ, as being active in this contrast. Interestingly, the stronger the interaction between these two regions of the mentalizing network, the lower is the punishment of ingroup members for defecting. A PPI analysis showed that activity in the lTPJ modulates the connectivity between dmPFC and areas in the OFC and rLPFC. These areas, thus, appear to form a mentalizing network that shades punishment of ingroup members.

Neural findings suggest that disrupting the right dlPFC through repeated transcranial magnetic stimulation (rTMS) reduces the rejection rates of subjects (often called deviant subjects) in




Figure 23.20 Outgroup effects captured by third party punishment of outgroup members who play opportunistically against ingroup members who cooperate. All data shown is for the condition DC. The first row shows activation in the three main brain regions of interest using fMRI data. The second row shows the contrast estimates, relative to a baseline, in the three brain regions, for three different constellations BC, AC, and ABC. Bars denote standard errors of mean.

Source: Baumgartner, T., Götze, L., Gügler, R., and Fehr, E. (2012). "The mentalizing network orchestrates the impact of parochial altruism on social norm enforcement." *Human Brain Mapping* 33: 1452–69. With permission from John Wiley and Sons. Copyright © 2011 Wiley-Liss, Inc.




Figure 23.21 Ingroup effects captured by third party punishment of ingroup members who play opportunistically against outgroup members who cooperate. All data shown is for the condition DC. The first row shows activation in the three main brain regions of interest using fMRI data. The second row shows the contrast estimates, relative to a baseline, in the three brain regions, for three different constellations BC, AC, and ABC. Bars denote standard errors of mean.

Source: Baumgartner, T., Götze, L., Gügler, R., and Fehr, E. (2012). "The mentalizing network orchestrates the impact of parochial altruism on social norm enforcement." *Human Brain Mapping* 33: 1452–69. With permission from John Wiley and Sons. Copyright © 2011 Wiley-Liss, Inc.

ultimatum games. Such disruption does not influence the ability of the deviant subjects to judge the fairness of offers (Knoch et al., 2006; Knoch et al., 2008; Wout et al., 2005). However, when rTMS is applied to the left dlPFC, it produces no differences in rejection rates relative to a sham treatment.

Baumgartner et al. (2011) combine fMRI methods with rTMS in order to reduce the limitations of these methods when they are applied separately. In their experiments, subjects play an ultimatum game. The proposer has an endowment of CHF 20 and he proposes a split to the responder. The offers are restricted to the set {4, 6, 8, 10}, so 4 is the most unfair offer and 10 the most fair because it results in an equal share of the pie for both players. Low frequency, 1Hz, rTMS is applied for 15 minutes to the right or left dlPFC of responders, prior to their decision to reject or accept offers (offline rTMS). This disrupts the functioning of the dlPFC; deviant subjects whose right dlPFC is disrupted are expected to reject unfair offers less frequently.

The results are shown in Figure 23.22. The rejection rates increase when the proposers make smaller offers. Disruption of the right dlPFC by rTMS leads to a much lower rate of rejection, which is consistent with the earlier results. Overall, the percentage rejection rates for offers of 4, 6, 8, and 10 were, respectively, 62, 25, 1, and 1. The percentage rejection rates for the unfair offers of 4 and 6, under right rTMS, were 45 and 13, while the corresponding figures for left rTMS were 79 and 35. The differences are statistically significant.

Several empirical studies report the following findings. (1) In the processing of unfair offers, two main brain regions, the anterior insula and the dACC, are reported to be activated (Sanfey et al., 2003; Tabibnia et al., 2008). (2) These brain regions are also known to be involved in the processing of negative emotions and in modulating conflict between cognitive and behavioral tasks (Baumgartner et al., 2009; Herwig et al., 2007).




Figure 23.22 Mean rejection rates $\pm$ s.e.m. for different treatment groups and for different offers.

Source: Reprinted by permission from Macmillan Publishers Ltd. *Nature Neuroscience* 14: 1468–74. Baumgartner, T., Knoch, D., Hotz, P., Eisenegger, C., and Fehr, E. (2011). "Dorsolateral and ventromedial prefrontal cortex orchestrate normative choice."

The empirical evidence shows that disrupting the left dlPFC does not reduce rejection rates but disrupting the right dlPFC does reduce rejection rates; see Figure 23.22. So the regions listed in (1) and (2) above, are expected to be activated in the contrast between these two disruption treatments when evaluating fair and unfair offers. However, no differential activation of these regions is found in the two treatments. The only brain region that is differentially activated in the two treatments in the fair/unfair contrast is the posterior vmPFC, a region typically associated with positive stimuli. However, unfair offers are negative stimuli. This suggests that distinct areas in the vmPFC are associated with the processing of positive and negative stimuli, a result that is confirmed by the authors.

Knoch et al. (2009) examine if disruption of the right lateral prefrontal cortex (rlPFC) by rTMS alters behavior in a 15 period repeated trust game in two treatment conditions. Trustees could either transfer back nothing, or transfer a quarter, or equalize payoffs by their back transfer. In the anonymous condition, the trustor does not learn of the trustee's past behavior. However, in the reputation condition, he can observe their behavior in the last three periods. The results are shown in Figure 23.23 for the following three cases. (i) Disruption of right dlPFC. (ii) Disruption of left dlPFC. (iii) A sham condition.

In the anonymous condition, there are no differences in the back transfers by the trustee in the three cases because there is no incentive to build a reputation. However, in the reputation condition, disruption to the right dlPFC causes lower back transfers relative to the other two cases. Such disruption appears to increase the temptation to reap current rewards relative to investing in one's reputation. It turns out that trustees who engaged in reputation building ended up with relatively higher payoffs. For instance, trustees who always chose to equalize payoffs, faced a 71% chance of getting a 10 point investment from the trustor, while the corresponding probability for




Figure 23.23 Relative back transfers in the two treatment conditions with varying degrees of brain function disruption.

Source: Knoch, D., Schneider, F., Schunk, D., Hohmann, M., and Fehr, E. (2009). "Disrupting the prefrontal cortex diminishes the human ability to build a good reputation." *Proceedings of the National Academy of Sciences* 106: 20895–9.

those who returned nothing was less than 6%. It is important to note that disruption to the right dlPFC did not alter the fairness judgments of subjects. This was inferred from questions that subjects were asked immediately after the experiment to judge the fairness of alternative levels of back transfers.

Individuals with lesions to the vmPFC have been found to be more sensitive to fairness concerns because they are more likely to reject unfair offers in ultimatum games (Koenigs and Tranel, 2007). However, when asked to respond to hypothetical unfair offers, their answers are very

similar to normal individuals without lesions (Leland and Grafman, 2005; Krajbich et al., 2009). Individuals with lesions in the vmPFC also exhibit diminished social preferences in their roles as proposers in dictator games, and trustees in trust games (Krajbich et al., 2009). The difference in hypothetical and actual choices indicate that the vmPFC is not involved in the detection of fairness, but it is involved in the expression of social preferences.

The amygdala has also been implicated in social preferences. For instance, individuals who have lesions in the amygdala react less to untrustworthy behavior (Koscik and Tranel, 2011). The amygdala also plays a role in determining punishments for different crimes (Buckholtz et al., 2008). In an ultimatum game, administration of a benzodiazepine that lowers amygdala activity also reduced the rejection rates (Gospic et al., 2011).

The size of the amygdala has been found to be associated with the size of one's social network and social behavior in nonhuman primates (Lewis and Barton, 2006). Bickart et al. (2011) conduct a similar analysis for humans. The regression results are shown in Figure 23.24, using two alternative but highly correlated measures of social networks. These are the number of people in the social network and the number of embedded networks (the number of distinct groups of contacts). Amygdala volume is measured using MRI data. Since the amygdala volume decreases with age, it is controlled for in the empirical analysis. Linear regression lines for both measures of social networks are plotted in Figure 23.24, and the regression coefficient is shown as the number B . A significantly positive relation between amygdala volume and social networks was found using both measures. Bickart et al. (2014) show that different parts of the amygdala are responsible for different aspects of social behavior.

Altruistic behavior has been observed in a wide range of other human activities. For instance, one may be slightly more altruistic when being observed by others. There is some evidence that there is increased activity in the left ventral striatum when one's charitable donations are observed by others (Izuma et al., 2010). Hsu et al. (2008) consider trade-offs between equity and efficiency. They find that a brain area, the putamen, responds to efficiency considerations while the insula




Figure 23.24 Two different measures of social networks plotted against amygdala size for two age groups: young participants (black circles) and older participants (gray triangles).

Source: Reprinted by permission from Macmillan Publishers Ltd. *Nature Neuroscience* 14: 163–4 (2011), "Amygdala volume and social network size in humans," Kevin C. Bickart, Christopher I. Wright, Rebecca J. Dautoff, Bradford C. Dickerson, and Lisa Feldman Barrett.

responds to inequity considerations. Zaki and Mitchell (2011) also provide behavioral and neural support to efficiency considerations. They implicated the OFC in efficiency choices. The activation of the insula in equity considerations, which is also found in several other studies, indicates that fairness in behavior has a basis in emotional processing. Indeed, the very first fMRI study of ultimatum games found increased activity in the anterior insula and the ACC in the fair/unfair offers contrast (Sanfey et al., 2003). In redistribution of wealth from rich to poor individuals, increased activity was observed in the lateral frontal gyrus and the anterior insula (Dawes et al., 2012). However, when expressing egalitarian preferences, only the insular cortex activations were observed.

23.5.4 Human virtues and the brain

Chapter 7 is devoted to the study of human virtues. Is there a neural basis for human virtues? Consider, for instance, the act of making a promise and keeping it, even at a cost in terms of material payoffs. This is considered to be a virtuous act in most cultures, yet it is at odds with a model based purely on self-regarding preferences. Indeed, oral and non-binding promises might have been the precursor of modern legally enforceable contracts that underpin economic exchange. However, the presence of enforceable contracts has not diminished the importance of promises in creating favorable expectations, and a great deal of modern commerce and exchange is possibly based on oral promises. In repeated games, promises can be motivated by reputational concerns. However, we know from Chapter 7 that promises and other human virtues can also be motivated by intrinsic considerations (Charness and Dufwenberg, 2006; Vanberg, 2008); we explore the neural basis of some of these considerations now.

Baumgartner et al. (2009) consider a modified two-player trust game in which the trustor (player A) can send any or part of an endowment of 2 MUs to a trustee (player B); B has no endowment. The experimenter multiplies the amount sent by A by a factor of 5 and passes it on to player B. Player B then sequentially decides to send back either half of the MUs to A, or nothing. Each set of players play three consecutive trust games. In one set of trials, player B cannot make an upfront promise on his back transfers (before A makes his decision). In another set of trials, B can choose the *strength of his promise*, i.e., promise upfront that he will send back half the MUs to A *always*, *mostly*, *sometimes*, or *never*. Almost all players choose their promise with the maximum strength, i.e., “always.”

The data revealed two clusters of players among those who had originally promised to send back a transfer. The first cluster does not send back any MUs (dishonest subjects), while the second cluster keeps its promise of sending back half the MUs (honest subjects). These clusters behaved in an identical manner, irrespective of whether the trials allowed a promise to be made or not. Thus, honest player Bs returned half their MUs to A, irrespective of having made a promise or not, while the dishonest player Bs never returned any MUs to A. An additional feature of these experiments relative to the earlier literature on deception in neuroscience studies is the introduction of an anticipation stage.¹⁹

Once player B has made a promise, there is a time interval of a few seconds that creates an anticipation in the mind of player B if player A will trust him. During the anticipation period, useful brain data can be gathered. Player Bs who intend to break their promise are likely to exhibit different brain activity relative to those who make a sincere promise that they intend to keep. This

¹⁹ For a survey of earlier neuroscience studies of deception that do not allow for promises made under free will or issues of anticipation, see Spence et al. (2004) and Sip et al. (2008).

is because dishonest players may face emotional conflict from making a promise that they do not intend to keep.

An important aim of this study is to be able to predict whether player Bs will subsequently honor their promise by looking at their brain activity at the time of making their decision. In other studies, such emotional conflict has led to increased activity in the anterior cingulate cortex (ACC) and the amygdala (Botvinick et al., 1999; Phillips et al., 2003; Sanfey et al., 2003). Furthermore, is the anticipation period for dishonest player Bs more uncertain and stressful as revealed by their fMRI scans? Finally, when dishonest player Bs break their promise, they may fear the negative consequences if their dishonesty is discovered. In previous studies, such deception has led to increased activation of the ACC.²⁰

These anticipated results are largely confirmed. In the anticipation stage, dishonest player Bs exhibited relatively greater brain activity in the right anterior insula and right inferior frontal gyrus (IFG), as compared to a baseline condition. In the decision stage, when player A makes his decision, dishonest player Bs exhibited relatively greater brain activity in the ACC as compared to a baseline condition. In another decision stage just prior to the decision of player Bs to honor their promise, when they are reminded of their original promise, dishonest player Bs show relatively greater activity in the left amygdala.²¹ Using brain activity at the time of making a promise, activation in the anterior insula, ACC, and IFG, can successfully predict if player B intends to break his promise subsequently.

23.5.5 Neural interplay of guilt and social preferences

Prosocial behavior may be caused by feelings of *guilt* that arise from antisocial behavior. We have considered issues of guilt in Part 4 of the book, in the section on psychological game theory. Here, we consider the neural correlates of guilt by describing the experiments of Chang et al. (2011), who collect fMRI data from individuals who play the trust game. They give a useful definition of guilt (p. 560): “Guilt can be conceptualized as a negative emotional state associated with the violation of a personal moral rule or a social standard (Haidt, 2003) and is particularly salient when one believes they have inflicted harm, loss, or distress on a relationship partner, for example when one fails to live up to the expectations of others (Baumeister et al., 1994).”

Using the notation of Part 4, let b_A^1 be the first order beliefs of the trustor (player A) about the amount of money that the trustee (player B) will return to him. The trustee’s second order beliefs about b_A^1 are denoted by b_B^2 and his payoff is given by

$$\pi_B = \begin{cases} m_B - \theta (b_B^2 - r) & \text{if } r \leq b_B^2 \\ m_B & \text{if } r > b_B^2 \end{cases},$$

where m_B is the final monetary amount in the ownership of player B at the end of the game, r is the amount returned by B to A, and $\theta > 0$ is a guilt-aversion parameter. Notice that guilt arises if B’s action (r) falls short of the second order beliefs of B about the action (b_B^2). Player B chooses r in order to maximize his utility. Furthermore, the beliefs are endogenous and, in

²⁰ For a focus on the neural systems underlying the trustors’ behavior in trust games see Delgado et al. (2005), King-Casas et al. (2005) and Krueger et al. (2007). van den Bos et al. (2009) show that player Bs who abuse the trust of player As in trust games, typically show enhanced activity in the vmPFC and the posterior cingulate cortex (PCC).

²¹ In this and the previous sentence, the word “relatively” captures statistical significance either at the 5% or the 1% level, based on the treatments.

equilibrium, they must concur with the actions. If the guilt-aversion parameter is high enough, in the sense that $\theta > 1$ and $r < b_B^2$, then by making an extra unit of back transfer relative to keeping it, the trustee increases his own payoff. Hence, guilt aversion may enhance prosocial behavior. On the other hand, if guilt aversion is low ($\theta < 1$) the trustee does not send back anything.

Participants in the experiments played the one-shot trust game multiple times. The first and second order beliefs of players were also elicited before they made their respective decisions. The trustees underwent fMRI scans while they made their decisions. The behavioral results are in line with the typical trust game experiments. Some amount of money was always sent by the trustors. The trustees were typically able to predict the expectations of the trustors fairly accurately. The actual amounts returned by them, r , were close to these expectations, which is consistent with the guilt-aversion hypothesis.

The fMRI results are as follows. The main contrast is between two groups of trustees for whom $r = b_B^2$ and $r < b_B^2$, respectively. When $r = b_B^2$, there is increased activity in the insula, supplementary motor area (SMA), dorsal anterior cingulate (dACC), dorsolateral prefrontal cortex (DLPFC), and the temporal parietal junction (TPJ). In contrast when $r < b_B^2$, there is increased activity in the ventromedial prefrontal cortex vmPFC, bilateral nucleus accumbens (NAcc), and dorsomedial prefrontal cortex (dmPFC).²² The individual differences between the players also give interesting insights. Trustees who reported that they would have felt guilty if they had returned less money showed more activity in the insula and in the SMA when they contributed the expected amount ($r = b_B^2$). By comparison, participants who said that they would not feel guilty by returning less than the expected amount, showed increased activity in the NAcc when they contributed lower amounts ($r < b_B^2$). Thus, activity in the insula appears to be correlated with fairer choices.

23.6 The neuroeconomics of time preferences

Animals, even primates, who are most closely related, find it difficult to delay gratification (unless the animal is pre-programmed to do so, e.g., seasonal food storage among squirrels). As Loewenstein et al. (2008) describe: “Even after long periods of training, our nearest evolutionary relatives have measured discount functions that fall in value to nearly zero after a delay of about one minute. For example, Stevens et al. (2005) report that cotton-top tamarin monkeys are unable to wait more than eight seconds to triple the value of an immediately available food reward.” In humans, individuals whose prefrontal cortex is damaged, say, by accident or surgery, are influenced more by the availability of immediate reward (Lhermitte et al., 1986). In particular, damage to the mPFC and the extent of the damage are both positively correlated with the choice of immediate over later rewards (Sellitto et al., 2010). These observations provide clues to the potential neural basis of intertemporal choice.

Let $c_0 = (c_0, c_1, \dots, c_T)$ denote the *temporal consumption profile* of the individual at time $t = 0$. Under the (β, δ) form of hyperbolic discounting, or quasi-hyperbolic discounting (see Part 3), the temporal stream $(c_0, c_1, \dots, c_T)$ is evaluated in the following manner (Laibson, 1997):

$$U^0(c_0, c_1, \dots, c_T) = u(c_0) + \beta \sum_{t=1}^T \delta^t u(c_t), \quad 0 < \beta < 1, \delta > 0. \quad (23.3)$$

²² Trustees in trust games who abuse trust show greater activity in the vmPFC and in the posterior cingulate region; see van den Bos et al. (2009).

By a multiplicative scaling, (23.3) can be written as

$$V^0 = \left(\frac{1}{\beta} - 1 \right) u(c_0) + \sum_{t=0}^T \delta^t u(c_t). \quad (23.4)$$

The first and the second terms in (23.4) can be thought of as describing, respectively, the β system and the δ system.

McClure et al. (2004) measure brain activity when participants in their experiments make a set of intertemporal choices between a smaller–earlier reward (\$ s with a time delay d) and a larger–later reward (\$ L with a time delay D); $L > s$ and $D > d$. Values of $d = 0$ (immediate gratification) and $d > 0$ were considered, while the rewards ranged from \$5 to \$40. The authors hypothesize that the limbic system is driven relatively by immediate rewards rather than future rewards while the anterior and dorsolateral regions of the prefrontal cortex are driven relatively more by future rewards.

In terms of the (β, δ) preferences of quasi-hyperbolic discounting, McClure et al. (2004) hypothesized that β is controlled by the limbic system, while δ is controlled by the prefrontal cortex. It turned out that when the choice (s, d) was made, parts of the limbic system, which the authors call β -areas, the ventral striatum, orbitofrontal cortex and the medial prefrontal cortex, were disproportionately activated. In contrast, when difficult choices are given, the δ -areas, such as the prefrontal and parietal cortex are disproportionately activated. This is because difficult choices are evaluated using the cognitive system. The δ -areas were also disproportionately activated when individuals chose the (L, D) option. The frontal and parietal cortex is activated for all (s, d) and (L, D) options; these are regions responsible for abstract planning and cognition.

In the McClure et al. (2004) experiments, the rewards were gift certificates, so one potential drawback of the approach was that gift certificates did not necessarily represent immediate rewards. This criticism was addressed in a later study by McClure et al. (2007) that we now describe. Subjects in their experiments abstained from drinks for 3 hours prior to the experiment. A choice between dated rewards was offered; the rewards being squirts of juice or water. The typical choice was of the following form. Choose (x, D) or (x', D') , where x and x' are both either squirts of juice or squirts of water and D, D' were, respectively, time delays expressed in minutes. In the experiments, the values of (x, x') used were (1, 2), (1, 3), and (2, 3) while the values of (D, D') used were (0, 1), (0, 5), (10, 11), (10, 15), (20, 21), (20, 25).

The behavioral data is shown in Figure 23.25, where the vertical axis is the probability of choosing the earlier reward, x , and D is the earlier time. The probability of choosing the earlier reward is always higher when the early time, D , is closer to the present; these results are consistent with models of present-biased preferences such as hyperbolic discounting. In the right hand side panel in Figure 23.25, it is clear that the results are mainly driven by the cases where the time delay difference $D' - D$ is 5 minutes; there is a smaller effect when $D' - D$ is 1 minute.

The fMRI data is shown in Figure 23.26. It is not based on a region of interest analysis, but rather on a whole brain analysis in which the BOLD signal measure in each voxel is regressed on a number of variables such as the reward magnitude and time delay. The brain regions that were most activated when subjects chose the earlier option were classified as the β areas (top panel of Figure 23.26). These brain regions include the nucleus accumbens (NAcc), subgenual cingulate cortex (SGC), medial orbitofrontal cortex (mOFC), posterior cingulate cortex (PCC), and precuneus. As in McClure et al. (2004), these regions lie within the limbic system and the paralimbic




Figure 23.25 The behavioral data.

Source: McClure, S. M., Ericson, K. M., Laibson, D. I., Loewenstein, G., and Cohen, J. D. (2007). "Time discounting from primary rewards." *Journal of Neuroscience* 27: 5796–804.




Figure 23.26 Activation in the β areas and the δ areas in the brain as rewards with varying delays are presented to the subjects.

Source: McClure, S. M., Ericson, K. M., Laibson, D. I., Loewenstein, G., and Cohen, J. D. (2007). "Time discounting for primary rewards." *Journal of Neuroscience* 27: 5796–804.

cortex. In contrast, the brain regions that are correlated with the larger-later choices, the δ areas (bottom panel of Figure 23.26), include regions associated with higher-level cognitive functions such as the posterior cingulate cortex (PCC), bilateral areas in the posterior parietal cortex (pPar), bilateral areas in the anterior insula (Ant Ins), and several regions in the dorsolateral prefrontal cortex (dlPFC) (Brodmann areas 9, 44, 46, and 10). Other areas that were activated included primary visual and motor cortex areas, possibly because subjects made choices by looking at a visual display. In comparing the results in McClure et al. (2004) and McClure et al. (2007), there is a greater match between the two in the δ -areas, but a lower match in the β -areas, suggesting that the β -areas might be context and task dependent.

Other studies have questioned the findings of McClure et al. (2004, 2007). Glimcher et al. (2007) do not find increased activation of the limbic and paralimbic systems in the case of immediate rewards. Kable and Glimcher (2007) find that the β -areas in McClure et al. (2007) are activated when the delayed reward, x' , changes, thus, they question if these areas are exclusively associated with immediate rewards. For some possible explanations for these contrary findings, see Loewenstein et al. (2008).

A second class of models that explains the choice between immediate and delayed rewards invokes conflict between two brain systems for control. Two main types of models in this class are the planner–doer model of Thaler and Shefrin (1981) and the dual-selves model of Fudenberg and Levine (2006); see Parts 3 and 6 for the details. Hare et al. (2009) confront subjects with 50 different kinds of foods and ask them to rate them on health and taste. This allows the identification of a reference food for each subject that is neutral on taste and health grounds. In a second task, subjects are asked to choose between each of the 49 food items and the reference food item; fMRI is used to monitor their brain activity as they make their choices. The subjects were then classified into two groups. Self-controllers (SC) use health and taste criteria to make choices while the non self-controllers (NSC) decided on the basis of taste alone.

The behavioral results show that the subjects in the SC group were more successful in avoiding liked-unhealthy foods relative to the NSC group. The SC group also chose liked-healthy foods more frequently than the NSC group. The dlPFC shows greater activity in the choices of SC subjects who avoid unhealthy foods and also shows greater activity when the reference food item is chosen over the tempting food items. Thus, the dlPFC appears to play the role of a patient long-run planner. This interpretation is supported in the results of Figner et al. (2010) who find that temporary disruption of the dlPFC by TMS induces greater impatience in subjects. Relative behavioral impulsivity in adolescents is also associated with slower development of the neural brain circuitry in the prefrontal cortex (Sowell et al., 2003; Casey et al. 2008). This provides some neural basis for why younger people might make more impatient choices, which has economic implications for them such as higher car insurance premiums.

23.7 The neuroeconomics of strategic interaction

In this section, we consider the neural evidence on strategic interaction. The direct cognitive evidence from eyetracking studies, typically using the MOUSELAB software, that supports level- k and cognitive hierarchy (CH) models is considered in Part 4 and is not repeated here (Crawford, 2008; Wang et al., 2010).

In Bhatt and Camerer (2005), subjects play eight, two-player, dominance solvable games with 2–4 strategies. Players engage in three tasks in random order. In the *choice task* (C), they choose their own strategy. In the *first order belief task* (B), they guess the strategy choice of the opponent. In the *second order belief task* (2B), they guess the opponent's first order beliefs about them. In the classical game theoretic equilibrium, choices, first order beliefs, and second order beliefs, should align, i.e., beliefs should be confirmed in equilibrium and choices should be optimal given beliefs. Bhatt and Camerer hypothesize that when subjects play according to equilibrium, neural activity in the three tasks C, B, and 2B should be similar. However, when the choices of subjects are not in equilibrium, say, because they follow level- k or CH reasoning (see Part 4), then dissimilar areas of the brain are likely to be activated. For instance, nonstrategic or level-0 players might not spend much time in forming beliefs. Consistent with these predictions, the authors find that in task C, when subjects were out of equilibrium, there is relatively greater activity in the mPFC and

dlPFC. In contrast, for equilibrium subjects, brain activity was confined largely to similar brain areas, except for a small area in the ventral striatum.

In game theory, it is critical to determine what *theory of the mind* players have about others, i.e., how players believe others feel, believe, and choose. One line of research seeks to identify such *mentalizing* on the part of players by comparing their behavior when they play against other humans, relative to playing against computer programs. The hypothesis is that humans may need to mentalize against other humans, but not computers; this is confirmed in several neuroeconomic studies (Gallagher et al., 2002; McCabe et al., 2001; Coricelli and Nagel, 2009).²³

In Coricelli and Nagel (2009), subjects play a *p*-beauty contest. They simultaneously choose a number in the interval $[0, 100]$ and the winning individual's guess is closest to p times the average guess of all individuals. For $p < 1$, in a Nash equilibrium, all subjects are predicted to choose 0. By contrast, a level- k individual chooses $p^k 50$, $k = 0, 1, 2, \dots$ (see Part 4 for details). For instance, when $p = \frac{2}{3}$, the optimal guess of a (i) level-0 individual is 50, (ii) level-1 individual is $\frac{2}{3} 50 = 33.3$, (iii) level-2 individual is $(\frac{2}{3})^2 50 = 22.2$, and so on. In the human condition, subjects were matched with other humans and in the computer condition, they were matched with computers who were pre-programmed to choose a number in the interval $[0, 100]$. Various values of p were used.

Figure 23.27 shows the results (the letter M in the figure is to be read as the probability p). The behavioral results are shown in panel A for a representative individual in each category—low level of reasoning (level-1 against a human opponent) and a high level of reasoning (level-2 against a human opponent) for varying values of p (data shown by filled-in dots). Although their behavior was different against human opponents, both types chose to behave like level-1 individuals against the computer opponent (data shown in triangles). The level-1 prediction is shown as the dotted line with triangles and the level-2 prediction as the dotted line with dots; in panel A, below $p = 1$, the dotted line with dots is below the dotted line with triangles, and above $p = 1$, the line with triangles cross over. The data for the two types of individuals for a particular value of $p = \frac{2}{3}$ is shown in the bottom figure in panel B. Overall, across all individuals, the majority of choices were consistent with level-1 to level-3 players and only 5% of the players were higher than level-3.

Panel B in Figure 23.27 shows the fMRI data. In the human versus computer contrast, enhanced activity was found in the mPFC, rostral anterior cingulate cortex (rACC), superior temporal sulcus (STS), posterior cingulate cortex, and bilateral temporoparietal junction (TPJ). However, the crucial finding came from comparing the data for low and high level reasoning subjects. In this case, in the contrast between the human and the computer treatments, activity in the mPFC was found in the high level reasoners but not in the low level reasoners. Two areas in the mPFC stood out. A more dorsal area with coordinates $(x, y, z) = (0, 48, 24)$ and a more ventral area with coordinates $(x, y, z) = (9, 48, -6)$. The area $(0, 48, 24)$ has been found to be important when humans take a third person perspective (Amodio and Frith, 2006; D'Argembeau et al., 2007) and the area $(9, 48, -6)$ has been found important in thinking about like-minded other people (Mitchell et al., 2005; Mitchell et al., 2006; Jenkins et al., 2008). While the *p*-beauty contest is a simultaneous move game, the mPFC has also been implicated in a theory of the mind when players make sequential choices (Hampton et al., 2008).

In contrast, the fMRI evidence indicates that the main brain area engaged in the decisions of low level reasoners appears to be the rACC that has been implicated in self-referential thinking in social cognitive tasks (Moran et al., 2006). Two areas of the brain showed increased activity, irrespective of the level of reasoning of the players; these are the TPJ and the superior temporal sulcus.

²³ One may conjecture that humans need to mentalize against computers too because humans wrote the software that determines the computer's choices. However, whether humans do so or not is an empirical question.




Figure 23.27 Results in the p-beauty contest game. Behavioral results (panel a) and fMRI results (panel b). M is to be read as the probability p .

Source: Coricelli, G. and Nagel, R. (2009). "Neural correlates of depth of strategic reasoning in medial prefrontal cortex." *Proceedings of the National Academy of Sciences* 106(23): 9163–8.

The TPJ has been implicated in a theory of the mind in other studies (Saxe and Wexler, 2005), while the superior temporal sulcus has been implicated in assigning intentions to the actions of others and establishing causality (Brunet et al., 2000; Castelli et al., 2000). However, several other empirical studies do not necessarily support the conclusions arrived at here. For instance, Kuo et al. (2009) show that in other contexts, the mPFC does not correlate with steps of reasoning in others. Similarly, the TPJ has been implicated in other studies in tasks that are unrelated to the theory of the mind (Decety and Lamm, 2007; Mitchell, 2008).

Bhatt et al. (2010) consider issues of deception or *management of social image* as with poker players who are engaged in bluffing. They use a two-player bargaining environment in which players must form first and second order beliefs about models of the mind. In the experiments, a buyer and a seller play 60 rounds of a bargaining game. At the beginning of each round, the buyer's private valuation, $v \in \{1, 2, \dots, 10\}$, is privately revealed to the buyer. The buyer then suggests a price to the seller, s , who is free to ignore the suggestion if he wishes to. Finally, the seller asks for a price, p . If $v \geq p$, then trade takes place with respective payoffs to the seller and

the buyer equal to p and $v - p$. Otherwise, no trade takes place and both parties receive 0. Since the interests of the two parties are completely non-aligned, the prediction of standard economic analysis is that the buyer will babble and the seller will pay no attention to the suggestions.

A regression of the suggested price, s , on the buyer's valuation, v , reveals three kinds of buyers. *Conservationists* play with their cards close to their chest and their suggestions do not reveal valuations; their behavior is closest to the standard game theoretic prediction. *Incrementalists* demonstrated a strong positive correlation between v and s ; this behavior is potentially efficiency enhancing if believed by the seller. Finally, *strategists* reveal a significant negative correlation between v and s . This behavior is consistent with buyers who are level-2 players who believe that the seller is level-1.²⁴

As compared to the other types, there is relatively greater activity in the dlPFC of strategists when they make their suggestions. Furthermore, when their realized value, v , is higher, strategists need to bluff the most (because v and s are strongly negative correlated); in this case they exhibit greater activity in the superior temporal sulcus (STS). In a later paper, Bhatt et al. (2012) show that the absolute gap between the actual price asked by the seller, p , and the buyer's suggested price, s , reflected the degree of suspicion that sellers have about the suggestions. There was relatively greater activation in the amygdala when more suspicious sellers choose prices.

23.8 Pharmacoeconomics: an application to social effects of oxytocin

There has been much interest in studying the economic effects of *neuromodulator systems*. In classical synaptic transmission, described in Section 23.2, information is transmitted from a single neuron to another neuron. In a different kind of information transmission, a small group of neurons can secrete neuromodulators in a wide and diffuse area of the brain, simultaneously influencing information transmission to many neurons; these effects can be long lasting. Examples of neuromodulators include neurotransmitters such as *dopamine*, *serotonin*, *norepinephrine*, as well as hormones such as *testosterone* and *oxytocin*. Crockett and Fehr (2014) provide an excellent and accessible survey on the economic effects of these neuromodulators. In this section, we confine ourselves to a brief outline of some of the social effects of oxytocin.

In their seminal study using a trust game experiment, Kosfeld et al. (2005) gave one group of subjects a single dose of intranasal oxytocin, while the other group received a placebo. Both players were initially endowed with 12 MUs. The trustor/investor could pass on any amount from the set $\{0, 4, 8, 12\}$ to the trustee; the experimenter tripled the amount before it was received by the trustee. The aim was to examine the effect of oxytocin on the behavior of trustors and trustees. In a second set of experiments, the risk experiments, trustees played against an inanimate random mechanism. This random mechanism replicated the trustees' decision in trust experiments, say, by randomly choosing from the distribution of back transfers in the trust game.

The transfers in the trust and risk experiments are shown, respectively, in panels a and b in Figure 23.28. Filled-in bars and blank bars represent, respectively, subjects who were given oxytocin and the placebo. In the trust experiments, subjects who are given oxytocin make, on average, 17% higher transfers relative to the placebo group. The maximum transfer of 12 MUs was made by 45% of the subjects in the oxytocin group and 6% in the placebo group. At the bottom end

²⁴ Level-1 sellers will look at the past history in judging current suggestions; if strategists behaved like conservationists, then a low suggestion will be interpreted as a bluff, hence, level-2 strategists will mix their suggestions.




Figure 23.28 Transfers in the trust experiments (panel a) and risk experiments (panel b). Filled-in bars represent subjects who were given oxytocin and the blank bars represent subjects who were given a placebo.

Source: Reprinted by permission from Macmillan Publishers Ltd. *Nature* 435: 673–6. Kosfeld, Michael, Markus Heinrichs, Paul J. Zak, Urs Fischbacher, and Ernst Fehr (2005). "Oxytocin increases trust in humans."

of transfers, 21% of the oxytocin group made transfers less than 8 MUs while 45% of the placebo group made similarly low transfers.

The authors then engage in an exemplary experimental design that further illuminates the role of oxytocin. Do the trustors transfer more in the oxytocin group because they perceive "lower risk in a social interaction" or just "lower risk"? The risk experiments that do not involve any social interaction are able to answer this question. In the risk experiments, there is no significant difference in the transfers made between the oxytocin and placebo groups. Thus, oxytocin appears to enhance trust only in social interactions. Could the results be caused by other factors such as moods and calmness differences? In order to test this, measures of moods and calmness were taken before and after the administration of oxytocin but no significant differences were found. Does oxytocin cause a general increase in prosocial behavior? In order to test this, one can examine the back transfers of trustees, which were not found to be statistically different among the oxytocin and the placebo groups. Thus, oxytocin does not lead to a general increase in prosocial behavior, but, specifically, an improvement in trusting behavior. Is the observed effect of oxytocin caused by a general increase in optimism? In order to test this, the beliefs of trustors about the back transfers of trustees in the oxytocin and placebo groups are elicited but no differences are found in these beliefs. Thus, the results are not driven by an increase in optimism.

Subsequent research has examined several social effects of oxytocin. In addition to trusting behavior, these include interpersonal relations, generosity, cooperation, perceptions of trustworthiness, attractiveness, approachability, attachment, and socio-emotional responses, such as envy and social aversion; the interested reader can look up the associated references in the meta-study

by Bartz et al. (2011). Bartz et al. (2011) also note that the overall social effects of oxytocin are small and heavily context dependent. Their solution is to advocate greater recognition of heterogeneous frames and contexts as well as individual-level heterogeneity. For instance, the prosocial effects of oxytocin are moderated, disappear, or reverse if one's opponents are portrayed as untrustworthy or are perceived to be outgroup members (Mikolajczak et al., 2010; Declerck et al., 2010; De Dreu et al., 2010).

Israel et al. (2014) examine the effect of oxytocin on the ability to detect deception, based on facial expression, body language, tone of voice, and other non-verbal clues. Two groups of subjects are shown clips of the game show "Friend or Foe" in which two individuals engage in a non-binding discussion prior to playing a prisoner's dilemma type of game; each player in the game has two strategies, friend (cooperate) or foe (defect). The clip is paused just before the outcome is revealed and subjects are asked to guess the outcome. One group is exposed to oxytocin and the second group is exposed to a placebo. Oxytocin subjects were more likely to be deceived as compared to placebo subjects, i.e., their accuracy in predicting deception was lower. However, oxytocin and placebo subjects made guesses of similar accuracy when predicting the cooperative behavior of others. Placebo subjects, but not oxytocin subjects, were able to predict the decisions of the contestants better than chance. Thus, oxytocin appears to reduce social vigilance that can reduce their payoffs.

REFERENCES FOR PART 9

- al-Nowaihi, A., Bradley, I., and Dhami, S. (2008). A note on the utility function under prospect theory. *Economics Letters* 99(2): 337–9.
- al-Nowaihi, A. and Dhami, S. (2006). A simple derivation of Prelec's probability weighting function. *Journal of Mathematical Psychology* 50(6): 521–4.
- Amodio, D. M. and Frith, C. D. (2006). Meeting of minds: the medial frontal cortex and social cognition. *Nature Reviews Neuroscience* 7: 268–77.
- Andreoni, J. and Bernheim, B. (2009). Social image and the 50–50 norm: a theoretical and experimental analysis of audience effects. *Econometrica* 77: 1607–36.
- Bartz, J. A., Zaki, J., Bolger, N., and Ochsner, K. N. (2011). Social effects of oxytocin in humans: context and person matter. *Trends in Cognitive Sciences* 15(7): 301–9.
- Bault, N., Joffily, M., Rustichini, A., and Coricelli, G. (2011). Medial prefrontal cortex and striatum mediate the influence of social comparison on the decision process. *Proceedings of the National Academy of Sciences* 108: 16044–9.
- Baumeister, R. F., Stillwell, A. M., and Heatherton, T. F. (1994). Guilt: an interpersonal approach. *Psychological Bulletin* 115: 243–67.
- Baumgartner, T., Fischbacher, U., Feierabend, A., Lutz, K., and Fehr, E. (2009). The neural circuitry of a broken promise. *Neuron* 64: 756–70.
- Baumgartner, T., Goette, L., Gügler, R., and Fehr, E. (2012). The mentalizing network orchestrates the impact of parochial altruism on social norm enforcement. *Human Brain Mapping* 33: 1452–69.
- Baumgartner, T., Knoch, D., Hotz, P., Eisenegger, C., and Fehr, E. (2011). Dorsolateral and ventromedial prefrontal cortex orchestrate normative choice. *Nature Neuroscience* 14: 1468–74.
- Bear, M. F., Connors, B. W., and Paradiso, M. A. (2007). *Neuroscience: Exploring the Brain*. Lippincott Williams and Wilkins.
- Bernheim, B. D. (2008). Neuroeconomics: a sober (but hopeful) appraisal. NBER Working Paper No. 13954.
- Berns, G., Capra, C., Chappelow, J., Moore, S., and Noussair, C. (2008). Nonlinear neurobiological probability weighting functions for aversive outcomes. *NeuroImage* 39: 2047–57.
- Berns, G. S., McClure S. M., Pagnoni G., and Montague P. R. (2001). Predictability modulates human response to reward. *Journal of Neuroscience* 21: 2793–8.
- Bhatt, M. and Camerer, C. F. (2005). Self-referential thinking and equilibrium as states of mind in games: fMRI evidence. *Games and Economic Behavior* 52: 424–59.
- Bhatt, M., Lohrenz, T., Camerer, C., and Montague, R. (2010). Neural signatures of strategic types in a two-person bargaining game. *Proceedings of the National Academy of Sciences* 107(46): 19720–5.
- Bhatt, M., Lohrenz, T., Camerer, C., and Montague, R. (2012). Distinct contributions of the amygdala and parahippocampal gyrus to suspicion and uncertainty in a repeated bargaining game. *Proceedings of the National Academy of Sciences* 109(22): 8728–33.
- Bickart, K. C., Dickerson, B. C., and Barrett, L. F. (2014). The amygdala as a hub in brain networks that support social life. *Neuropsychologia* 63: 235–48.
- Bickart, K. C., Wright, C. I., Dautoff, R. J., Dickerson, B. C., and Barrett, L. F. (2011). Amygdala volume and social network size in humans. *Nature Neuroscience* 14: 163–4.
- Blood, A. J. and Zatorre, R. J. (2001). Intensely pleasurable responses to music correlate with activity in brain regions implicated in reward and emotion. *Proceedings of the National Academy of Sciences USA* 98: 11818–23.
- Bolton, G. E. and Ockenfels, A. (2000). ERC: a theory of equity, reciprocity, and competition. *American Economic Review* 90(1): 166–93.
- Botvinick, M., Nystrom, L. E., Fissell, K., Carter, C. S., and Cohen, J. D. (1999). Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature* 402: 179–81.
- Bowles, S. and Gintis, H. (2004). The evolution of strong reciprocity. *Theoretical Population Biology* 65: 17–28.
- Bowles, S. and Gintis, H. (2011). *A Cooperative Species: Human Reciprocity and Its Evolution*. Princeton, NJ: Princeton University Press.
- Boyd, R., Gintis, H., Bowles, S., and Richerson, P. J. (2003). Evolution of altruistic punishment.

- Proceedings of the National Academy of Sciences USA* 100: 3531–5.
- Breedlove, S. M., Watson, N. V., and Rosenzweig, M. R. (2014). *Biological Psychology: An Introduction to Behavioral, Cognitive and Clinical Neuroscience*. Sinauer Associates, Inc. Publishers.
- Breiter, H. C., Aharon, I., Kahneman, D., Dale, A., and Shizgal, P. (2001). Functional imaging of neural responses to expectancy and experience of monetary gains and losses. *Neuron* 30: 619–39.
- Brodmann, K. (1909). *Vergleichende Lokalisationstheorie der Grosshirnrinde* (in German). Leipzig: Johann Ambrosius Barth. English translation by Gary, L. (1999). *Brodmann's Localisation in the Cerebral Cortex*. Imperial College Press: World Scientific Publishers.
- Brunet, E., Sarfati, Y., Hardy-Bayle, M. C., and Decety, J. (2000). A PET investigation of the attribution of intentions with a nonverbal task. *NeuroImage* 11: 157–66.
- Buckholtz, J. W., Asplund C. L., and Dux, P. E. (2008). The neural correlates of third-party punishment. *Neuron* 60(5): 930–40.
- Busemeyer, J. R. and Diederich, A. (2014). Estimation and testing of computational psychological models. Chapter 4 in P. W. Glimcher and E. Fehr (eds.), *Neuroeconomics*. Elsevier Inc.
- Camerer, C. F. (2007). Neuroeconomics: using neuroscience to make economic predictions. *Economic Journal* 117: C26–C42.
- Camerer, C. F. (2008). The case for mindful economics. In A. Caplin and A. Schotter (eds.), *The Foundations of Positive and Normative Economics*. Oxford: Oxford University Press.
- Camerer, C. F. (2013). Review essay about *Foundations of Neuroeconomic Analysis* by Paul Glimcher. *Journal of Economic Literature* 51(4): 1155–82.
- Camerer, C. F., Cohen, J., Fehr, E., Glimcher, P., and Laibson, D. (2015). Neuroeconomics. In: J. Kagel and A. Roth (eds.), *Handbook of Experimental Economics*. Forthcoming.
- Camerer, C. F., Loewenstein, G., and Prelec, D. (2004). Neuroeconomics: why economics needs brains. *Scandinavian Journal of Economics* 106: 555–79.
- Camerer, C. F., Loewenstein, G., and Prelec, D. (2005). Neuroeconomics: how neuroscience can inform economics. *Journal of Economic Literature* 43: 9–64.
- Caplin, A. and Dean, M. (2008). Axiomatic neuroeconomics. In P. W. Glimcher, C. F. Camerer, E. Fehr, and R. A. Poldrack (eds.), *Neuroeconomics: Decision Making and the Brain*. London and San Diego: Elsevier, Academic Press, pp. 21–32.
- Caplin, A., Dean, M., Glimcher, P. W., and Rutledge, R. B. (2010). Measuring beliefs and rewards: a neuroeconomic approach. *Quarterly Journal of Economics* 125(3): 923–60.
- Cappelen, A. W., Eichele, T., Hugdahl, K., Karsten Specht, K., Sørensen, E. Ø., and Tungodden, B. (2014). Equity theory and fair inequality: a neuroeconomic study. *Proceedings of the National Academy of Sciences USA* 111(43): 15368–72.
- Casey, B. J., Getz, S., and Galvan, A. (2008). The adolescent brain. *Development Review* 28: 62–77.
- Castelli, F., Happe, E., Frith, U., and Frith, C. (2000). Movement and mind: a functional imaging study of perception and interpretation of complex intentional movement patterns. *NeuroImage* 12: 314–25.
- Chang, L. J., Smith, A., Dufwenberg, M., and Sanfey, A. G. (2011). Triangulating the neural, psychological, and economic bases of guilt aversion. *Neuron* 70(3): 560–72.
- Charness, G. and Dufwenberg, M. (2006). Promises and partnership. *Econometrica* 74(6): 1579–1601.
- Chib, V. S., Rangel, A., Shimojo, S., and O'Doherty, J. P. (2009). Evidence for a common representation of decision values for dissimilar goods in human ventromedial prefrontal cortex. *Journal of Neuroscience* 29: 12315–20.
- Clithero, J. A. and Rangel, A. (2013). Combining response times and choice data using a neuroeconomic model of the decision process improves out-of-sample predictions. Unpublished manuscript.
- Coricelli, G. and Nagel, R. (2009). Neural correlates of depth of strategic reasoning in medial prefrontal cortex. *Proceedings of the National Academy of Sciences* 106(23): 9163–8.
- Crawford, V. P. (2008). Look-ups as the windows of the strategic soul: studying cognition via information search in game experiments. In A. Caplin and A. Schotter (eds.), *The Foundations of Positive and Normative Economics*. New York: Oxford University Press, pp. 249–80.
- Crockett, M. J. and Fehr, E. (2014). Pharmacology of economic and social decision making.

- In P. W. Glimcher and E. Fehr (eds.), *Neuroeconomics*. Elsevier Inc, pp. 259–79.
- Cunningham, W. A., Kesek, A., and Mowrer, S. M. (2009). Distinct orbitofrontal regions encode stimulus and choice valuation. *Journal of Cognitive Neuroscience* 21, 1956–66.
- D'Argembeau, A., Ruby, P., Collette, F., et al. (2007). Distinct regions of the medial prefrontal cortex are associated with self-referential processing and perspective taking. *Journal of Cognitive Neuroscience* 19: 935–44.
- Dawes, C. T., Loewen, P. J., Schreiber, D., et al. (2012). Neural basis of egalitarian behavior. *Proceedings of the National Academy of Sciences USA* 109(17): 6479–83.
- de Araujo, I. E., Rolls, E. T., Kringelbach, M. L., McGlone, F., and Phillips, N. (2003). Taste-olfactory convergence, and the representation of the pleasantness of flavour, in the human brain. *European Journal of Neuroscience* 18: 2059–68.
- De Dreu, C. K., Greer, L. L., Handgraaf, M. J., et al. (2010). The neuropeptide oxytocin regulates parochial altruism in intergroup conflict among humans. *Science* 328: 1408–11.
- De Martino, B., Camerer, C. F., and Adolphs R. (2010). Amygdala damage eliminates monetary loss aversion. *Proceedings of the National Academy of Sciences USA* 107: 3788–92.
- De Martino, B., Kumaran, D., Holt, B., and Dolan, R. J. (2009). The neurobiology of reference-dependent value computations. *Journal of Neuroscience* 29: 3833–42.
- De Martino, B., Kumaran, D., Seymour, B., and Dolan R. J. (2006). Frames, biases, and rational decision-making in the human brain. *Science* 313: 684–7.
- Decety, J. and Lamm, C. (2007). The role of the right temporoparietal junction in social interaction: how low-level computational processes contribute to meta-cognition. *Neuroscientist* 13: 580–93.
- Declerck, C. H., Boone, C., and Kiyonari, T. (2010). Oxytocin and cooperation under uncertainty: the moderating influence of incentives and social information. *Hormones and Behavior* 57: 368–74.
- Delgado, M. R., Frank, R. H., and Phelps, E. A. (2005). Perceptions of moral character modulate the neural systems of reward during the trust game. *Nature Neuroscience* 8: 1611–18.
- Delgado, M. R., Locke, H. M., Stenger, V. A., and Fiez, J. A. (2003). Dorsal striatum responses to reward and punishment: effects of valence and magnitude manipulations. *Cognitive, Affective, and Behavioral Neuroscience* 3(1): 27–38.
- Delgado, M. R., Nystrom, L. E., Fissell, C., Noll, D. C., and Fiez, J. A. (2000). Tracking the hemodynamic responses to reward and punishment in the striatum. *Journal of Neurophysiology* 84: 3072–7.
- De Quervain, D., Fischbacher, U., Treyer, V., et al. (2004). The neural basis of altruistic punishment. *Science* 305(5688): 1254–8.
- Falk, A., Fehr, E., and Fischbacher, U. (2008). Testing theories of fairness and reciprocity—intentions matter. *Games and Economic Behavior* 62(1): 287–303.
- Fecteau, S., Knoch, D., Fregni, E., Sultani, N., Boggio, P., and Pascual-Leone, A. (2007). Diminishing risk-taking behavior by modulating activity in the prefrontal cortex: a direct current stimulation study. *Journal of Neuroscience* 27: 12500–5.
- Fehr, E. and Camerer, C. F. (2007). Social neuroeconomics: the neural circuitry of social preferences. *Trends in Cognitive Sciences* 11: 419–27.
- Fehr, E. and Fischbacher, U. (2003). The nature of human altruism. *Nature* 425: 785–91.
- Fehr, E., Fischbacher, U., and Gächter, S. (2002). Strong reciprocity, human cooperation, and the enforcement of social norms. *Human Nature* 13(1): 1–25.
- Fehr, E. and Gächter, S. (2002). Altruistic punishment in humans. *Nature* 415: 137–40.
- Fehr, E. and Krajbich, I. (2014). Social preferences and the brain. Chapter 11 in P. W. Glimcher and E. Fehr (eds.), *Neuroeconomics*. Elsevier Inc.
- Fehr, E. and Rangel, A. (2011). Neuroeconomic foundations of economic choice—recent advances. *Journal of Economic Perspectives* 25(4): 3–30.
- Fehr, E. and Schmidt K. M. (1999). A theory of fairness, competition and cooperation. *Quarterly Journal of Economics* 114(3): 817–68.
- Figner, B., Knoch, D., Johnson, E. J., et al. (2010). Lateral pre-frontal cortex and self-control in intertemporal choice. *Nature Neuroscience* 13: 538–39.
- Firth, C. D. and Frith, U. (2007). Social cognition in humans. *Current Biology* 17: R724–R732.
- Fritsch, G. and Hitzig, E. (1870). Ueber die elektrische Erregbarkeit des Grosshirns [On the

- electrical excitability of the cerebrum]. *Archiv für Anatomie und Physiologie* 37: 300–32.
- Fudenberg, D. and Levine, D. (2006). A dual self model of impulse control. *American Economic Review* 95: 1449–76.
- Gallagher H. L. and Frith C. D. (2003). Functional imaging of “theory of mind.” *Trends in Cognitive Science* 7: 77–83.
- Gallagher, H. L., Jack, A. I., Poepstorff, A., and Frith, C. D. (2002). Imaging the intentional stance in a competitive game. *NeuroImage* 16: 814–21.
- Gazzaniga, M. S. (2009). *The Cognitive Neurosciences*. Cambridge, MA: MIT Press.
- Glimcher, P. (2011). *Foundations of Neuroeconomic Analysis*. Oxford: Oxford University Press.
- Glimcher, P. W., Camerer, C. F., Fehr, E., and Pol-drack, R. A. (2009). *Neuroeconomics*. Academic Press, Elsevier Inc.
- Glimcher, P. W. and Fehr, E. (eds.). (2014). *Neuroeconomics*. Elsevier Inc.
- Glimcher, P. W., Kable J., and Louie, K. (2007). Neuroeconomic studies of impulsivity: now or just as soon as possible? *American Economic Review* 97: 142–47.
- Goette, L., Huffman, D., and Meier, S. (2006). The impact of group membership on cooperation and norm enforcement: Evidence using random assignment to real social groups. *American Economic Review* 96: 212–16.
- Gospic, K., Mohlin, E., Fransson, P., Petrovic, P., Johansson, M., and Ingvar, M. (2011). Limbic justice: amygdala involvement in immediate rejection in the ultimatum game. *PLoS Biology* 9(5): e1001054.
- Gul, F. and Pesendorfer, W. (2008). The case for mindless economics. In A. Caplin and A. Schotter (eds.), *The Foundations of Positive and Normative Economics*. Oxford: Oxford University Press.
- Haidt, J. (2003). The moral emotions. In R. J. Davidson, K. R. Scherer, and H. H. Goldsmith (eds.), *Handbook of Affective Sciences* Oxford: Oxford University Press, pp. 852–70.
- Hampton, A., Bossaerts, P., and O’Doherty, J. (2008). Neural correlates of mentalizing-related computations during strategic interactions in humans. *Proceedings of the National Academy of Sciences* 105(18): 6741–6.
- Hare, T., Camerer, C., and Rangel, A. (2009). Self-control in decision-making involves modulation of the vmPFC valuation system. *Science* 324: 646–8.
- Hare, T. A., O’Doherty, J., Camerer, C. F., Schultz, W., and Rangel, A. (2008). Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. *Journal of Neuroscience* 28: 5623–30.
- Harlow, J. M. (1868). Recovery from the Passage of an Iron Bar through the Head. *Publications of the Massachusetts Medical Society* 2: 327–47. Reprinted as Recovery from the Passage of an Iron Bar through the Head (David Clapp and Son, 1869).
- Herwig, U., Baumgartner, T., Kaffenberger, T., et al. (2007). Modulation of anticipatory emotion and perception processing by cognitive control. *Neuroimage* 37: 652–62.
- Hollerman, J. R., Tremblay, L., and Schultz, W. (1998). Influence of reward expectation on behavior-related neuronal activity in primate striatum. *Journal of Neurophysiology* 80(2): 947–63.
- Houser, D. and McCabe, K. (2014). Experimental economics and experimental game theory. Chapter 2 in P. W. Glimcher and E. Fehr (eds.), *Neuroeconomics*. Elsevier Inc.
- Hsu, M., Anen, C., and Quartz, S. R. (2008). The right and the good: distributive justice and neural encoding of equity and efficiency. *Science* 320: 1092–5.
- Hsu, M., Krajbich, I., Zhao, C., and Camerer, C. F. (2009). Neural response to reward anticipation under risk is nonlinear in probabilities. *Journal of Neuroscience* 29(7): 2231–7.
- Huettel, S. A., Song, A. W., and McCarthy, G. (2009). *Functional Magnetic Resonance Imaging*. 2nd edition. Sunderland, MA: Sinauer Associates.
- Hutzler, F. (2014). Reverse inference is not a fallacy per se: cognitive processes can be inferred from functional imaging data. *NeuroImage* 84(1): 1061–9.
- Israel, S., Hart, E., and Winter, E. (2014). Oxytocin decreases accuracy in the perception of social deception. *Psychological Science* 25(1): 293–5.
- Izuma, K., Saito, D. N., and Sadato, N. (2008). Processing of social and monetary rewards in the human striatum. *Neuron* 58: 284–94.
- Izuma, K., Saito, D. N., and Sadato, N. (2010). Processing of the incentive for social approval in

- the ventral striatum during charitable donation. *Journal of Cognitive Neuroscience* 22: 621–31.
- Jenison, R. L., Rangel, A., Oya, H., Kawasaki, H. and Howard, M. A. (2011). Value encoding in single neurons in the human amygdala during decision making. *Journal of Neuroscience* 31(1): 331–8.
- Jenkins, A. C., Macrae, C. N., and Mitchell, J. P. (2008). Repetition suppression of ventromedial prefrontal activity during judgments of self and others. *Proceedings of the National Academy of Sciences USA* 105: 4507–12.
- Kable, J. W. and Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nature Neuroscience* 10: 1625–33.
- Kable, J. W. and Glimcher, P. W. (2009). The neurobiology of decision: consensus and controversy. *Neuron* 63: 733–45.
- Kahn, I., Yeshurun, Y., Rotshtein, P., Fried, I., Ben-Bashat, D., and Hendler, T. (2002). The role of the amygdala in signaling prospective outcome of choice. *Neuron* 33(6): 983–94.
- Kahneman, D. (2011). *Thinking Fast and Slow*. New York: Farrar, Strauss, Giroux.
- Kahneman, D. and Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica* 47: 263–91.
- Keysers, C. et al. (2004). A touching sight: SII/PV activation during the observation and experience of touch. *Neuron* 42: 335–46.
- King-Casas, B., Tomlin, D., Anen, C., Camerer, C. F., Quartz, S. R., and Montague, P. R. (2005). Getting to know you: reputation and trust in a two person economic exchange. *Science* 308: 78–83.
- Knoch, D., Nitsche, M. A., Fischbacher, U., Eisenegger, C., Pascual-Leone, A., and Fehr, E. (2008). Studying the neurobiology of social interaction with transcranial direct current stimulation: the example of punishing unfairness. *Cerebral Cortex* 18: 1987–90.
- Knoch, D., Pascual-Leone, A., Meyer, K., Treyer, V., and Fehr, E. (2006). Diminishing reciprocal fairness by disrupting the right prefrontal cortex. *Science* 314: 829–32.
- Knoch, D., Schneider, F., Schunk, D., Hohmann, M., and Fehr, E. (2009). Disrupting the prefrontal cortex diminishes the human ability to build a good reputation. *Proceedings of the National Academy of Sciences* 106: 20895–9.
- Knutson, B., Adams, C. M., Fong, G. W., and Hommer, D. (2001). Anticipation of increasing monetary reward selectively recruits nucleus accumbens. *Journal of Neuroscience* 21: RC159(1–5).
- Knutson, B., Westdorp, A., Kaiser, E., and Hommer, D. (2000). FMRI visualization of brain activity during a monetary incentive delay task. *NeuroImage* 12: 20–27.
- Knutson, B., Wimmer, G. E., Rick, S., Hallon, N. G., Prelec, D., and Loewenstein, G. (2008). Neural antecedents of the endowment effect. *Neuron* 58: 814–22.
- Koenigs, M. and Tranel, D. (2007). Irrational economic decision-making after ventromedial prefrontal damage: evidence from the ultimatum game. *Journal of Neuroscience* 27(4): 951–6.
- Koscik, T. R. and Tranel, D. (2011). The human amygdala is necessary for developing and expressing normal interpersonal trust. *Neuropsychologia* 49: 602–11.
- Kosfeld, M., Heinrichs, M., Zak, P. J., Fischbacher, U., and Fehr, E. (2005). Oxytocin increases trust in humans. *Nature* 435: 673–6.
- Krajchich, I., Adolphs, R., Tranel, D., Denburg, N., and Camerer, C. (2009). Economic games quantify diminished sense of guilt in patients with damage to the prefrontal cortex. *Journal of Neuroscience* 29(7): 2188–92.
- Krueger, F., McCabe, K., Moll, J., et al. (2007). Neural correlates of trust. *Proceedings of the National Academy of Science USA* 104: 20084–9.
- Kuhnen, C. M. and Knutson, B. (2005). The neural basis of financial risk taking. *Neuron* 47: 763–70.
- Kuo, W. J., Sjöström, T., Chen, Y. P., Wang, Y. H., and Huang, C. Y. (2009). Intuition and deliberation: Two systems for strategizing in the brain. *Science* 324: 519–22.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics* 112(2): 443–78.
- Leland, J. W. and Grafman, J. (2005). Experimental tests of the somatic marker hypothesis. *Games and Economic Behavior* 52(2): 386–409.
- Levy, I., Snell, J., Nelson, A. J., Rustichini, A., and Glimcher, P. W. (2010). The neural representation of subjective value under risk and ambiguity. *Journal of Neurophysiology* 103(2): 1036–47.
- Lewis, K. P. and Barton, R. A. (2006). Amygdala size and hypothalamus size predict social play frequency in nonhuman primates: a compar-

- ative analysis using independent contrasts. *Journal of Comparative Psychology* 120: 31–7.
- Lhermitte, F., Pillon, B., and Serdaru, M. (1986). Human anatomy and the frontal lobes. Part I: Imitation and utilization behavior: a neuropsychological study of 75 patients. *Annals Neurology* 19: 326–34.
- Lin, A., Adolphs, R., and Rangel, A. (2012). Social and monetary reward learning engage overlapping neural substrates. *Social Cognitive and Affective Neuroscience* 7(3): 274–81.
- Litt, A., Plassmann, H., Shiv, B., and Rangel, A. (2010). Dissociating valuation from motivation, attentional, and arousal signals during decision-making. *Cerebral Cortex* 21: 95–102.
- Loewenstein, G., Rick, S., and Cohen, J. D. (2008). Neuroeconomics. *Annual Review of Psychology* 59: 647–72.
- Loewenstein, G., Weber, E. U., Hsee, C. K., and Welch, N. (2001). Risk as feelings. *Psychological Bulletin* 127: 267–86.
- Louie, K. and De Martino, B. (2014). The neurobiology of context dependent valuation and choice. Chapter 24 in P. W. Glimcher and E. Fehr (eds.), *Neuroeconomics*. Elsevier Inc.
- Luce R. D. (2001). Reduction invariance and Prelec's weighting functions. *Journal of Mathematical Psychology* 45(1): 167–79.
- McCabe, K., Houser, D., Ryan, L., Smith, V., and Trouard, T. (2001). "A functional imaging study of cooperation in two-person reciprocal exchange". *Proceedings of the National Academy of Sciences of the United States of America* 98(20): 11832–5.
- McClure S. M., Ericson K. M., Laibson D. I., Loewenstein G., and Cohen J. D. (2007). Time discounting for primary rewards. *Journal of Neuroscience* 27: 5796–804.
- McClure, S., Laibson, D., Lowenstein, G., and Cohen, J. (2004). Separate neural systems value immediate and delayed monetary rewards, *Science* 306: 503–7.
- Mason W. A., Capitanio J. P., Machado C. J., Mendoza S. P., and Amaral D. G. (2006). Amygdalectomy and responsiveness to novelty in rhesus monkeys (*Macaca mulatta*): generality and individual consistency of effects. *Emotion* 6: 73–81.
- Mikolajczak, M., Gross, J. J., Lane, A., Corneille, O., de Timary, P., and Luminet, O. (2010). Oxytocin makes people trusting, not gullible. *Psychological Science* 21: 1072–4.
- Mitchell, J. P. (2008). Activity in right temporo-parietal junction is not selective for theory-of-mind. *Cerebral Cortex* 18: 262–71.
- Mitchell, J. P., Banaji, M. R., and Macrae, C. N. (2005). The link between social cognition and self-referential thought in medial prefrontal cortex. *Journal of Cognitive Neuroscience* 17: 1306–15.
- Mitchell, J. P., Macrae, C. N., and Banaji, M. R. (2006). Dissociable medial prefrontal contributions to judgments of similar and dissimilar others. *Neuron* 50: 655–63.
- Mohr, P. N. C., Biele, G., and Heekeren, H. R. (2010). Neural processing of risk. *Journal of Neuroscience* 30: 6613–19.
- Moran, J. M., Macrae, C. N., Heatherton, T. F., Wyland, C. L., and Kelley, W. M. (2006). Neuroanatomical evidence for distinct cognitive and affective components of self. *Journal of Cognitive Neuroscience* 18: 1586–94.
- Morrison, I., Lloyd, D., di Pellegrino, G., and Roberts, N. (2004). Vicarious responses to pain in anterior cingulate cortex: is empathy a multisensory issue? *Cognitive, Affective and Behavioral Neuroscience* 4: 270–8.
- Murray, E. A., Gaffan, E. A., and Flint, R. W., Jr. (1996). Anterior rhinal cortex and amygdala: dissociation of their contributions to memory and food preference in rhesus monkeys. *Behavioral Neuroscience* 110: 30–42.
- O'Doherty, J. P. (2004). Reward representations and reward-related learning in the human brain: insights from neuroimaging. *Current Opinion in Neurobiology* 14: 769–76.
- Padoa-Schioppa, C. (2009). Range-adapting representation of economic value in the orbitofrontal cortex. *Journal of Neuroscience* 29: 14004–14.
- Padoa-Schioppa, C. and Assad, J. A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature* 441(7090): 223–6.
- Paulus, M. P. and Frank, L. R. (2006). Anterior cingulate activity modulates nonlinear decision weight function of uncertain prospects. *Neuroimage* 30: 668–77.
- Penfield, W. and Rasmussen, T. (1950). *The Cerebral Cortex of Man: A Clinical Study of Localization of Function*. Oxford: Macmillan.
- Phillips, M. L., Drevets, W. C., Rauch, S. L., and Lane, R. (2003). Neurobiology of emo-

- tion perception I: the neural basis of normal emotion perception. *Biological Psychiatry* 54: 504–14.
- Plassmann, H., O'Doherty, J., and Rangel, A. (2007). Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *Journal of Neuroscience* 27: 9984–8.
- Plassmann, H., O'Doherty, J., and Rangel, A. (2010). Appetitive and aversive goal values are encoded in the medial orbitofrontal cortex at the time of decision making. *Journal of Neuroscience* 30: 10799–808.
- Plassmann, H., O'Doherty, J., Shiv, B., and Rangel, A. (2008). Marketing actions can modulate neural representations of experienced pleasantness. *PNAS* 105(3): 1050–4.
- Poldrack, R. A. (2006). Can cognitive processes be inferred from neuroimaging data? *Trends in Cognitive Science* 10(2): 59–63.
- Prelec, D. (1998). The probability weighting function. *Econometrica* 66(3): 497–527.
- Preuschoff, K., Bossaerts, P., and Quartz, S. R. (2006). Neural differentiation of expected reward and risk in human subcortical structures. *Neuron* 51: 381–90.
- Rabin, M. (1993). Incorporating fairness into game theory and economics. *American Economic Review* 83: 1281–302.
- Rangel, A., Camerer, C., and Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience* 9: 545–56.
- Rangel, A. and Hare, T. (2010). Neural computations associated with goal-directed choice. *Current Opinion in Neurobiology* 20: 262–70.
- Rilling, J. K., Sanfey, A. G., Aronson, J. A., Nystrom, L. E., and Cohen, J. D. (2004). The neural correlates of theory of mind within interpersonal interactions. *NeuroImage* 22: 1694–703.
- Roiser, J. P., De Martino, B., Tan, G. C. Y., et al. (2009). A genetically mediated bias in decision making driven by failure of amygdala control. *Journal of Neuroscience* 29(18): 5985–991.
- Ruff, C. C. and Huettel, S. C. (2014). Experimental methods in cognitive neuroscience. Chapter 6 in P. W. Glimcher and E. Fehr (eds.), *Neuroeconomics*. Elsevier Inc.
- Rushworth, M. F., Mars, R. B., and Summerfield, C. (2009). General mechanisms for making decisions? *Current Opinion Neurobiology* 19: 75–83.
- Sanfey, A. G., Rilling, J. K., Aronson, J. A., Nystrom, L. E., and Cohen, J. D. (2003). The neural basis of economic decision-making in the ultimatum game. *Science* 300: 1755–8.
- Saxe, R. and Wexler, A. (2005). Making sense of another mind: the role of the right temporoparietal junction. *Neuropsychologia* 43: 1391–9.
- Seymour, B., Singer, T., and Dolan, R. (2007). The neurobiology of punishment. *Nature Reviews Neuroscience* 8: 300–11.
- Sellitto, M., Ciaramelli, E., and Pellegrino, G. D. (2010). Myopic discounting of future rewards after medial orbitofrontal damage in humans. *Journal of Neuroscience* 30: 16429–36.
- Simons, D. J. and Chabris, C. F. (1999). Gorillas in our midst: sustained inattention blindness for dynamic events. *Perception* 28: 1059–74.
- Singer, T., Seymour, B., O'Doherty, J., Kaube, H., Dolan, R. J., Frith, C. D. (2004). Empathy for pain involves the affective but not sensory components of pain. *Science* 303: 1157–62.
- Singer, T., Seymour, B., O'Doherty, J. P., Stephan, K. E., Dolan, R. J., Frith, C. D. (2006). Empathic neural responses are modulated by the perceived fairness of others. *Nature* 439: 466–9.
- Sip, K. E., Roepstorff, A., McGregor, W., and Frith, C. D. (2008). Detecting deception: the scope and limits. *Trends in Cognitive Sciences* 12: 48–53.
- Smith, D. V., Hayden, B. Y., Truong, T. K., Song, A. W., Platt, M. L., and Huettel, S. A. (2010). Distinct value signals in anterior and posterior ventromedial prefrontal cortex. *Journal of Neuroscience* 30: 2490–5.
- Sokol-Hessner, P., Delgado, M. R., Hsu, M., Camerer, C. F., and Phelps, E. A. (2009). Thinking like a trader: cognitive re-appraisal and loss-aversion. *Proceedings of National Academy of Sciences* 106: 5035–40.
- Sokol-Hessner, P., Hsu, M., Delgado, M. R., Camerer, C., and Phelps, E. A. (2009). Getting some perspective: neural correlates of regulation during decision-making. Program No. 102.20, Society for Neuroscience.
- Sowell, E. R., Peterson, B. S., Thompson, P. M., Welcome, S. E., Henkenius, A. L., and Toga, A. W. (2003). Mapping cortical change across the human life span. *A. W. Nature Neuroscience* 6(3): 309–15.
- Spence, S. A., Hunter, M. D., Farrow, T. F., et al. (2004). A cognitive neurobiological account

- of deception: evidence from functional neuroimaging. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 359: 1755–62.
- Stevens J. R., Hallinan E. V., and Hauser M. D. (2005). The ecology and evolution of patience in two NewWorld monkeys. *Biology Letters* 1: 223–26.
- Stott, H. P. (2006). Choosing from cumulative prospect theory's functional menagerie. *Journal of Risk and Uncertainty* 32(2): 101–30.
- Tabibnia, G., Satpute, A. B., and Lieberman, M. D. (2008). The sunny side of fairness: preference for fairness activates reward circuitry (and disregarding unfairness activates self-control circuitry). *Psychological Science* 19(4): 339–47.
- Takahashi, H., Matsui, H., Camerer, C. F., et al. (2010). Dopamine D1 receptors and nonlinear probability weighting in risky choice. *Journal of Neuroscience* 30: 16567–72.
- Talmi, D., Hurlmann, R., Patin, A., and Dolan, R. J. (2010). Framing effect following bilateral amygdala lesion. *Neuropsychologia* 48: 1823–7.
- Thaler, R. H. and Shefrin, H. M. (1981). An economic theory of self-control. *Journal of Political Economy* 89(2): 392–406.
- Tobler, P. N., O'Doherty, J. P., Dolan, R. J., and Schultz, W. (2007). Reward value coding distinct from risk attitude-related uncertainty coding in human reward systems. *Journal of Neurophysiology* 97(2): 1621–32.
- Tom, S. M., Fox, C. R., Trepel, C., and Poldrack, R. A. (2007). The neural basis of loss aversion in decision-making under risk. *Science* 315: 515–18.
- Tricomi, E., Rangel, A., Camerer, C. F., and O'Doherty, J. P. (2010). Neural evidence for inequality-averse social preferences. *Nature* 463: 1089–91.
- Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4): 297–323.
- van den Bos, W., van Dijk, E., Westenberg, M., Rombouts, S. A. R. B., and Crone, E. A. (2009). What motivates repayment? Neural correlates of reciprocity in the Trust Game. *Social Cognitive & Affective Neuroscience* 4: 294–304.
- Van Overwalle, F. (2009). Social cognition and the brain: a meta-analysis. *Human Brain Mapping* 30(3): 829–58.
- Vanberg, C. (2008). Why do people keep their promises? An experimental test of two explanations. *Econometrica* 76(6): 1476–80.
- Wallis, J. (2007). Orbitofrontal cortex and its contributions to decision-making. *Annual Reviews in Neuroscience* 30: 31–56.
- Wang, J. T., Spezio, M., and Camerer, C. F. (2010). Pinocchio's pupil: using eyetracking and pupil dilation to understand truth telling and deception in sender-receiver games. *American Economic Review* 100(3): 984–1007.
- Wassermann, E. M., Epstein, C. M., Ziemann, U., Walsh, V., Paus, T., Lisanby, S. H. (eds.). (2008). *The Oxford Handbook of Transcranial Stimulation*. Oxford: Oxford University Press.
- Webb, R. (2015). The dynamics of stochastic choice. Unpublished manuscript.
- Wicker, B. et al. (2003). Both of us disgusted in my insula: the common neural basis of seeing and feeling disgust. *Neuron* 40: 655–64.
- Woodford, M. (2014). Stochastic choice: an optimizing neuroeconomic model. *American Economic Review* 104(5): 495–500.
- Wout, M. van't, Kahn, R. S., Sanfey, A. G., and Aleman, A. (2005). Repetitive transcranial magnetic stimulation over the right dorsolateral prefrontal cortex affects strategic decision-making. *NeuroReport* 16: 1849–52.
- Yacubian, J., Glascher, J., Schroeder, K., Sommer, T., Braus, D. F., and Buchel, C. (2006). Dissociable systems for gain- and loss-related value predictions and errors of prediction in the human brain. *Journal of Neuroscience* 26: 9530–7.
- Zak, P. J. (2004). Neuroeconomics. *Philosophical Transactions of the Royal Society B* 359: 1737–48.
- Zaki, J. and Mitchell, J. (2011). Equitable decision making is associated with neural markers of subjective value. *Proceedings of the National Academy of Sciences USA* 108: 19761–6.
- Zink, C. F., Tong, Y., Chen, Q., Bassett, D. S., Stein, J. L., and Meyer-Lindenberg, A. (2008). Know your place: neural processing of social hierarchy in humans. *Neuron* 58: 273–83.

APPENDIX ON GAME THEORY

A.1 Introduction

This appendix introduces the notation and some of the basics of game theory in an informal manner, with particular emphasis on some of the games that are considered in this book. It is not a substitute for a course in game theory, nor intended to teach formal game theory as such. The idea is to provide interdisciplinary readers and undergraduate economics students with a practical starter-guide that enables them to read this book. The material is provided in a relatively condensed form, without proofs, but with recourse to simple examples. Readers who are interested in a formal treatment of game theory can consult several excellent sources.¹

A.2 Normal form games

The set of players is $N = \{1, 2, \dots, n\}$. Each player $i \in N$ can play n_i *pure strategies*, thus, the pure strategy set of player i , $S_i = \{s_{i1}, s_{i2}, \dots, s_{ini}\}$. To minimize notation, we shall denote a generic element of the set S_i by s_i . Elements of S_i can be abstract entities, but for most of the applications we have $s_i \in \mathbb{R}$.² A *pure strategy profile* is given by $s = (s_1, s_2, \dots, s_n)$; this is an element of the set $S = \times_i S_i$. It is sometimes convenient to write the pure strategy profile as $s = (s_i, s_{-i})$ where s_{-i} is the pure strategy profile of all players other than player i .

The utility function of the i th player is given by the mapping $u_i : S \rightarrow \mathbb{R}$. We summarize the list of all payoff functions by $u = (u_1, u_2, \dots, u_n)$.³ An n -player *normal form* or *strategic form* game is denoted by $G(N, S, u)$. *Actions* and *strategies* coincide in one-shot simultaneous move games, as in the static prisoner's dilemma game. However, in sequential move games (see below) these two entities differ.

It is possible that players might wish to randomize between their pure strategy, much like a tennis player randomizing his service to the right and to the left of the opponent or a football player kicking a penalty to the right or to the left of a goalkeeper. This motivates the concept of a *mixed strategy*.

Player i 's *mixed strategy*, σ_i , is simply a probability distribution over the pure strategy set, S_i . We denote by $\sigma_i(s_i)$ the probability with which player i plays the pure strategy $s_i \in S_i$. Sometimes

¹ There is a range of excellent books on the subject. For instance, at a basic level, Gibbons (1992), Dixit and Skeath (2004), Osborne (2004), and Watson (2013) provide lucid treatments that are rich in applications. Fudenberg and Tirole (1991a) is a more advanced, but very concise and authoritative, book that would be suitable for those who prefer a more rigorous approach. A small sample of other advanced books that may be consulted is Maschler et al. (2013), Osborne and Rubinstein (1994), Myerson (1991), and Kreps (1990).

² When this may create potential ambiguity we denote the generic element of S_i by s_{ij} , $j = 1, 2, \dots, n_i$.

³ At this point, we impose no restrictions on the compactness of the strategy space or the continuity or shape (e.g., quasiconcavity) of the payoff functions. Such restrictions will, however, be needed when we consider sufficient conditions for the existence of an equilibrium below.

we may need to use the slightly more cumbersome notation $\sigma_i(s_{ij})$ for the probability with which player i plays the pure strategy $s_{ij} \in S_i$. Any pure strategy that is played with a non-zero probability under σ_i is said to belong to its *support*. In particular, a pure strategy $s_i \in S_i$ for player i , simply corresponds to a degenerate mixed strategy, $\sigma_i(s_i) = 1$. The set of all mixed strategies of player i is denoted by Σ_i . Formally,

$$\Sigma_i = \left\{ \sigma_i = (\sigma_i(s_{i1}), \sigma_i(s_{i2}), \dots, \sigma_i(s_{in_i})) : \sigma_i(s_{ij}) \geq 0 \forall j \text{ and } \sum_{j=1}^{n_i} \sigma_i(s_{ij}) = 1 \right\}. \quad (\text{A.1})$$

The mixed strategy profile of the game $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_n)$ is an element of the set $\Sigma = \Sigma_1 \times \Sigma_2 \times \dots \times \Sigma_n$. It will sometimes be convenient to express the mixed strategy profile by $\sigma = (\sigma_i, \sigma_{-i})$ where σ_{-i} is the mixed strategy profile of all players other than player i . In actual practice, one only observes a particular pure strategy played by a player, so the observer must infer the underlying mixed strategy of the player. Whether players can do so is an empirical question. Part 4 considers these issues in some detail.

Suppose that each player's randomization over the pure strategies is independent of that of the other player. Hence, the joint probability of any strategy profile $s = (s_1, s_2, \dots, s_n) \in S$ is $\prod_{i=1}^n \sigma_i(s_i)$. Thus, assuming the expected utility property, the expected payoff of player i from the mixed strategy profile $\sigma = (\sigma_i, \sigma_{-i})$ is given by

$$\pi_i(\sigma_i, \sigma_{-i}) = \sum_{s \in S} \left[\prod_{i=1}^n \sigma_i(s_i) \right] u_i(s_i, s_{-i}). \quad (\text{A.2})$$

In (A.2), the own strategy, σ_i , or the strategy of others, σ_{-i} , can be possibly degenerate. This allows us to compute $\pi_i(s_i, \sigma_{-i})$, $\pi_i(\sigma_i, s_{-i})$ and $\pi_i(s_i, s_{-i}) \equiv u_i(s_i, s_{-i})$.

Definition A.1 (*Best response*): Suppose that players satisfy the expected utility property.

(a) Denote the mixed strategy profile of all players other than player i by σ_{-i} . Then σ_i^* is said to be a *best response* (sometimes called *best reply*) of player i if

$$\pi_i(\sigma_i^*, \sigma_{-i}) \geq \pi_i(s_i, \sigma_{-i}) \text{ for all } s_i \in S_i. \quad (\text{A.3})$$

Denote the set of all best responses to σ_{-i} by $BR_i(\sigma_{-i})$.

(b) (*Nash equilibrium*) A strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*, \dots, \sigma_n^*)$ is a *Nash equilibrium* if for all $i \in N$ we have $\sigma_i^* \in BR_i(\sigma_{-i}^*)$.

In Definition A.1, the best response and the Nash equilibrium need not be unique. Also note that in (A.3) we need to check only if there are any profitable “pure strategy deviations” of player i . Due to the linearity of the payoffs in (A.2), if (A.3) holds then so does $\pi_i(\sigma_i^*, \sigma_{-i}) \geq \pi_i(\sigma_i, \sigma_{-i})$ for all $\sigma_i \in \Sigma_i$ [The reader should check this.]

All pure strategies in the support of a mixed strategy must give identical expected payoffs. Suppose that two pure strategies in the support of a mixed strategy give different expected payoffs. Then the individual can improve the payoff by eliminating the pure strategy with the lower expected payoff. Thus, it follows that a *strict best response* in which the inequality in (A.3) is strict can only be in pure strategies. We now give the strict version of Definition A.1.

Definition A.2 (*Strict best response*): If $\pi_i(s_i^*, \sigma_{-i}) > \pi_i(s_i, \sigma_{-i})$ for all $s_i \neq s_i^*$ then the *best response* is *unique*, i.e., $BR_i(\sigma_{-i}) = \{s_i^*\}$.

Remark A.1 (Strict Nash equilibrium): The strategy profile $s^* = (s_1^*, s_2^*, \dots, s_n^*)$ is a strict Nash equilibrium if $BR_i(s_{-i}^*) = \{s_i^*\}$ for all $i \in N$.

There is no guarantee that a strict Nash equilibrium will exist (see, for instance, the matching pennies game below).

An alternative equilibrium concept can be given in terms of dominated strategies.

Definition A.3

(a) (Strict dominance): A pure strategy s_i is “strictly dominated” for player i if there exists some other mixed strategy $\sigma_i \in \Sigma_i$ such that

$$\pi_i(\sigma_i, s_{-i}) > \pi_i(s_i, s_{-i}) \text{ for all } s_{-i} \in S_{-i}. \quad (\text{A.4})$$

(b) (Weak dominance): If the inequality in (A.4) is weak but for at least one $s_{-i} \in S_{-i}$ it is strict then s_i is a “weakly dominated” strategy for player i .

Note that due to the linearity of the payoffs in (A.2), if (A.4) holds for all $s_{-i} \in S_{-i}$ then it also does hold for all $\sigma_{-i} \in \Sigma_{-i}$. The reason that we check for dominance of a pure strategy against a mixed strategy in (A.4) is given in the next example.

Example A.1 A pure strategy which is not dominated by any pure strategy can nevertheless be strictly dominated by a mixed strategy. As an example, consider the following payoff matrix.

	L	R
U	8, .	0, .
M	0, .	8, .
D	3, .	3, .

The row player (player 1) has three strategies, U, M, D and the column player (player 2) has two strategies, L, R. The payoffs of player 1 are shown in the payoff matrix while the payoffs of the column player are suppressed because they are not relevant to the argument. For player 1, none of the pure strategies is dominated by any other pure strategy. Now consider a mixed strategy for the row player in which U and M are each played with probability 1/2. The payoff from this mixed strategy, given any strategy of the opponent is $4 > 3$, hence, it dominates the pure strategy D.

It is obvious that any mixed strategy whose support contains a dominated pure strategy, is itself dominated by another possibly mixed strategy—one that eschews the play of the dominated strategy. The next example makes a more subtle point that even if a mixed strategy does not contain any dominated pure strategy in its support, it could still be dominated by another mixed strategy.

Example A.2 Suppose that we alter the normal form of the game in Example A.1 to the following (continuing to suppress the payoffs of player 2).

	L	R
U	8, .	−6, .
M	−6, .	8, .
D	2, .	2, .

None of the pure strategies of the row player is dominated. Consider a mixed strategy in which the row player plays U, M with equal probabilities. Irrespective of the strategy of the column player, the expected payoff from this mixed strategy is 1 which is dominated by the pure strategy D .

Suppose that we were to iteratively delete strictly dominated strategies of players. The simplest case occurs when one round of deletion of strictly dominated strategies leads to a unique strategy profile $s^* = (s_1^*, s_2^*, \dots, s_n^*)$. It follows that for all i , $\pi_i(s_i^*, s_{-i}^*) > \pi_i(s_i, s_{-i}^*)$ for all $s_i \neq s_i^*$ and $s_i \in S_i$. Thus, s^* is the unique Nash equilibrium of the game. One can analogously show that if one round of deletion of weakly dominated strategies leads to a unique surviving strategy profile, it too must be a Nash equilibrium. These results can be generalized significantly as shown in the next proposition; the reader can consult any standard game theory text for the proof.

Proposition A.1 *Suppose that the successive iterated deletion of strictly dominated strategies leads to a unique strategy profile $s^* = (s_1^*, s_2^*, \dots, s_n^*)$. Then s^* is the unique Nash equilibrium of the game. The result is invariant to the order of deletion of the strictly dominated strategies.⁴*

Proposition A.1 demonstrates the importance of the concept of iterated deletion of dominated strategies by showing its relationship with a Nash equilibrium. A few comments are in order. First, as one engages in more rounds of iterated deletion of dominated strategies, the degree of rationality that we impose on players grows too. Indeed, we need *common knowledge of rationality*⁵ and of the *structure of the game* for these results to hold. Second, not all games can be solved by iterated deletion of dominated strategies, yet they may have a Nash equilibrium, hence, the latter concept is less restrictive. Third, the concept of iterated deletion of dominated strategies bears an important relation to the concept of *rationalizability* introduced by Bernheim (1984) and Pearce (1984). The basic idea of rationalizability is to narrow down the set of strategies that rational players will play as best responses, given their beliefs. The beliefs themselves are not arbitrary. The assumption of common knowledge of rationality ensures that rational players believe that other rational players would also play a best response to some reasonable beliefs that they hold about them and so on; rationalizable strategies are those that survive this infinite regress. In two-player games, the set of rationalizable strategies and those that survive iterated strict dominance, coincide. However, for more general games, rationalizable strategies are a subset of the strategies that survive iterated deletion of strictly dominated strategies.

It is often said that a Nash equilibrium gives a *consistent prediction* of how the game will be played. The reason is that if each player can correctly guess the strategy of the others in a Nash equilibrium, then none of the players has any incentive to deviate from their own strategies in a Nash equilibrium. Whether players actually do so is an empirical question that is examined in Parts 4 and 5. The situation gets murkier if there are multiple Nash equilibria. How do players coordinate their beliefs on a particular Nash equilibrium in this case? Evidence on this issue is examined in Part 4.

⁴ This conclusion does not hold for the iterated deletion of weakly dominated strategies (IDWDS) where the order of deletion may be important. Further, not all Nash equilibria may survive IDWDS (the reader may wish to construct simple examples to show this).

⁵ In a two-player game with players A and B, there is common knowledge of an event X if the following infinite regress holds. Player A knows X. Player A knows that B knows X. Player A knows that B knows that A knows that B knows X and so on. A similar statement holds for player B. By contrast, *mutual knowledge of rationality* means that each player knows that he or she is rational and that other players are rational.

Although this is ultimately an empirical question that we address in the book, why should a purely deductive equilibrium concept such as a Nash equilibrium hold in static, simultaneous move games, or in novel situations with new players? In a strict Nash equilibrium with N players, say, $s^* = (s_1^*, s_2^*, \dots, s_n^*)$, we imagine player i asking questions such as these: “My opponents are playing their part of the Nash equilibrium, s_{-i}^* , can I do any better by deviating to one of my other pure strategies $s_i \in S_i$?” But why should someone be thinking along these lines when they do not know what strategies their opponents are choosing? Perhaps there is an obvious way to play the game where such a deductive approach may be useful. But what if there is no obvious way to play the game? There are two observations in this respect.

1. The concept of a Nash equilibrium (Definition A.1) does not place any restrictions on the beliefs of players about what others will do. Restrictions on beliefs that lead to a Nash equilibrium outcome are the subject matter of the *epistemic foundations of a Nash equilibrium*. The epistemic conditions for a Nash equilibrium turn out to be extremely demanding. Aumann and Brandenburger (1995) proved the following result. Suppose that we have mutual knowledge of the payoffs, mutual knowledge of rationality of players, common knowledge of the beliefs (or conjectures) each player has about the others, and common priors. Then the common conjectures about each player j by the other players agree and constitute a Nash equilibrium of the game. Polak (1999) showed that if mutual knowledge of the payoffs is strengthened to common knowledge of payoffs then the Aumann–Brandenburger conditions imply common knowledge of rationality. Recent research offers an unsettled view of the epistemic foundations; see for instance, Gintis (2009). For our purposes, it is important to note that if a Nash equilibrium is violated by the empirical evidence, then the epistemic conditions are also empirically refuted.
2. A commonly used defence of Nash equilibrium is that it is the outcome of a learning process that takes place over a “sufficiently long period of time”. In these learning processes, players could be responding to the average play of the opponent in the past, or the play of the opponent in the previous period (as in a *Cournot adjustment process*), or the past level of some aggregate variable. At a purely theoretical level, (i) convergence of play to a Nash equilibrium is not always guaranteed using these learning processes, and (ii) when convergence takes place, there is no guarantee that the underlying learning process is plausible; see, for instance, Fudenberg and Tirole (1991a, pp. 25–9).

Even when convergence does take place, a valid question to ask is: Does it take place in a “sensible amount of time” that is representative of actual real-world interactions?⁶ On an empirical level, when we use field data, can we be sure that the data is not contaminated by repeated game effects in what is essentially a static game? By contrast, such repeated game effects can be controlled within an experimental setting by ensuring that subjects are not repeatedly matched and do not communicate. These are some of the issues that must be addressed by learning-based justifications of a Nash equilibrium.

Rationalizability, on the other hand, does not make predictions based on any equilibrium concept. However, if the set of possible strategy profiles that are deemed to be reasonable under rationalizable strategies is too large then the predictive power of theory is weak.

⁶ My long time co-author, Ali al-Nowaihi, attended a conference at a leading university where the speaker claimed that, in his model, he was pleased to report that convergence to a Nash equilibrium occurs in 10^{83} periods. Ali asked him if he knew the number of particles in the observable universe (an estimate is 10^{80})? Clearly such justifications of a Nash equilibrium are not convincing.

Does a Nash equilibrium exist? In many examples that we consider, the pure strategy sets of players, S_i , are finite (this is not always the case, e.g., in Cournot games where output choices are continuous⁷). For such games, the existence of a Nash equilibrium, possibly in mixed strategies, is always guaranteed. The proofs require standard fixed point arguments so we omit them; see for instance Fudenberg and Tirole (1991a).

Theorem A.1 (Nash, 1950): Consider the normal form game $G(N, S, u)$ where N is a finite set. Suppose that, for all i , the strategy sets S_i are finite. Then there exists at least one Nash equilibrium in mixed strategies.

In some commonly used games, such as under Cournot competition, the strategy spaces of each player are chosen from some continuous interval, hence, S_i is not finite. Thus, Theorem A.1 is not applicable. Notice that Theorem A.1 only provides sufficient conditions, so while a Nash equilibrium might still exist when the finiteness of strategy spaces is violated, it is not guaranteed. Another theorem is often helpful in this case for finding a *pure strategy* Nash equilibrium.

Theorem A.2 Consider the normal form game $G(N, S, u)$ where N is a finite set. Suppose that for all i , the strategy sets S_i are non-empty, convex, and compact subsets of some Euclidean space. If for all i , π_i is continuous in $s = (s_1, s_2, \dots, s_n)$ and quasi-concave in own strategy, s_i , then a pure strategy Nash equilibrium exists.⁸

Example A.3 (Mixed strategy Nash equilibrium in the battle of the sexes game).

	O	T
O	2, 1	0, 0
T	0, 0	1, 2

A couple can either go to the Opera (strategy O) or the Theater (strategy T). They only like going out together so the payoff from going separately to the two events is zero for both. The husband (row player) prefers the Opera while the wife (column player) prefers the Theater. There are two pure strategy Nash equilibria, (O, O) and (T, T). There is also one mixed strategy equilibrium, which can be found as follows. Denote the mixed strategies of the husband and wife respectively by σ_H and σ_W . The husband's mixed strategy σ_H involves the play of O, T with respective probabilities $p, 1 - p$ while under σ_W the strategies O, T are played with respective probabilities $q, 1 - q$. We wish to find the best response of each player, given the mixed strategy of the other. The husband's payoffs from the two pure strategies are:

$$\pi_H(O, \sigma_W) = 2q; \quad \pi_H(T, \sigma_W) = 1 - q \quad (\text{A.5})$$

Recall that if two pure strategies belong to the support of a mixed strategy, they should generate identical expected payoffs (otherwise the one that generates a lower payoff can be eliminated to improve payoffs). The two payoffs in (A.5) are equal when $2q = 1 - q \Leftrightarrow q = 1/3$. For the husband, when $q > 1/3$, the strategy Opera dominates and when $q < 1/3$, the strategy

⁷ In a Cournot game, two firms simultaneously choose their quantity levels. The sum of these quantities determines the market supply and the market demand curve is downward sloping. The two firms may or may not have identical costs.

⁸ For proofs see Debreu (1952) and Glicksberg (1952). A proof is also sketched in Fudenberg and Tirole (1991a).




Figure A.1 Best response functions for the battle of sexes game.

Theater dominates. However, when $q = 1/3$ then the husband is indifferent between the two. Denote the best response of the husband, given the wife's strategy by $\sigma_H^* \in BR^H(\sigma_W)$. Since $p + 1 - p = 1$, let us summarize σ_H as a choice of p alone (analogously for the wife, by q alone). The optimal value of $p = p^*(q)$ for any q is, thus, summarized by:

$$p^*(q) = \begin{cases} 1 & \text{if } q > 1/3 \\ [0, 1] & \text{if } q = 1/3. \\ 0 & \text{if } q < 1/3 \end{cases} \quad (\text{A.6})$$

Proceeding analogously, one can derive the best reply of the wife:

$$q^*(p) = \begin{cases} 1 & \text{if } p > 2/3 \\ [0, 1] & \text{if } p = 2/3. \\ 0 & \text{if } p < 2/3 \end{cases} \quad (\text{A.7})$$

The two best reply functions are shown in Figure A.1. On the vertical axis we have $p^*(q)$, the probability with which the husband plays O, as a function of the probability, q , with which the wife plays O. $p^*(q)$, the husband's best response, as described in (A.6), can then be drawn as the bold continuous step function. On the horizontal axis we have $q^*(p)$, the optimal probability with which the wife plays O as a function of p , the husband's probability of playing O. The best response of the wife, $q^*(p)$, is shown as the bold, dotted step function. The two best response functions intersect at the two pure strategy Nash equilibria, (0, 0) and (1, 1). They also intersect at the unique interior point (2/3, 1/3), which is the (unique) Nash equilibrium in mixed strategies.

A.3 A basic taxonomy of some normal form games

Consider the following game with two players and two actions for each player. The row player chooses from the actions U, D and the column player chooses from the actions L, R . The payoffs corresponding to each strategy profile are shown in the following payoff matrix.

	L	R
U	a_{11}, b_{11}	a_{12}, b_{12}
D	a_{21}, b_{21}	a_{22}, b_{22}

1. If (U, L) and (D, R) are Nash equilibria, then we have a *coordination game*. In such games, (U, L) is said to *payoff dominate* (D, R) if $a_{11} \geq a_{22}$ and $b_{11} \geq b_{22}$ with at least one of these inequalities being strict. This enables us to Pareto rank the two equilibria: (U, L) is Pareto superior to (D, R) . Under these conditions, if we observe that (U, L) is not played, then we have a *coordination failure*. These definitions go through for any finite game.
2. Consider a coordination game in which (U, L) and (D, R) are strict Nash equilibria (i.e., $a_{11} > a_{21}$, $b_{11} > b_{12}$, $a_{22} > a_{12}$, $b_{22} > b_{21}$), but in addition, we have $a_{12} < a_{21}$ and $b_{21} < b_{12}$ then we have a *stag-hunt game*. For instance,

	L	R
U	2, 2	0, 1
D	1, 0	1, 1

Payoff maximizing players want the equilibrium (U, L) to be played. However, there might be uncertainty about the strategy of the opponent. For instance, if the row player plays U and the column player plays R , the row player gets the lowest possible payoff. Hence, this game is sometimes also known as an *assurance game*; each player would like to be assured that the other will play their part of the *payoff dominant equilibrium*, (U, L) . In the absence of such assurance, players might wish to play a *secure option*. The *maximin outcome* or the *secure option* (see formal definition below) arises when the outcome is (D, R) . In this case, each player is guaranteed a minimum payoff of 1, irrespective of the action of the opponent.

3. Suppose that (U, L) and (D, R) are Nash equilibria. If $a_{11} > a_{22}$ and $b_{11} < b_{22}$ then this is known as a *battle of sexes game* (BoS) that we used to illustrate the computation of mixed strategy equilibrium, above. The row player would like the Nash equilibrium (U, L) to be played while the column player prefers the Nash equilibrium (D, R) to be played.

A.4 Two-person zero-sum games

A fair deal of the testing of the *mixed strategy Nash equilibrium* (MSE) has been carried out in finite, two-player zero-sum games. In such games, for any pair of strategies of the two players, (s_1, s_2) , the sum of the payoffs of the two players is zero, i.e., $\pi_1(s_1, s_2) + \pi_2(s_1, s_2) = 0$. The *minimax outcome* in such games can be explained as follows. Suppose that the mixed strategies of the two players are σ_1, σ_2 and the payoff functions of the players are $\pi_1(\sigma_1, \sigma_2)$, $\pi_2(\sigma_1, \sigma_2)$. The *minimax value* to player 1 is then defined as

$$v_1 = \min_{\sigma_2} \left(\max_{\sigma_1} \pi_1(\sigma_1, \sigma_2) \right).$$

Essentially first fix the strategy of player 2 and maximize the payoff of player 1 with a choice of his strategy, σ_1 . Then minimize the result with a choice of the strategy of player 2, σ_2 . Clearly, player 1 can never be held to a payoff lower than this.

The *maximin value* to player 1 can be similarly calculated in two steps as follows.

$$V_1 = \max_{\sigma_1} \left(\min_{\sigma_2} \pi_1(\sigma_1, \sigma_2) \right).$$

von Neumann (1928) showed that for zero-sum games, $v_1 = V_1$. Analogously for player 2, $v_2 = V_2$. Let the mixed strategy Nash equilibrium for the game be denoted by (σ_1^*, σ_2^*) . For finite, zero-sum games, the minimax, maximin, and MSE outcomes coincide and $v_1 = V_1 = \pi_1(\sigma_1^*, \sigma_2^*)$ and $v_2 = V_2 = \pi_2(\sigma_1^*, \sigma_2^*)$. Thus, support for or rejection of minimax in these games is respectively support for or rejection of MSE.

A.5 Selection criteria among Nash equilibria: payoff dominance and risk dominance

In the presence of multiple Nash equilibria, one needs to invoke extra refinements to predict which equilibrium might be chosen. An important refinement of Nash equilibrium, based on purely deductive reasoning, is Harsanyi and Selten's *payoff dominance* and *risk dominance criteria*. Consider the two-player coordination game given by:

1\2	L	R
L	a_{11}, b_{11}	a_{12}, b_{12}
R	a_{21}, b_{21}	a_{22}, b_{22}

We assume that (L, L) and (R, R) are strict Nash equilibria (i.e., $a_{11} > a_{21}; b_{11} > b_{12}; a_{22} > a_{12}; b_{22} > b_{21}$). Further, both players get a higher payoff under one of the two Nash equilibria so that one equilibrium *payoff dominates* the other. For instance, (R, R) *payoff dominates* (L, L) if $a_{22} \geq a_{11}$ and $b_{22} \geq b_{11}$ with at least one of these two inequalities being strict. Can one make an argument in favor of (L, L) , on the grounds that it is less risky in some sense? The risk dominance criteria explores this question.

Suppose that player 1 (the row player) believes that player 2 (the column player) plays R with probability q and L with probability $1 - q$. Player 2 believes that player 1 plays R with probability p and L with probability $1 - p$. If $q = 1$ then player 1 should play R because (R, R) is a Nash equilibrium. For $q = 0$, player 1 should play L because (L, L) is a Nash equilibrium. This motivates the next definition.

Definition A.4 Let $\underline{q}$ be the minimum probability of play of R by player 2 such that it is optimal for player 1 to play R . Analogously let $\underline{p}$ be the minimum probability of play of R by player 1 such that it is optimal for player 2 to play R .

Definition A.5 (*Risk factor and risk dominance*): The risk factor of the equilibrium (R, R) is defined as $\min \{ \underline{q}, \underline{p} \}$. Similarly the risk factor of the equilibrium (L, L) is defined as $\min \{ 1 - \underline{q}, 1 - \underline{p} \}$. A risk dominant equilibrium is one whose risk factor is lowest.

Using Definition A.5, (L, L) is the risk dominant equilibrium if

$$\min \{ \underline{q}, \underline{p} \} \geq \min \{ 1 - \underline{q}, 1 - \underline{p} \}. \quad (\text{A.8})$$

Consider the equilibrium (R, R) first. It is optimal for the row player to play R if

$$a_{21}(1 - q) + a_{22}q \geq a_{11}(1 - q) + a_{12}q,$$

hence,

$$\underline{q} = \frac{a_{21} - a_{11}}{a_{21} + a_{12} - a_{11} - a_{22}}. \quad (\text{A.9})$$

Analogously, we can find that

$$\underline{p} = \frac{b_{12} - b_{11}}{b_{12} + b_{21} - b_{11} - b_{22}}. \quad (\text{A.10})$$

Using (A.8), (A.9), (A.10), (L, L) is the risk dominant equilibrium if

$$(a_{11} - a_{21})(b_{11} - b_{12}) \geq (a_{22} - a_{12})(b_{22} - b_{21}). \quad (\text{A.11})$$

Expression (A.11) gives a quick way of finding the risk dominant equilibrium. Each player is less likely to deviate if the equilibrium payoff is relatively greater as compared to deviating to an alternative strategy. The risk dominant equilibrium maximizes the product of the gain from unilateral deviation. There is a particularly simple criteria to check for risk dominance in symmetric 2×2 games such as the following.

	L	R
L	a, a	c, d
R	d, c	b, b

Using (A.11) (L, L) is risk dominant relative to (R, R) if

$$(a - d)(a - d) \geq (b - c)(b - c) \Leftrightarrow a - d \geq b - c. \quad (\text{A.12})$$

Consider the following symmetric coordination game from Fudenberg and Levine (1998).

	L	R
L	$2, 2$	$-a, 0$
R	$0, -a$	$1, 1$

In this game, there are two pure strategy Nash equilibria, (L, L) and (R, R) . The reader may wish to find the conditions under which (L, L) is a risk dominant equilibrium relative to (R, R) .




Figure A.2 Two examples of extensive form games.

A.6 Extensive form games

In an *extensive form game*, players move in sequence, having observed the history of the game so far. Consider two examples of extensive form games in Figure A.2 in panels A and B used in the experimental tests of, respectively, Beard and Beil (1994) and Schotter et al. (1994). The empirical evidence for these games is considered in Part 4.

1. The game in Figure A.2-A is a 2-player *dynamic game of perfect information*. Player 1 moves first, choosing among two possible strategies L, R. If player 1 plays R, then player 2 gets an opportunity to choose one among two possible strategies, l, r. The payoffs are shown at the bottom of the game. When making a move, each player is perfectly certain which decision node he is at.
2. The game in Figure A.2-B is a *dynamic game of imperfect information*. Player 1 can play T and end the game or he can play L and hand over the choice of moves to player 2. Player 2 can then play T and end the game or play L, in which case both players play a simultaneous move battle of sexes game, BoS (see Section A.3 above), in which each player chooses among one of two possible strategies, M, B. The dotted line indicates that when player 2 chooses his action, M or B, he does not, in the simultaneous move BoS game, know the action chosen by player 1; this is the sense in which information is imperfect.

Consider a slightly more formal structure of extensive form games. There is a finite number of players, $i = 1, 2, \dots, n$ who typically move in sequence although at some stage of the game, a subset of the players may move simultaneously. There is a finite set of (decision) nodes $T = \{t_1, t_2, \dots, t_K\}$. For instance, in the extensive form in Figure A.2-B, there are five nodes at which players decide on their choice of actions. A binary relation $<$ is defined on T that takes the form of a *strict precedence relation*. Thus, $t_i < t_j$ means that the node t_i precedes the node t_j . Since this relation is strict, it is asymmetric so we cannot have $t_i < t_i$ but the relation is transitive, i.e., $t_i < t_j$ and $t_j < t_k$ imply $t_i < t_k$. It may help to think of several generations of a family as the motivation for this structure.

It will be useful, for exposition purposes, to classify the nodes into the following distinct categories. By $P(t)$ we denote the set of *predecessors* of node t and by $S(t)$ the set of *successors* of node t . Formally,

$$P(t) = \{t_j \in T : t_j \prec t\} \text{ and } S(t) = \{t_j \in T : t \prec t_j\}.$$

There is an *initial node* in each panel of Figure A.2 corresponding to the move by player 1 which is not preceded by any node. We call this the set of initial nodes $I = \{t \in T : P(t) = \phi\}$, where ϕ denotes the empty set. In this book we typically assume that the set I is a singleton set. Also in each panel of Figure A.2, there are a set of *final nodes* immediately followed by the payoffs but no successor nodes. We call these the *terminal nodes* $Z = \{t \in T : S(t) = \phi\}$. Other than at the terminal nodes, some player makes a decision at every other node. Thus, we call $D = T \setminus Z$ the set of decision nodes.

Drawing the analogy with successive generations of humans, any human has a unique parent. Analogously, any node other than the set of initial nodes, $t \in T \setminus I$, has an immediate unique predecessor node, $\rho(t)$, so $\rho(t) \prec t$. Thus, if $t_k \prec t$ and $t_k \neq \rho(t)$ then it must be that $t_k \prec \rho(t)$ [The reader should try to show this.]. This is an important implication because it means that starting from any node $t \in T$, by tracing the immediate predecessor of each node $(t, \rho(t), \rho(\rho(t)), \dots)$ we will eventually reach one of the initial nodes. Analogously, starting from an initial node, there is a unique path to any node t . In particular, there can be no cycles, and so we get a tree structure or an *arborescence*.

At any decision node $d \in D$, a subset of the players, or a special player, *nature*, listed as player $n + 1$, can make a move. We say that player $i(d) \subseteq \{1, 2, \dots, n, n + 1\}$ moves at the decision node $d \in D$. The set of actions available at a decision node $d \in D$ is denoted by $A(d)$.

A critical component in the description of extensive form games is the specification of *information sets*. All decision nodes $d \in D$ are partitioned into information sets. Denote by H the set of all possible information sets. A typical element of H is denoted by h . Two nodes, d_1, d_2 belong to the same information set h if the following three conditions are satisfied. (i) Each of the nodes does not belong to the set of predecessors of the other ($d_1 \notin P(d_2)$ and $d_2 \notin P(d_1)$). (ii) The player who moves at the two nodes is identical ($i(d_1) = i(d_2)$) and this player does not know with certainty, which of the two nodes he or she is at. (iii) The set of actions at the two nodes are identical ($A(d_1) = A(d_2)$).

Extensive form games in which each information set $h \in H$ is a singleton are known as *dynamic games of perfect information*. An example is the game in Figure A.2-A as are many other games examined in this book such as the *centipede games* which provide important empirical insights on refinements of Nash equilibrium. By contrast, when there is at least one $h \in H$ that is not a singleton then we have a *dynamic game of imperfect information*, as in Figure A.2-B. In Figure A.2-B when players play the BoS game, the two nodes at which player 2 may move and play M or B are in the same information set (check that this satisfies the conditions of an information set).

The final component, the payoff function, u_i , gives the payoff of each of the players $i \in N$ at each of the terminal nodes in T .

Dynamic games of perfect information can be solved by *backward induction*, which can be explained as follows. Assume that there is common knowledge of the game, of the payoff structure, and of the rationality of players. Begin at the bottom of the tree with all nodes whose successors are terminal nodes. The player with the move at any of these nodes chooses an action that maximizes his payoff. Fold back the game tree, replacing each of these nodes with the payoffs achieved in the previous calculation and call this the new terminal node. Now start with the predecessor of all these terminal nodes and repeat the previous step, folding back the game tree again. Continue to do so until one reaches an initial node of the tree. Backward induction can eliminate strategies that are weakly dominated, hence, it can also remove Nash equilibria that are based on such strategies.

Let us illustrate the method of backward induction for the dynamic game of perfect information in Figure A.2-A. There are two Nash equilibria (R,r) and (L,l) (check that the strategy of each player is a best response to the strategy of the other player). However, the equilibrium (L,l) relies on a *non-credible threat* by player 2 to play l if he gets the move (because the payoff of player 2 from the action r is higher). The backward induction equilibrium eliminates precisely these sorts of non-credible threats. Using the method outlined above, the reader can check that the equilibrium backward induction outcome for the extensive form game in Figure A.2-A is uniquely given by (R,r).

For dynamic games of imperfect information, we need to define the concept of a *proper subgame*. Formally, a proper subgame is (i) a node t that belongs to an information set that is a singleton ($h(t) = \{t\}$) and, (2) all successors $S(t)$ of node t such that for any $t_j \in S(t)$, the information set containing t_j , $h(t_j)$, is a *proper subset* of $S(t)$ (i.e., $h(t_j) \subseteq S(t)$). Consider Figure A.2-B. The third node at which the two players play a simultaneous move BoS game is an example of a proper subgame. The definition of a proper subgame allows us to consider the game starting at the decision node t as a game in its own right, which motivates the next definition. We do not define strategies in extensive form games formally, but the reader can think of such strategies as complete contingent plans of action that specify the action which a player will take at each possible node, when called upon to make a move.

Definition A.6 (*Subgame perfect Nash equilibrium*): A Nash equilibrium of an extensive form game is “subgame perfect” if the restriction of strategies to any subgame, constitute a Nash equilibrium for that subgame.⁹

As an illustration of a subgame perfect equilibrium, apply Definition A.6 to the extensive form game in Figure A.2-B. The game starting at the third node when player 1 has the move and both players play a simultaneous move BoS game is a proper subgame (check that it satisfies the definition of a proper subgame). In the BoS game, there are two equilibria in pure strategies: (M, M) and (B, B). There is also a mixed strategy equilibrium in which player 1 plays M with probability 2/3 and player 2 plays M with probability 1/3 for an expected payoff to each player of 2. The mixed strategy (expected payoff of 2) is dominated by the two pure strategy equilibria (which give a payoff of either 6 or 3), hence, we rule it out by a dominance argument.

We now need to roll back the game tree for each of the two cases that arise with a play of (M, M) and (B, B). Thus, replace the third node respectively with payoffs of (6, 3) and (3, 6). In both cases when the tree is rolled back to the move of player 2, he finds it optimal to play L. In the first case (where (M, M) is the outcome in the BoS game) player 1 plays L at the first node so that in the subgame perfect outcome, the payoffs are (6, 3). In the second case (where (B, B) is the outcome in the BoS game) player 1 finds it optimal to play T, so the payoffs in the subgame perfect outcome are (4, 4).

In some cases, we may need a refinement of subgame perfection in order to pin down beliefs at *out of equilibrium actions*. The relevant refinement has two main components.

1. For each information set $h \in H$, a specification of a mixed strategy that will be played at all nodes in that information set. Denote by σ the resulting strategy profile for the entire extensive form game when one pools this information for each possible information set.
2. For each information set $h \in H$, a specification of the beliefs about the node in h where the player makes a move. Thus, for any information set $h \in H$ and for any node $t \in h$, we may denote the belief of player $i(h)$ who moves at this information set by $\mu(t) \geq 0$ such

⁹ If an extensive form game has no proper subgames, then a Nash equilibrium of the game is also subgame perfect.

that $\sum_{t \in h} \mu(t) = 1$. Denote the set of such beliefs for the entire game, considered across all possible information sets, by μ .

Definition A.7 (*Perfect Bayesian equilibrium (PBE)*; Fudenberg and Tirole, 1991b): A perfect Bayesian equilibrium (PBE) is a pair of strategy vector and belief vector for the entire game, (σ, μ) , such that the following conditions hold.

- (i) (*Sequential rationality*): At any information set $h \in H$ the action of player $i(h)$ who moves at this information set maximizes his payoff given (i) his belief at each node in this information set, $\mu(t)$, $t \in h$, and (ii) the subsequent strategies of the players as listed in σ .
- (ii) (*Bayes' rule applied along the path of play*): Consider the set of information sets that are reached with positive probability, conditional on the actions, σ . On these information sets, we require a mutual consistency of beliefs with actions in the sense that beliefs are computed from the strategies given in Bayes' rule.

One problem with the definition of a PBE in Definition A.7 is that it is silent about beliefs at those information sets that are *off the path of play*, i.e., never reached with strictly positive probability, given the equilibrium strategies; see Definition A.9 below. The reason is that Bayes' rule does not allow us to specify any probabilities at such nodes (see the discussion following Definition A.9); technically the denominator in Bayes' rule becomes zero, hence, the ratio is undefined. In order to ensure that all information sets are reached with a positive probability, a *sequential equilibrium* focuses on sequences of strictly mixed strategies $\{\sigma_n\}$ in which all actions are played with some probability such that $\lim_{n \rightarrow \infty} \sigma_n = \sigma$.¹⁰ The corresponding beliefs are found by Bayes' rule, which are now defined at each node, and given by $\{\mu_n\}$ such that $\lim_{n \rightarrow \infty} \mu_n = \mu$. A sequential equilibrium is found by $\lim_{n \rightarrow \infty} (\sigma_n, \mu_n) = (\sigma, \mu)$ and it is subgame perfect. However, in the limiting case, Bayes' rule cannot be applied to information sets that are off the path of play. In applications, the PBE is often simpler to apply as compared to a sequential equilibrium. Further refinements, such as the Cho–Kreps intuitive criteria (see signaling games below) are needed to put more restrictions on such beliefs.

To see some of these equilibrium concepts in action, consider the two games in Figure A.3 that are due to Kreps (1990).

Consider first the game in Figure A.3-A. The game begins with a move by player 1, who plays either A or D. If player 2 gets the move, he can either end the game by playing action a or can play




Figure A.3 An illustration of sequential rationality with two examples.

¹⁰ The sequential equilibrium concept was introduced by Kreps and Wilson (1982). The basic idea is very similar to the notion of a *trembling hand perfect equilibrium* due to Selten (1975).

action b to give the move to player 3. If given the move, player 3 does not know whether he is at the left or the right hand side node, thus both nodes belong to the same information set, hence, they are connected with a dotted line. Player 3, the only player whose information set is not a singleton, can play either L or R. Subgame perfection cannot refine the set of Nash equilibria here because there are no proper subgames.

One Nash equilibrium of the game is (D, a, L). The reader can check that the action of each player is optimal given the actions of the other players. In the equilibrium (D, a, L), the action a of player 2 is *off the path of play* so we may wonder if it is reasonable? Given that player 3 is supposed to play L, player 2 if given the move should never play a; instead, he should play b to get a payoff of 5 instead of 3. But then the profile (D, b, L) is no longer an equilibrium because player 1 would wish to deviate to A (however, in that case (A, b, L) is no longer an equilibrium because player 3 would wish to deviate to R).

Let us now use sequential rationality to show that the choice of a by player 2 in the equilibrium (D, a, L) is unreasonable. Since the information sets of players 1 and 2 are singletons, they must place probability equal to one of being at their respective nodes if given a chance to move. However, there are two nodes in player 3's information set. Suppose that player 3 assigns a probability of 1 to being at the left node. First note that the action profile for this game is mutually consistent with the beliefs of players (Definition A.7(ii)). Given the prior play of the game in the strategy profile (D, a, L), Bayes' rule recommends player 3 to assign probability 1 of being at the left hand side node. Given that the subsequent path of play specifies that player 3 will play L, player 2 should play b not a. Thus, the play of a by player 2 is not sequentially rational.

Consider the strategy profile (A, b, L) and the belief of player 3 that he is at the left node in his information set with probability 0.95. It is easily checked that the actions of the players are sequentially rational but the condition in Definition A.7(ii) is violated. Given the path of play leading up to the choice by player 3, consistent with Bayes' rule, player 3 should assign probability 1 of being at the right node in the information set. Thus, a PBE can be used to rule out nonsensical equilibria such as (A, b, L).

Let us now consider the game in Figure A.3-B. The extensive form features of this game should by now be apparent to the reader. In order to apply the definition of a PBE, we need to specify the strategy at each information set and the beliefs of players where the information sets are not singletons. Player 2's information set is not a singleton. Suppose that we wish to check for the sequential rationality of the strategy profile (L, d, D). The strategies of players 2 and 3, in this case, are off the path of play, hence, we cannot use Bayes' rule to pin down the precise beliefs that player 2 holds at any of the nodes in his information set. Rather, any beliefs will do (see below). So let us suppose that player 2 assigns a probability of 0.5 of being at any of his nodes. The payoff to player 2 from playing u is $0.5(8) + 0.5(1) = 4.5$ and the payoff from playing d is $0.5(4) + 0.5(2) = 3.0$. Clearly the play of d by player 2 in the strategy profile (L, d, D) is not sequentially rational.

A.7 Games of incomplete information

In the games analyzed so far, there is complete information in the sense that each player knows the *types* of other players. Incomplete information arises when players are not sure of the types of other players. For instance, in a Cournot oligopoly, a firm's type could be its privately known cost; each type then corresponds to a different level of privately known cost.

The set of types of any of the finite number of players $i \in N$ is given by $T_i = \{t_{i1}, t_{i2}, \dots, t_{i\tau_i}\}$; τ_i is the number of possible types of player i , which can potentially be an infinite number, although

for most applications we assume a finite number of types. We take types to differ in their payoff functions. For instance, in a Cournot oligopoly, the different types of the same firm may have different marginal costs. The Cartesian product of types of all players is given by the set $T = T_1 \times T_2 \times \dots \times T_n$. The prior probability of type $t_i \in T_i$ is given by $p(t_i)$, which is common knowledge among all players. The prior probability of a particular vector of types $(t_1, t_2, \dots, t_n)$ is given by $p(t_1, t_2, \dots, t_n)$ and it is also common knowledge. For any player i , the set of types of all other players belongs to the set $T_{-i} = \times_{j \neq i} T_j$; a typical member of T_{-i} is denoted by $t_{-i} = (t_1, \dots, t_{i-1}, t_{i+1}, \dots, t_n)$. The set of actions of any player $i \in N$ is given by S_i and the same action set S_i is available to all types of the player.

Each player has incomplete information about the exact type of the opponents, although he/she is aware of the prior probability, p , of the vector of types of the other players. In this sense, this is a *game of incomplete information* with common priors.¹¹ Following Harsanyi, we assume that nature moves first and draws some vector of types $\mathbf{t} = (t_1, t_2, \dots, t_n) \in T$. To each player $i \in N$, nature privately reveals the corresponding type t_i that was drawn but does not give this information to any other player. This *changes a game of incomplete information into a game of imperfect information*.

A.7.1 Finite Bayesian games of incomplete information

In static games of incomplete information, the players choose their actions simultaneously. The types of the players may be correlated, hence, the probability assigned by player i that the type vector of other players is t_{-i} is given by Bayes' rule

$$\mu(t_{-i} | t_i) = \frac{p(t_{-i}, t_i)}{p(t_i)} = \frac{p(t_{-i}, t_i)}{\sum_{t_{-i} \in T_{-i}} p(t_{-i}, t_i)}.$$

When types are independent, then $\mu(t_{-i} | t_i)$ does not depend on t_i and $\mu(t_{-i} | t_i) = \mu(t_{-i} | t'_i)$ for all $t_i, t'_i \in T_i$.

A finite Bayesian game is denoted by $G(N, S, \pi, T, \mu)$ in which the number of players is finite and their strategy sets are also finite. In normal form games, S_i simply gives the set of actions available to player i , but in extensive form games, the elements of S_i specify what the player does at each information set. The strategy profiles of all players other than player i are given by $S_{-i} = \times_{j \neq i} S_j$. μ , which is common to all players, gives the set of probability distributions over T . The payoff of player i , $\pi_i : S \times T \rightarrow \mathbb{R}$, thus, the payoffs depend on the actions of all the players as well as the types. We shall denote by $\pi_i(s_i, s_{-i}; t_i, t_{-i})$ the payoff of a type t_i of player i from the pure strategy profile (s_i, s_{-i}) when the types of the other players are given by t_{-i} . Denote the vector of payoffs of all players by $\pi = (\pi_1, \pi_2, \dots, \pi_n)$.

A mixed strategy for player i , σ_i , specifies the probability with which each of the pure strategies in the set S_i is played by each of the possible types of the player in the set T_i . Thus, $\sigma_i : T_i \rightarrow \Sigma_i$ and $\sigma_i(s_i | t_i)$ is the probability with which type t_i of player i plays strategy s_i . The mixed strategy profile of all players other than player i is denoted by $\sigma_{-i} : T_{-i} \rightarrow \times_{j \neq i} \Delta S_j$, where ΔS_j denotes probability distributions over the set S_j . We denote the probability that some realization of the types of other players $t_{-i} \in T_{-i}$ play some pure strategy profile $s_{-i} \in S_{-i}$ by $\sigma_{-i}(s_{-i} | t_{-i})$.

¹¹ The common prior assumption is a strong one. Gintis (2009) argues that traditional game theory provides weak justification for this assumption. He proposes a correlating device, such as social norms, which may induce a set of common priors.

A *Bayesian Nash equilibrium* (BNE) of the game $G(N, S, \pi, T, \mu)$ specifies an optimal strategy for each type of each player conditional on the beliefs. The reason is that since players only know their own types but not the types of other players, each player needs to take into account the strategy of each of the possible types of the other players.

Definition A.8 (*Bayesian Nash equilibrium*): Consider the static Bayesian game $\{N, S, \pi, T, \mu\}$. The strategy vector (σ_i, σ_{-i}) is a Bayesian Nash equilibrium if for all $i \in N$, all types $t_i \in T_i$ and all actions s_i^* such that $\sigma_i(s_i^* | t_i) > 0$ we have

$$s_i^* \in \arg \max_{s_i \in S_i} \sum_{t_{-i} \in T_{-i}} \mu(t_{-i}, t_i) \sum_{s_{-i} \in S_{-i}} \sigma_{-i}(s_{-i} | t_{-i}) \pi_i(s_i, s_{-i}; t_i, t_{-i}). \quad (\text{A.13})$$

The Nash equilibrium is a special case of a BNE when each player has a single type that is public information, thus, the BNE imposes relatively more stringent rationality requirements. For instance, the evidence in Chapter 19 on judgment heuristics shows that players typically do not use Bayes' rule. This per se is not fatal to the concept of a BNE; one may replace μ by non-Bayesian beliefs in (A.13). However, there are other reasons why a BNE might not be able to explain the evidence. For instance, it is linear in beliefs, and does not have reference dependence and loss aversion. By contrast, the evidence reviewed in Part 1 on behavioral decision theory suggests that non-linear probability weighting, reference dependence, and loss aversion, are important features of human behavior.

A.7.2 Signaling games

We motivate signaling games by considering the game used in the experiments of Brandts and Holt (1992) in Figure A.4; for the experimental results on this game, see Chapter 12.

Signaling games are two-player games with a *sender* (Player 1) and a *receiver* (Player 2). The sender knows his own type but the receiver does not, hence, this is a game of incomplete information. The set of possible types of the sender is $T = \{t_1, t_2, \dots, t_n\}$. The receiver has prior beliefs about the types in T ; the prior probability of type $t_i \in T_i$ is given by $p(t_i)$. As in static games of incomplete information, we use the device of an external player, nature, who chooses the type of the sender and reveals it only to the sender. The sender can then send a set of signals $S_1 = \{s_1, s_2, \dots, s_m\}$. The receiver observes the signal (but not the type of the sender) and then chooses from a set of actions $S_2 = \{a_1, a_2, \dots, a_q\}$. Generic elements of S_1 and S_2 are denoted respectively by s and a .




Figure A.4 The signaling game in Brandts and Holt (1992).

Source: With permission from the American Economic Association.

In the signaling game in Figure A.4, the sender has two types $T = \{L, H\}$ and the initial move by nature is represented by the hollow circle in the middle of each figure. The initial beliefs of the receiver are that Type L is chosen with probability $1/3$ (game proceeds to the top of the figure) and type H with probability $2/3$ (game proceeds to the lower end of the figure). The sender can send two possible signals $S_1 = \{A, B\}$. The receiver can take two possible actions $S_2 = \{C, D\}$. Since the receiver does not know the type of the sender, having observed any signal, say, A, the receiver does not know if play has proceeded to the upper right hand node or the lower right hand node. This imperfect information is captured by the two information sets of the receiver that are connected by dotted lines. The payoffs are shown at the terminal nodes.

Signals and actions can be in pure or in mixed strategies. We denote by $\sigma_1(s, t)$ the mixed strategy of the sender. It is to be interpreted as the probability with which a sender of type t sends signal s . Each type is required to send at least one signal, so for all $t \in T$ we have $\sum_{s \in S_1} \sigma_1(s, t) = 1$. We denote by $\sigma_2(a, s)$ the mixed strategy of the receiver; it is the conditional probability with which the receiver plays action a after observing the signal, s . We assume that on observing any signal, receivers do take some action, i.e., for all $s \in S_1$ we have $\sum_{a \in S_2} \sigma_2(a, s) = 1$.

Payoffs accrue at the end of the game, at the terminal nodes, once the receiver has taken his action. The payoff of the sender is given by $\pi_1(t, s, a)$; it depends on his type, signal, and the action subsequently taken by the receiver. The payoff of the receiver is given by $\pi_2(t, s, a)$.

On receiving a signal, the receiver updates his prior beliefs about the possible types of the sender by using Bayes' rule to form the appropriate posterior. $\mu(t | s)$ is the posterior probability that an observed signal s originates from type t of the sender. In conformity with Bayes' rule, it is given by

$$\mu(t | s) = \frac{p(t)\sigma_1(s, t)}{\sum_{t' \in T} p(t')\sigma_1(s, t')}, \quad (\text{A.14})$$

where the numerator gives the probability that the signal s comes from type t and the denominator gives the a-priori probability that the signal is s .

Definition A.9 (“On the path of play” and “off the path of play”): A signal s is said to be along the path of play if the denominator in (A.14) is strictly positive, i.e., non-zero. Conversely if the denominator is zero then the signal s is off the path of play; such a signal is not sent by any of the types.

Since division by zero is undefined, we cannot use Bayes' rule to pin down posterior beliefs for off the path of play signals. In this case, any assignment of probabilities to off the path play can be made (this may not hold when we consider refinements of a perfect Bayesian equilibrium; see definition below).

Definition A.10 (Perfect Bayesian equilibrium (PBE) in signaling games): A perfect Bayesian equilibrium (PBE) in a signaling game is the triple $(\sigma_1, \sigma_2, \mu)$ that satisfies the following.

- (a) If any signal is played by the sender with strictly positive probability then it must be a maximizing choice for the sender. Thus, if $\sigma_1(s^*, t) > 0$ then

$$s^* = \arg \max_{s \in S_1} \sum \sigma_2(a, s) \pi_1(t, s, a). \quad (\text{A.15})$$

- (b) On receiving a signal s , if any action a is played with strictly positive probability by the receiver then it must maximize the receiver's expected utility. Thus, if $\sigma_2(a^*, s) > 0$ when signal s is observed, then

$$a^* = \arg \max \sum_{t \in T} \mu(t | s) \pi_2(t, s, a). \quad (\text{A.16})$$

- (c) The posterior beliefs in (A.16) are constructed by using Bayes' rule as given in (A.14) for signals that are "on the path of play."
 (d) For signals that are "off the path of play," some beliefs must be assigned such that (A.16) holds.

From (A.15), (A.16), the reader can immediately see that sequential rationality is built into the definition of an equilibrium. The sender plays optimally given the subsequent equilibrium action of the receiver and the receiver plays optimally given the equilibrium signal sent by the sender. Furthermore, in computing the optimal action, the receiver uses Bayes' rule along the path of play—thus, actions and beliefs must be mutually consistent. However, and as noted earlier, the equilibrium may impose beliefs on signals that are off the path of play, which may not be reasonable; this leads to a refinement of the equilibrium that we consider below. The rationality requirements in a PBE of a signaling game are even higher than those of a BNE and these mainly come about from an imposition of sequential rationality that was missing in a BNE.

Three kinds of equilibria are of typical interest in signaling games. In a *separating equilibrium*, all types of the sender send distinct signals, hence, immediately revealing their types to the receiver. The receiver then assigns a posterior probability of 1 to the corresponding type of a sender after observing a particular signal, which is consistent with Bayes' rule. In a *pooling equilibrium*, all types of the sender send identical signals, hence, as required by the Bayes' rule, the posterior beliefs of the receiver are identical to the prior beliefs. In a *hybrid equilibrium*, the situation is intermediate between a separating and a pooling equilibrium. For instance, one of the types may choose to separate (in strategy), while the other plays a mixed strategy—choosing to separate with some probability and to pool with the complementary probability. The posterior beliefs of the receiver then reflect this probability mixture appropriately.

A drawback of the PBE concept in Definition A.10(d) is that it does not require out of equilibrium beliefs to be reasonable enough. One natural restriction is suggested by Cho and Kreps (1987). Suppose that for some type of the sender $t \in T$ there is a signal s that strictly dominates another signal s' in the following sense

$$\min_{a \in S_2} \pi_1(t, s, a) > \max_{a \in S_2} \pi_1(t, s', a), \quad (\text{A.17})$$

i.e., the maximum payoff that the sender gets from sending the signal s' (irrespective of the actions of the receiver) is lower than the minimum payoff from signal s . The *simple dominance* criteria then requires that on observing the signal s' , even if it is out of equilibrium, the receiver should assign probability zero that it was sent by a sender of type t .

A second refinement, due to Cho and Kreps (1987), is *equilibrium dominance*. Suppose that the equilibrium signal of a type t sender is s^* and the equilibrium action of the receiver is a^* . Then signal s' is *equilibrium dominated* for the type t sender if

$$\pi_1(t, s^*, a^*) > \max_{a \in S_2} \pi_1(t, s', a). \quad (\text{A.18})$$

Equilibrium dominance in (A.18) provides an important restriction on out of equilibrium beliefs, which we state next.

Definition A.11 (*Intuitive criterion; Cho and Kreps, 1987*): Suppose that the signal s' is off the equilibrium path (as defined above). In any equilibrium, if (A.18) holds then the receiver should assign a posterior probability of zero that the sender of type t will send the signal s' . This is known as the intuitive criterion or the equilibrium dominance criterion.

We do not pursue further refinements of games with incomplete information, partly because the empirical support for such refinements is quite weak; see Chapter 12 on behavioral game theory.

Let us now apply these ideas to determine the equilibrium outcomes in the signaling game in Figure A.4. The game has only two equilibria and these are both pooling equilibria (we leave it as an exercise to the reader to show that there are no separating equilibria). The equilibrium responses are shown by the bold lines in Figure A.4. (1) In the left panel, in equilibrium, both types of senders play B. The receiver responds to B by playing C and to the non-equilibrium signal A by playing D. (2) In the right panel, both types of senders send the signal A and the receiver plays C in response to the equilibrium signal A and D in response to the non-equilibrium signal B.

In each pooling equilibrium, the posterior beliefs of the receiver are identical to prior beliefs. Bayes' rule permits us to place any restrictions on the out of equilibrium beliefs. One set of out of equilibrium beliefs that can support these two pooling equilibria (see Definition A.10(d)) is the following. Whenever the receiver observes an out of equilibrium signal by the sender, it is assumed that the sender is of type L. The reader can check that the stated equilibrium actions and beliefs constitute a PBE of the signaling game (check all components of Definition A.10).

To see if these equilibria are reasonable, consider the pooling equilibrium in the right panel. The out of equilibrium beliefs require that if the receiver observes signal B, he should assume that it is sent by a type L sender. We can use the Cho–Kreps (1987) *equilibrium dominance* criteria, described above, to rule out this pooling equilibrium. In the right panel, type H of the sender gets an equilibrium payoff of 120, while type L gets an equilibrium payoff of 140. If instead, type L decided to send an out of equilibrium signal, B, then the highest payoff that he could get is always lower than 140. Hence, type L will never want to send an out of equilibrium signal. However, the out of equilibrium signal is not equilibrium dominated for type H because type H can get a maximum of $140 > 120$. Thus, the out of equilibrium beliefs of player 2 are not intuitive or reasonable. Player 2 should have beliefs such that an out of equilibrium signal of B could only come from type H. In that case, the optimal response of the receiver to signal B is to play C. But then B rather than A is optimal for type H, breaking down the pooling equilibrium in which both types of senders send signal A. This is the reason why the right panel is labeled as the unintuitive equilibrium. There is no problem with the pooling equilibrium in the LHS panel. It survives equilibrium dominance (the reader should check this¹²).

In signaling games, a signal s is dominated for type t of a player, if there is another signal s' , such that the lowest possible payoff of type t from the signal s' is greater than the highest possible payoff from signal s ; we may term this as *simple dominance*. Equilibrium dominance is stronger than *simple dominance*. The concept of simple dominance cannot be used to eliminate the *unintuitive equilibrium* in the right panel of Figure A.4. This is because, for type L, neither of the two actions A or B dominates the other in the sense of (A.17). Refinements of PBE that are stronger than equilibrium dominance such as *divinity* and *strategic stability* can also be used to rule out the unintuitive equilibrium on the RHS. Hence, experimental tests that show subjects do not play the equilibrium in the right panel are also consistent with these other refinements.

¹² The reader should also check that there are no separating equilibria.

A.8 Correlated equilibrium

Consider an n -player *strategic form* game, $G(N, S, u)$, where probability distributions over S_i are denoted by Σ_i , given in (A.1), and probability distributions over S are denoted by Σ . Choose a particular member of Σ , say, q . Then for any $s \in S$, $q(s)$ is the probability of the strategy profile s , which implies $\sum_{s \in S} q(s) = 1$.

Definition A.12 (*Correlated equilibrium; Aumann, 1974*): A probability distribution $q \in \Sigma$ is a *correlated equilibrium* if $\forall i$ and $\forall s_i \in S_i$, with $q(s_i) > 0$:

$$\sum_{s_{-i} \in S_{-i}} q(s_{-i} | s_i) \pi_i(s'_i, s_{-i}) \leq \sum_{s_{-i} \in S_{-i}} q(s_{-i} | s_i) \pi_i(s_i, s_{-i}), \text{ for all } s'_i \in S_i.$$

Definition A.12 can be motivated as follows.

1. Suppose that all players know that an element $s \in S$ will be drawn using the probability distribution q .
2. Once $s = (s_1, \dots, s_n)$ is drawn, then player $i \in N$ is told of his pure strategy, $s_i \in S_i$, but not about the remaining strategy profile, $s_{-i} \in S_{-i}$.
3. q is a correlated equilibrium if all players $i \in N$ wish to follow their recommended strategy s_i and not change to some other $s'_i \in S_i$, when they know that the strategies are drawn with respect to the underlying distribution, q .

We now define the more general notion of *coarse correlated equilibrium*. To understand the intuition, let us now change the three steps above as follows:

1. Suppose that all players know that an element $s \in S$ will be drawn using the probability distribution q .
2. Any player $i \in N$ can opt out at this stage and play any $s'_i \in S_i$.
3. Suppose that the set of players who do not wish to opt out is $N' \subseteq N$. A strategy profile $s = (s_1, \dots, s_n)$ is drawn according to q . All players in N' are now committed to playing their part in the profile s (so player $i \in N'$ must play $s_i \in S_i$).
4. q is a coarse correlated equilibrium if no players prefer to opt out in stage 2, i.e., $N' = N$.

Definition A.13 (*Coarse correlated equilibrium; Moulin and Vial, 1978*): A probability distribution $q \in \Sigma$ is a *coarse correlated equilibrium* if $\forall i$ and $\forall s'_i \in S_i$,

$$\sum_{s_{-i} \in S_{-i}} q_{-i}(s_{-i}) \pi_i(s'_i, s_{-i}) \leq \sum_{s \in S} q(s) \pi_i(s). \quad (\text{A.19})$$

The LHS of (A.19) is the expected payoff of player i if he opts out in stage 2 and the RHS is his payoff if he chooses to opt in. If q is a correlated equilibrium then it is also a coarse correlated equilibrium. In the special case of a product distribution, i.e., when $q(s) = \prod_{i=1}^n q_i(s_i)$, the following three concepts coincide: correlated equilibrium, coarse correlated equilibrium, and Nash equilibrium.

REFERENCES

- Aumann, R. J. (1974). Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics* 1: 67–96.
- Aumann, R. J. and Brandenburger, A. (1995). Epistemic conditions for Nash equilibrium. *Econometrica* 63: 1161–80.
- Beard, T. R. and Beil, R. O. Jr. (1994). Do people rely on the self-interested maximization of others: an experimental test. *Management Science* 40: 252–62.
- Bernheim, D. (1984). Rationalizable strategic behavior. *Econometrica* 52: 1007–28.
- Brandts, J. and Holt, C. H. (1992). An experimental test of equilibrium dominance in signaling games. *American Economic Review* 82: 1350–65.
- Cho, I.-K. and Kreps, D. M. (1987). Signaling games and stable equilibria. *Quarterly Journal of Economics* 102: 179–221.
- Debreu, G. (1952). A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences* 38: 886–93.
- Dixit, A. and Skeath, S. (2004). *Games of Strategy*. 2nd edition. New York: W. W. Norton.
- Fudenberg, D. and Levine, D. K. (1998). *The Theory of Learning in Games*. Cambridge, MA: MIT Press.
- Fudenberg, D. and Tirole, J. (1991a). *Game Theory*. Cambridge, MA: MIT Press.
- Fudenberg, D. and Tirole, J. (1991b). Perfect Bayesian and sequential equilibrium. *Journal of Economic Theory* 53: 236–60.
- Gibbons, R. (1992). *Game Theory for Applied Economists*. Princeton, NJ: Princeton University Press.
- Gintis, H. (2009). *The Bounds of Reason: Game Theory and the Unification of Behavioral Sciences*. Princeton, NJ: Princeton University Press.
- Glicksberg, I. L. (1952). A further generalization of the Kakutani fixed point theorem, with application to Nash equilibrium. *Proceedings of the American Mathematical Society* 3: 170–4.
- Kreps, D. M. (1990). *A Course in Microeconomic Theory*. Princeton, NJ: Princeton University Press.
- Kreps, D. and Wilson, R. (1982). Sequential equilibria. *Econometrica* 50: 863–94.
- Maschler, M., Solan, E., and Zamir, S. (2013). *Game Theory*. Cambridge: Cambridge University Press.
- Moulin, H. and Vial, J. P. (1978). Strategically zero-sum games: the class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory* 7: 201–21.
- Myerson, R. B. (1991). *Game Theory: Analysis of Conflict*. Cambridge, MA: Harvard University Press.
- Nash, J. F., Jr. (1950). Equilibrium points in n -person games. *Proceedings National Academy of Sciences* 36: 48–9.
- Osborne, M. J. (2004). *An Introduction to Game Theory*. Oxford: Oxford University Press.
- Osborne, M. J. and Rubinstein, A. (1994). *A Course in Game Theory*. Cambridge, MA: MIT Press.
- Pearce, D. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica* 52: 1008–50.
- Polak, B. (1999). Epistemic conditions for Nash equilibrium, and common knowledge of rationality. *Econometrica* 67: 673–6.
- Schotter, A., Weigelt, K., and Wilson, C. (1994). A laboratory investigation of multi-person rationality and presentation effects. *Games and Economic Behavior* 6: 445–68.
- Selten, R. (1975). A reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* 4: 25–55.
- von Neumann, J. (1928). Zur theorie der gesellschaftsspiele. *Mathematische Annalen* 100: 295–320.
- Watson, J. (2013). *Strategy: An Introduction to Game Theory*. 3rd edition. New York: W. W. Norton.

INDEX OF NAMED AUTHORS

- Abbink, K. 347, 358, 380, 381–3, 436
 Abdellaoui, M. A. 31, 92, 108–12, 127–9, 133, 140, 142–3, 144–8, 178–9, 183, 206, 211, 229, 286, 304–8, 619
 Abeler, J. 13, 112, 131, 167, 170, 346, 360–1, 459, 509, 1463
 Abreu, D. 734, 735, 736
 Acemoglu, D. 1347, 1408
 Ackert, L. F. 36, 416, 417
 Agranov, M. 878, 922
 Ahn, D. 299, 301
 Ahumada, H. A. 617
 Ainslie, G. W. 38, 583, 587, 595–8, 616
 Ajzen, I. 1345, 1376
 Akerlof, G. A. 61, 357, 455, 465–6, 473–4, 843, 863, 972, 1348, 1496
 Alesina, A. 1284, 1347, 1409
 Alexander, R. D. 395, 1040, 1066, 1078
 Alger, I. 1040, 1081, 1083, 1084
 Allais, M. 3, 10, 26, 83, 95, 96–7, 98, 101, 110, 976
 Allcott, H. 1409, 1632
 Allingham, M. G. 215, 243
 al-Nowaihi, A. 8, 24, 29, 30–2, 36, 38, 106, 112, 114–21, 127, 139, 150–2, 191–7, 198–9, 215–6, 243–7, 251, 254, 400, 403, 412, 415–23, 428–9, 602, 605–12, 620–5, 628, 632–8, 640–1, 880, 962, 974, 977, 980–1, 984, 1396, 1667–8, 1707
 Alós-Ferrer, C. 1159, 1167–8
 Amador, M. 1228, 1266, 1269, 1270
 Amaldoss, W. 715, 747, 748, 1125, 1127
 Andersen, S. 12, 97, 179, 344, 350, 594, 618
 Anderson, J. C. 13, 1386
 Andreoni, J. 36, 115, 199, 200–5, 247, 345, 351, 354–6, 368, 373, 377, 399, 431–3, 859, 1066, 1130, 1670
 Ang, J. S. 1498, 1533
 Angeletos, G.–M. 651, 655
 Angrist, J. 483, 554, 556, 560
 Anscombe, F. J. 91, 289, 290, 292
 Arad, A. 923, 924
 Ariely, D. 16, 46, 217–21, 482–3, 543–7, 550, 554–5, 689, 1345, 1371, 1454, 1472–3, 1598
 Arkes, H. R. 17, 1386, 1463, 1474
 Arrow, K. 193, 194, 293
 Arthur, W. B. 1092, 1098, 1102, 1156
 Ashraf, N. 3, 658, 691
 Atkinson, R. 1092, 1098
 Attanasi, G. 937, 950, 952
 Aumann, F. J. 1707, 1723
 Aumann, R. J. 91, 289, 290, 292, 729, 745, 848, 879–80, 962, 964, 976, 981, 1348
 Avery, C. 994, 997, 1001
 Axelrod, R. 342, 1039, 1055, 1077
 Ayton, P. 45, 1343, 1357, 1358, 1427
 Babcock, L. 33, 46, 264, 340, 720, 813, 1345, 1372–3
 Bacharach, M. 881, 952, 987
 Baillon, A. 286, 297, 299, 304–6, 308
 Baker, M. 1462, 1488, 1526, 1533, 1541, 1542
 Bandiera, O. 387, 481
 Bandura, A. 548, 1584
 Banerjee, A. V. 890, 1350, 1408, 1437
 Banks, J. S. 721, 822, 826
 Baran, N. M. 349, 386, 394
 Barber, B. M. 1354, 1486, 1493, 1496, 1501–2, 1517
 Barberis, N. C. 29, 55–6, 111, 199, 207, 216, 242, 263, 1355, 1455, 1480, 1486–7, 1491, 1494–6, 1500, 1520, 1525
 Barbey, A. K. 1345, 1349, 1377, 1378
 Bardsley, N. 3, 6, 8, 10, 12, 16, 17, 18–19, 762
 Barenstein, M. 689
 Bar-Hillel, M. 45, 46, 1343, 1345, 1352, 1375–6
 Barmettler, F. 12, 348, 387, 388
 Bartling, B. 439, 481, 482, 490, 522, 525, 526–7, 544–8
 Barton, R. A. 1683
 Bartz, J. A. 1648, 1694
 Basu, K. 892, 1440
 Bateman, I. 224, 617
 Battalio, R. C. 103, 104, 716, 767–8, 773, 783, 1094, 1106, 1116, 1120
 Battigalli, P. 36, 345, 354, 879, 927–8, 931, 933, 935–6, 938, 949–50, 960–1, 1233
 Baucells, M. 199, 619
 Baumeister, R. F. 1231, 1291, 1294, 1685
 Baumgartner, T. 1647, 1677, 1679, 1680–1, 1684
 Bazerman, M. H. 718, 810, 812, 991, 993, 1095, 1134
 Beard, T. R. 713, 724, 725, 726, 1713
 Becker, G. M. 4, 18
 Becker, G. S. 194, 219, 229, 1223, 1226, 1230, 1236–7, 1241–2, 1410
 Beil, R. O., Jr. 713, 724, 725, 726, 1713
 Bell, D. 101, 113, 154, 178, 181, 1245
 Bellemare, C. 207, 240, 346, 362, 363, 386, 403
 Bénabou, R. 17, 341, 345, 353, 477, 481, 507, 528, 531–3, 539, 552, 1228–9, 1261, 1264, 1272–3, 1289, 1311–2, 1316, 1319–21
 Benartzi, S. 29, 56, 144, 207, 215, 241–2, 257, 658, 689, 1454–5, 1470, 1481–4, 1493, 1499–1500, 1596
 Ben-David, I. 56, 1350, 1449, 1528
 Benhabib, J. 618, 625, 1080, 1292
 Benjamin, D. J. 454, 467–8, 477, 604, 1033, 1228, 1295, 1297, 1600
 Bentham, J. 3, 1228, 1242, 1270, 1600
 Benz, M. 346, 349, 358, 366, 394, 1279
 Benzion, U. 596, 597–9
 Bereby-Meyer, Y. 786, 1102
 Bergstrom, T. 1082, 1083

- Bernard, V. L. 1486, 1498, 1520
 Bernasconi, M. 108, 247
 Bernheim, B. D. 36, 354, 658, 914, 986, 1227, 1242, 1257, 1292, 1582, 1595, 1601, 1604, 1611–2, 1614–8, 1621, 1645, 1670, 1706
 Berns, G. S. 116, 1289, 1646, 1663, 1667
 Bertrand, M. 1436, 1488, 1527
 Beshears, J. 1346, 1460, 1463, 1483, 1579, 1580
 Bettman, J. R. 1425, 1488, 1527
 Bewley, T. F. 292–3, 297, 340, 346, 1473
 Beyth, R. 1346, 1385
 Bhatt, M. 1647, 1648, 1689, 1691, 1692
 Biais, B. 46, 54, 56, 1346, 1387–8, 1390, 1430, 1449, 1493
 Bicchieri, C. 976, 1178
 Bickart, K. C. 1647, 1683
 Billig, M. 454, 466
 Binmore, K. G. 59, 345, 351, 521, 714, 718–9, 738, 741, 749, 794–5, 798, 801, 802, 904, 1118, 1143, 1153
 Birnbaum, M. H. 31, 113, 114, 173, 174, 188, 189, 190, 191
 Bisin, A. 1068, 1080, 1292
 Bjørnerstedt, J. 1087, 1089, 1091
 Black, F. 1485, 1491, 1502
 Blanchard, O. J. 256, 723, 856, 857, 1494, 1503
 Blanchet, D. 1273, 1618
 Blaug M. 5, 6
 Blavatsky, P. R. 97, 217, 240
 Bleichrodt, H. 93, 112, 114, 145–7, 182–3, 188, 607
 Blinder, A. S. 860, 1272
 Bloomfield, R. 715, 749
 Blount, S. 401, 434, 438, 440
 Blume, A. 784, 1120
 Blume, L. E. 1159, 1164, 1174, 1182, 1183, 1186
 Bochet, O. 347, 376
 Bocskocsky, A. 1343, 1355
 Bodner, R. 1289
 Boebel, R. B. 715, 742, 895
 Bohnet, I. 369, 370, 483, 556
 Bolton, G. E. 12, 25, 35, 36, 339, 368, 370, 387, 399–402, 410–1, 416–7, 430, 433–4, 436–40, 441–2, 443, 446, 718, 719, 799, 800–1, 811, 1278, 1670
 Bonini, N. 1344, 1363
 Bordalo, P. 114, 184, 186
 Börgers, T. 1089, 1102, 1348, 1410, 1415
 Borges, B. 1348, 1424, 1427
 Bornstein, B. H. 57, 1373, 1449
 Bosch–Domènech, A. 43, 175, 716, 761–2, 876, 899–900, 1093, 1115
 Bowles, S. 34, 342–3, 386, 392, 528, 548, 1040, 1064, 1065–6, 1073, 1075, 1076, 1077, 1079, 1673
 Boyd, R. 1065–6, 1078, 1084, 1673
 Boyer, B. 263, 1494
 Brandenburger, A. 879, 962, 976, 981, 1707
 Brandts, J. 363, 374, 440, 441, 717, 721, 774–5, 783–5, 819–22, 845, 1719
 Brase, G. L. 1345, 1377, 1378
 Brav, A. 1488, 1534, 1541
 Breitmöser, Y. 41, 900
 Brocas, I. 645, 677–8, 918, 998, 1289, 1304, 1306–7
 Broseta, B. 765, 1124, 1138
 Brosnan, S. F. 215, 237–8
 Brown, G. 1093, 1107
 Brown, J. N. 58, 714, 739–42, 747–48, 750–1
 Brown, M. 36, 480, 510–11, 518–22
 Bruhin, A. 29, 113, 139, 177–8, 180
 runner, C. 1096, 1136
 Buchanan, J. M. 1439, 1583, 1623
 Budesu, D. V. 45, 300, 715, 743, 744, 1343, 1352–3
 Burchardi, K. B. 878, 900, 924
 Burke, M. A. 1159, 1184–5
 Burks, S. 301, 386, 483, 557
 Busemeyer, J. R. 207, 299, 1656
 Bush, R. R. 1092, 1098
 Cachon, G. P. 717, 775–6, 778, 1128
 Cadsby, C. B. 468, 721–2, 831–4
 Camerer, C. F. 3, 4, 12–7, 24, 29, 40–4, 51, 103–4, 146, 175–6, 179, 183, 199, 213, 216, 226, 256, 259, 262, 300, 349, 352, 366, 394–5, 715–7, 721, 738, 775–9, 816–7, 822, 830–1, 876, 884, 895, 898–903, 916, 925, 1001, 1037, 1094–5, 1121, 1122–30, 1342–6, 1351, 1355, 1364, 1378–9, 1383, 1454, 1469, 1528, 1580–2, 1593, 1595, 1596, 1598, 1604–5, 1606, 1644–5, 1646, 1647–8, 1649, 1650, 1655, 1658, 1661–3, 1671, 1672, 1689
 Cameron, L. A. 14, 350
 Caplin, A. 47, 922, 1245–7, 1250, 1582, 1604, 1605, 1663
 Cappelen, A. W. 461, 559, 1646, 1673
 Capra, C. M. 875, 892–4
 Cardenas, J. C. 482, 549
 Carillo, J. D. 645, 677, 678, 1289, 1304, 1306–9, 1320
 Carmon, Z. 219, 221
 Carpenter, J. P. 13, 348, 349–50, 375, 386, 389, 411, 483, 555, 559
 Cason, T. N. 722, 779–80, 843–8, 1617
 Casscells, W. 46, 1376, 1377
 Chabris, C. F. 596, 603, 1649
 Chakraborti, S. 755, 756
 Chakravarty, S. 300
 Chamberlin, E. H. 48, 49, 723, 724, 864–8, 871, 873
 Chang, L. J. 1647, 1685
 Chapman, G. B. 57, 175, 207, 221, 603, 619, 1373, 1449
 Charness, G. 36, 181, 346, 359, 360, 363, 367, 396, 401, 411, 431–3, 438, 440, 444–7, 448, 450, 454, 465, 469, 470, 483, 502, 560–1, 879, 913, 928, 952, 954–6, 996, 1000, 1344, 1362, 1363, 1684
 Chateaufneuf, A. 29, 112, 125, 148, 150, 292, 298
 Chatterjee, K. 719, 793, 807, 811
 Chaudhuri, A. 347, 375–6, 717, 778–81
 Cheema, A. 1455, 1461, 1482
 Chen, M. K. 29, 215, 236, 1296
 Chen, Y. 454, 468, 473, 1122
 Cheng, Y. 1498, 1533
 Cherepanov, V. 1583, 1622
 Cherry, T. L. 12, 350, 351, 1453
 Chetty, R. 1350, 1432, 1438–9, 1583, 1625–6, 1630
 Cheung, S. 202, 205
 Cheung, Y.–W. 1093, 1111–2
 Chew, S. H. 84, 92, 97, 101, 103, 148, 200, 202, 207, 288, 294, 298, 304
 Chmura, T. 1095, 1096, 1134–8
 Cho, I.–K. 820, 825, 1348, 1412–3, 1415, 1716, 1721–2
 Choi, J. J. 689, 1104, 1462, 1483, 1594, 1597
 Cinyabuguma, M. 347, 375
 Clark, A. E. 1273, 1278, 1279, 1284
 Cleave, B. L. 13, 348, 386

- Clotfelter, C. T. 45, 1343, 1352–3
 Coble, K. H. 206, 619
 Cohen, J. D. 1290, 1291, 1688
 Cohen, L. J. 1349, 1426, 1433
 Cohen, M. 125, 300, 301
 Cohn, A. 346, 349, 360, 388–9, 469
 Coleman, J. S. 1178, 1286
 Collins, R. 715, 745, 746
 Colman, A. M. 761, 763, 881, 912, 987, 1077, 1415
 Conlisk, J. 103, 104, 108
 Conte, A. 113, 180, 285, 299
 Cook, P. J. 45, 1343, 1352, 1353
 Cooper, D. J. 36, 717, 721, 774–5, 783–5, 822–3, 826–8
 Cooper, R. W. 716, 758–9, 769, 770–1, 852, 913
 Coricelli, G. 1647, 1648, 1690, 1691
 Cosmides, L. 47, 1064, 1288, 1290, 1345, 1377–8
 Costa–Gomes, M. A. 43, 877–8, 900, 903, 916–8, 921–2, 1001, 1128
 Cox, J. C. 250, 401, 431, 438, 449
 Crawford, V. P. 24, 41, 43–4, 216, 260, 262, 340, 765, 814, 877–8, 895–6, 900–2, 904–8, 914–6, 918–20, 921, 996–1002, 1063, 1096, 1124, 1138–42, 1689
 Crockett, M. J. 1648, 1692
 Cropper, M. L. 616, 617
 Crosron, R. T. A. 181, 368, 461, 718, 811, 859, 1343, 1353–4, 1356
 Crupi, V. 1344, 1364
 Cubitt, R. P. 207, 988
 Curley, S. P. 285, 301, 303
 Cutler, D. 55, 1486, 1499, 1520
 Czerlinski, J. 1348, 1424

 Dabora, E. M. 55, 1486, 1498–9
 Dal Bó, P. 375, 521, 785, 786
 Dalton, P. S. 1580, 1583–4, 1586, 1588, 1590, 1613
 Damasio, A. R. 1226, 1231–2, 1368
 Dana, J. 12, 345, 352–3
 Daniel, K. 1498, 1502, 1519, 1540
 Daniel, T. E. 720, 807, 808, 809
 Darai, D. 977, 978
 D'Argembeau, A. 1690
 Darley, J. 1346, 1391
 Dasgupta, P. 627, 644, 1073
 Datta, S. 1347, 1350, 1407, 1409, 1437
 Davis, D. D. 16, 49, 360, 404, 727, 864, 868, 869, 870, 871, 874
 Dawes, C. T. 434, 1684
 Dawes, R. M. 373, 377
 Dawkins, R. 385, 1039, 1066, 1290
 De Bondt, W. F. M. 55, 1486, 1497–8, 1500, 1520
 De Bruyn, A. 719, 799–800
 de Finetti, B. 83, 90, 91
 De Giorgi, E. G. 56, 263, 1494
 De Long, J. B. 55, 1486, 1492, 1502–3
 De Martino, B. 29, 1646, 1663–5, 1666
 de Meza, D. 29, 216, 272, 275, 276–7
 de Neufville, R. 83, 93
 De Neve, J. E. 1278, 1285
 De Quervain, D. J.–F. 54, 351, 1646, 1673–5, 1678
 de Sousa, J. 878, 923
 de Waal, F. B. M. 1297

 Deaton, A. 1229, 1275, 1280–2, 1285
 Decety, J. 1678, 1691
 Deci, E. L. 482, 542
 Dekel, E. 6, 25, 101, 730, 1084, 1228, 1258
 Delgado, M. R. 1663, 1665–6, 1675, 1685
 DellaVigna, S. 39, 481, 646, 684, 688–9, 1256, 1347, 1406, 1433, 1517
 Deneffe, D. 31, 83, 94, 112, 142, 146
 Denzau, A. T. 1347, 1407
 Dhami, S. 24, 29–32, 36, 38, 112–121, 127, 139, 150–1, 191–199, 215–6, 243–7, 251, 254, 400, 403, 412, 415–23, 428–9, 602, 605–12, 620–5, 628–41, 880, 962, 974, 977, 980–4, 1396, 1667–8
 Dhar, R. 221, 1354
 Di Tella, R. 1274, 1275, 1281, 1283
 Diamond, P. A. 39, 645, 649, 664–5, 669, 1462
 Dicks–Mireaux, L. D. L. 1462
 Diecidue, E. 127, 178, 206, 211
 Diener, E. 1277, 1281, 1283, 1285
 Dohmen, T. 181, 346, 364, 365, 509–10, 1401
 Dong, M. 1498, 1533
 Dorn, D. 1501, 1502
 Duesenberry, J. S. 32, 341, 1237, 1277
 Duffy, J. 722, 779, 786, 840, 841–2, 848–51, 872, 1092, 1101
 Duflo, E. 39, 690, 1408
 Dufwenberg, M. 36, 345, 354, 398, 400, 426, 428–30, 436–8, 454, 465, 732, 879, 927–56, 960–1, 1094, 1684
 Dyer, D. 993, 994

 Easterlin, R. A. 1229, 1274–5, 1277–8, 1283
 Eckel, C. C. 35, 181, 346, 368, 402, 403, 869
 Edwards, W. 10, 123, 1378, 1425
 Einhorn, H. J. 123, 146, 175–7, 179, 1349, 1427
 Einstein, A. 7–8
 Eisenman, R. 1344, 1365
 Eliaz, K. 207, 1250, 1608
 Ellingsen, T. 341, 403, 447, 454, 460–1, 527–8, 533, 536, 538–9, 879, 907, 954, 955–6, 958–9
 Ellison, G. 1085, 1160, 1168–9, 1200, 1203–4, 1435
 Ellsberg, D. 3, 235, 286
 Elster, J. 646, 925, 1178, 1223, 1459
 Engelmann, D. 214, 235, 349, 351, 396, 401, 432–3, 447, 521
 Englich, B. 46, 57, 1345, 1372–3, 1449
 Englmaier, F. 479, 490, 495, 510, 1350, 1442
 Epley, N. 1344, 1374–5
 Epper, T. 31, 108, 110, 114–5, 117–8, 122–3, 126–7, 136–7, 155, 176–7, 178–9, 180, 183, 199–200, 202, 204–6, 208–13, 216, 263–4, 606, 619, 627
 Epstein, L. G. 92, 148, 200, 202, 207, 297–8, 1264
 Epstein, S. 1369, 1399
 Erat, S. 453, 455, 461
 Erev, I. 31, 179, 240, 351, 521, 915, 1092, 1093, 1098, 1100–06, 1120, 1127, 1426
 Ergin, H. 297, 301
 Ericson, K. M. M. 40, 112, 154, 170–2, 1346, 1688
 Evans, G. W. 862, 863
 Evans, J. 1377–8, 1399, 1429
 Eyster, E. 972, 974, 995–7, 1001–2

- Fafchamps, M. 97, 594
 Falk, A. 12–15, 36, 131, 348, 351, 358–9, 363, 365, 385–6, 390, 401, 432, 435, 436–9, 446, 481–2, 509–10, 521, 522, 534, 538–9, 541, 542–3, 547, 551–2, 879, 938, 945–7, 949, 1670
 Fama, E. F. 1489, 1495, 1497, 1498, 1500–1, 1502, 1520
 Fang, H. M. 39, 617, 689
 Farber, H. S. 216, 257, 259, 260, 262
 Farhi, E. 1583, 1634
 Farrell, J. 840, 842, 907, 910, 913
 Fedorikhin, A. 1291, 1295
 Fehr, E. 3, 13, 25, 29, 31–6, 49–53, 213, 216, 258–9, 264, 270–2, 339–41, 344, 345–52, 357–60, 363, 366–8, 373–5, 378, 382, 386–8, 396–412, 416, 419, 422, 426–9, 432–9, 443–7, 450–2, 462–5, 478–84, 489, 495, 498, 506–8, 512, 522–6, 535, 540, 541–2, 723–4, 800, 836–7, 856–62, 1066, 1182, 1415, 1644–5, 1648–9, 1658, 1662, 1670–4, 1676, 1679–82, 1692–3
 Fehr–Duda, H. 32, 108, 110, 113–5, 117–8, 122–3, 126–7, 136–7, 177, 178–9, 180, 183, 199–200, 202, 204–6, 208–13, 216, 263–4
 Feld, L. P. 411, 483, 559
 Feltoich, N. 722, 777, 840, 841–2, 848–51, 1092, 1101, 1118–20
 Fessler, D. M. T. 12, 345, 352, 387, 1066
 Fey, M. 731, 733
 Figlewski, S. 1491, 1492
 Figner, B. 1647, 1689
 Finkelstein, A. 1350, 1440
 Finucane, M. L. 1344, 1368
 Fischbacher, U. 36, 53, 339–40, 344–5, 347, 349, 351, 373, 377–8, 396, 397, 401, 406, 432, 435–9, 446, 453, 457–60, 521, 879, 938, 945–7, 949, 1084, 1673–4, 1693
 Fischer, I. 45, 1343, 1357–8
 Fischer, S. 256, 723, 856–8, 862, 1494, 1503
 Fischhoff, B. 46, 302, 1344, 1346, 1368, 1373, 1384–6
 Fishburn, P. C. 83, 85, 101, 110, 113, 123, 145, 148, 181, 192, 586, 587, 593, 601
 Fisman, R. 356, 402
 Fleurbaey, M. 1273, 1618
 Föllmi–Heusi, F. 453, 457–60, 1084
 Forsythe, R. 345, 352, 356, 719, 771, 804, 805, 976
 Foster, D. P. 1097, 1153–8, 1200
 Fouraker, L. E. 10, 718, 789
 Fox, C. R. 146, 301–2, 303–4
 Frank, L. R. 1646, 1667
 Frascara, J. 56, 1449
 Frazzini, A. 1354, 1433
 Fréchette, G. R. 13, 348, 447, 785–6
 Frederick, S. 37, 583, 604, 616, 1293, 1342, 1347, 1397–8, 1400
 French, K. R. 301, 1481, 1497–8, 1500–1, 1520
 Frey, B. S. 551, 1272, 1275, 1276, 1279, 1281, 1287
 Friedman, D. 449, 1093, 1111–2
 Friedman, M. 6–7, 8, 10, 122, 816, 1489, 1502
 Frith, C. D. 1678, 1679, 1690
 Froot, K. A. 55, 1486, 1498–9
 Fryer, R. G. 29, 228, 483, 559, 560, 1097, 1148, 1406–7
 Fudenberg, D. 40, 627, 719, 803, 769, 822, 844, 963, 990, 998, 1066–7, 1085, 1087, 1091, 1093, 1108, 1110–1, 1117, 1130, 1154, 1156, 1259, 1288, 1298–9, 1300, 1303, 1473, 1647, 1689, 1703, 1707–8, 1716
 Fuster, A. 112, 154, 170–2, 1346
 Gabaix, X. 1350, 1432, 1435, 1443, 1446–8, 1583, 1634
 Gächter, S. 34, 36, 52–4, 222, 340, 344, 346–8, 351, 359–60, 373–8, 382–5, 434, 436, 481–2, 521, 540–1, 944, 951, 975, 1673
 Gaissmaier, W. 1425, 1429
 Galinsky, A. D. 46, 1345, 1373
 Gallagher, H. L. 1647, 1679, 1690
 Galton, F. 1346, 1395
 Gardner, J. 1276, 1277
 Garvin, S. 823, 826, 828
 Geanakoplos, J. 36, 345, 878, 925–6, 928, 930–1, 933, 960
 Genesove, D. 213, 226
 Gentzkow, M. 1406, 1437
 Georganas, S. 878, 921
 Gervais, S. 1361, 1488, 1529, 1534–5, 1540
 Ghirardato, P. 285, 291, 293, 298–9
 Ghosal, S. 1580, 1583, 1584, 1586, 1588, 1613
 Gibbons, J. D. 755, 756
 Gibbons, R. 527, 810, 812, 1703
 Gibson, R. 453, 457
 Gigerenzer, G. 47–8, 1339, 1344–5, 1348–9, 1362, 1377–8, 1416–9, 1422–30
 Gilboa, I. 6, 7, 13, 21, 90, 114, 187–8, 284, 288–9, 291–2, 297, 299, 300, 925
 Gilovich, T. 56, 1342–4, 1355, 1374–5, 1428, 1449
 Gintis, H. 2, 4, 6, 7, 21, 23, 25, 32, 34, 35, 40, 53, 54, 57, 91, 104, 108, 339, 341, 342–3, 349, 375, 384–6, 388, 402–3, 587, 602, 635, 711–12, 738, 879, 962, 963–4, 965–6, 978, 1038–41, 1048, 1053, 1064, 1065–6, 1067–8, 1072–7, 1079, 1410, 1673, 1707, 1718
 Girotto, V. 1346, 1376, 1378
 Glaeser, E. L. 688, 1599
 Glaser, M. 56, 1449, 1501, 1528
 Glazer, J. 734–5, 736–7
 Glimcher, P. W. 3, 25, 1231, 1644–5, 1647–49, 1662, 1689
 Gneezy, U. 17, 181, 207, 213, 215, 239–40, 346, 359, 361–2, 368, 439, 453, 455–7, 459, 460–1, 482–3, 545, 550–1, 556, 560–1, 732, 952, 1084, 1094
 Gode, D. K. 49, 724, 864, 872–3
 Goel, A. M. 1488, 1527
 Goeree, J. K. 403, 726, 727, 876, 881, 886, 890–2, 984–6, 992, 995, 1001, 1607
 Goerg, S. J. 509, 1135, 1137–8
 Goette, L. F. 213, 216, 258–9, 454, 470–2, 482, 554, 1677
 Goldfarb, A. 43, 877, 916
 Goldstein, D. G. 1348, 1419, 1422–4, 1426, 1596
 Goldstein, W. M. 123, 146, 176, 177, 179
 Gonzalez, M. 1346, 1377–8
 Gonzalez, R. 122, 123, 146
 Gospic, K. 1647, 1683
 Götte, L. 1679, 1680
 Gourville, J. T. 1474
 Graham, J. R. 1501, 1534
 Grant, S. 92, 207, 288
 Grätz, S. 977, 978

- Green, B. S. 1343, 1355
 Green, L. 596, 598, 604, 616, 664
 Grether, D. M. 20–1, 1376
 Griffin, D. 1346, 1379, 1380–2
 Grimm, V. 1097, 1148, 1150, 1151, 1152
 Grinblatt, M. 1354, 1498, 1501–2, 1540,
 Gross, P. 1346, 1391
 Grosskopf, B. 43, 900
 Grossman, P. J. 181, 368
 Grossman, S. J. 264, 492, 508, 826, 1401
 Gruber, J. H. 39, 681, 683, 690, 1242
 Guala, F. 10, 473, 1599
 Gügler, R. 1679, 1680
 Guilbault, R. L. 46, 56, 1346, 1386, 1449
 Gul, F. 40, 84, 103, 104, 114, 154, 178, 182–3, 263,
 297–8, 719, 803, 816, 1228, 1258, 1260–1, 1263–4,
 1581–2, 1599–1600, 1603–6, 1610, 1645
 Gürrer, Ö. 378, 379
 Gurevich, G. 216, 229
 Guryan, J. 45, 1343, 1355, 1356
 Güth, W. 33, 344, 349, 406, 718, 719, 794–5, 799, 811,
 877, 915, 1040, 1080–1
 Guyer, M. J. 984, 985
- Haendel, V. 596, 598
 Haigh, M. S. 97, 215, 227, 240
 Haile, P. A. 876, 890–1
 Haisken–DeNew, J. P. 1284
 Haisley, E. 1579, 1581, 1582, 1595, 1597, 1599–600,
 1601, 1612, 1617
 Haldane, J. B. S. 1073–4, 1075
 Halevy, Y. 108–9, 199, 208, 210, 295–6, 298, 300, 606,
 623, 625, 627, 633, 638
 Haley, K. J. 12, 345, 352, 387, 1066
 Hamilton, W. D. 342, 385, 1039, 1055, 1073, 1074–6
 Hammerstein, P. 1038, 1043–4, 1049, 1063, 1066, 1077
 Hands, D. W. 5, 6
 Hannan, R. L. 359, 386
 Hare, T. A. 1647, 1662–3, 1689
 Hargreaves Heap, S. P. 44, 369, 393, 395, 455, 473, 895
 Harris, C. 617, 651
 Harris, M. 1488, 1516, 1540
 Harrison, G. W. 108, 180, 351, 994
 Harsanyi, J. C. 57, 738, 767, 984, 1711, 1718
 Hart, O. D. 216, 264–5, 268, 270, 272, 278, 492, 508,
 519, 1401
 Hart, S. 1097, 1155, 1156
 Haruvy, E. 1096, 1142, 1143, 1145, 1147–8
 Heath, C. 56, 216, 251, 254, 301, 303, 1354, 1449, 1474
 Heaton, J. B. 1529, 1535
 Heckman, J. J. 12, 14, 15
 Helbing, 1039, 1068, 1072, 1178
 Helliwell, J. F. 1273, 1286
 Hennig–Schmidt, H. 944, 951
 Henrich, J. 34, 50, 345, 350–1, 463, 550, 556, 1072
 Herrmann, B. 53, 347–8, 374, 382–5, 411, 483, 558
 Herrnstein, R. J. 595, 596
 Hertwig, R. 16, 17, 1344, 1362, 1364
 Herweg, F. 112, 160, 161, 164, 216, 276–9, 281, 283
 Heukamp, F. H. 199, 619
 Hey, J. D. 30–1, 113, 179–80, 183, 285, 288, 298–300,
 307–8
- Hirsch, M. W. 1186, 1190, 1191
 Hirshleifer, D. A. 1433, 1508, 1517, 1528
 Ho, T.–H. 41, 103, 146, 183, 876–8, 901, 1094–5, 1120–2,
 1125, 1127–31, 1133
 Hofbauer, J. 1038, 1041, 1048, 1052, 1091
 Hoffman, E. 12, 350–1, 387, 482, 549, 801
 Hoffrage, U. 47, 1345, 1349, 1377–8, 1427
 Hogarth, R. M. 12, 16, 17, 129, 175, 301, 1349, 1427
 Holland, G. 214, 235
 Hollerman, J. R. 1675
 Holmstrom, B. 268, 481
 Holt, C. A. 16, 49, 108, 175, 360, 404, 721, 726–7, 819,
 820, 821–2, 864, 868–71, 874, 881, 984–6, 994, 1607
 Hong, H. 1487, 1493, 1517, 1521, 1526, 1528
 Hopkins, E. 414, 1092, 1102–3, 1121
 Horowitz, J. K. 214, 220
 Hossain, T. 1350, 1431
 Hotz, V. J. 1350, 1439
 Houser, D. 374, 559, 1676
 Houthakker, H. S. 1579
 Hsee, C. K. 178, 603, 1235
 Hsu, M. 1646, 1667–8, 1683
 Huang, M. 263, 1455, 1480, 1494
 Huberman, G. 1501, 1516
 Huck, S. 50, 369–70, 449, 968–9, 1093, 1102, 1113–4
 Huettel, S. A. 1655, 1658
 Huffman, D. 259, 365, 689
 Hurwicz, L. 285, 293, 872
- Ikenberry, D. 1498, 1534
 Ippolito, R. A. 55, 1495
 Iriberry, N. 41, 44, 356, 432, 877, 895, 998, 1001–2
 Isoni, A. 214, 230, 720, 813–15
 Israel, S. 1648, 1694
 Itoh, H. 478, 483, 485, 487, 490
 Izuma, K. 1646, 1669, 1683
- Jackson, M. O. 1097, 1148, 1182, 1406–7
 Jegadeesh, N. 55, 1486, 1498, 1520
 Jehiel, P. 966–9, 973, 1156
 Jenison, R. L. 1646, 1666
 Johannesson, M. 341, 403, 447, 454, 460–1, 482,
 527–8, 533, 536, 538–9, 553, 955
 Johnson, E. J. 108, 214, 219, 221, 230, 234, 721, 816–8,
 917, 1373, 1455, 1479, 1596, 1607
 Johnson, M. D. 346, 367
 Jonker, L. 1038, 1048, 1052
 Jovanovic, B. 1350, 1431
 Judd, K. L. 25, 651, 872
 Jullien, B. 148, 216, 263
 Jung, Y. J. 828, 829, 830–1
- Kable, J. W. 1647, 1662, 1689
 Kachelmeier, S. J. 175, 176
 Kagel, J. H. 36, 49, 104, 363, 435, 823, 826, 828, 864,
 868, 874, 992, 993, 994, 995, 997, 1001
 Kahneman, D. 4, 10, 19, 25, 26–7, 28–33, 41, 44–5, 46,
 47–8, 49, 54, 56–7, 80, 84, 95–6, 98, 107, 111, 115–6,
 122, 129–33, 136, 142–3, 145–6, 148, 151, 153, 173, 176,
 179, 181, 185, 190–3, 196, 199, 207, 213–4, 217–20, 222,
 224–6, 232–3, 247, 340, 603, 607, 624, 632, 634, 680,
 711, 776, 877, 881, 915, 976, 1228, 1271, 1296, 1339,

- Kahneman, D. (*cont.*)
 1342, 1343–4, 1345–7, 1349, 1351, 1357, 1361–5, 1368,
 1370–1, 1373–6, 1385–7, 1396–1401, 1410, 1416–17,
 1426–9, 1430–1, 1449, 1453–7, 1471–2, 1474, 1494,
 1596, 1600–1, 1608, 1644, 1667
- Kalai, E. 788, 1097, 1156–7
- Kamenica, E. 1581, 1602–3
- Kandori, M. 1158–9, 1162, 1164, 1167, 1200
- Karandikar, R. 1102, 1348, 1415
- Karlan, D. 349, 393, 1437
- Karlin, S. 1196, 1197
- Karni, E. 91, 108
- Kartik, N. 457, 460
- Katok, E. 441, 714, 734, 736
- Kaustia, M. 1354, 1355
- Kearney, M. S. 45, 1343, 1355, 1356
- Keloharju, M. 1354, 1501
- Keynes, J. M. 287, 896
- Khalmetski, K. 879, 954, 956–60
- King, M. A. 1425, 1462
- Kinney, W. R., Jr. 1498
- Kirby, K. N. 596, 598, 604, 616
- Kirchler, E. 56, 1449
- Kirchsteiger, G. 36, 50, 398, 436, 437, 438, 879, 933,
 937–8, 941, 943, 945–6
- Klein, P. G. 1399, 1534
- Klibanoff, P. 285, 292, 295–8
- Kliger, D. 216, 229
- Knetsch, J. L. 33, 219, 229, 1472
- Knez, M. 778, 779
- Knoch, D. 351, 1647, 1663, 1678, 1680–2
- Knutson, B. 1646, 1665, 1668, 1675
- Köbberling, V. 112, 140, 144, 145, 148, 150
- Kocher, M. G. 374, 1128
- Koehler, D. J. 301, 302, 601
- Koenigs, M. 1647, 1682
- Koessler, F. 969, 973
- Korniotis, G. M. 1501
- Koscik, T. R. 1647, 1683
- Kosfeld, M. 393, 482, 534, 538, 539, 542–3, 547, 559,
 1648, 1692, 1693
- Kosse, F. 618–19
- Köszegi, B. 30, 39, 112, 131, 153–9, 161–6, 168–72, 183,
 214, 220, 233, 257, 260, 263, 273, 276–8, 279–80, 645,
 664–5, 669, 681, 690, 956, 960, 1242, 1582, 1605,
 1609, 1610
- Kothiyal, A. 31, 148, 179, 286, 288, 299–300, 307
- Kotlikoff, L. J. 649, 658
- Krajbich, I. 1647, 1676, 1683
- Kranton, R. E. 61, 455, 465–6, 473, 474
- Kreps, D. M. 85, 202, 207, 525, 721, 733, 820, 825, 828,
 830, 831, 1093, 1108, 1110, 1246–7, 1258, 1260, 1605,
 1608, 1611, 1703, 1716, 1721–22
- Krishna, V. 988, 990, 998
- Krueger, J. I. 974, 977–8, 1396
- Kube, S. 346, 359, 361, 363, 364
- Kübler, D. 370, 890
- Kuhn, H. W. 274, 932
- Kuhn, T. S. 5, 6
- Kuhnen, C. M. 1665
- Kumar, A. 1354, 1501
- Kunreuther, H. 116, 129, 191–4, 196, 301, 1296
- Kuo, W. J. 1648, 1691
- Kurtz, K. J. 114, 188
- La Porta, R. 1498, 1520
- Lacetera, N. 482, 553, 1350, 1440, 1442–3, 1449
- Laibson, D. I. 38, 596, 606, 612, 615, 617, 632, 637, 645,
 651, 658, 660, 663–4, 1226, 1236, 1240, 1242, 1264,
 1350, 1432, 1435, 1633, 1647, 1686, 1688
- Lakatos, I. 5–6
- Lakonishok, J. 55, 1486, 1495, 1497–8, 1520
- Lakshminarayanan, V. 215, 236–7, 239
- Lamm, C. 1678, 1691
- Langer, E. J. 1488, 1528
- Larrick, R. P. 1386, 1527
- Laslier, J. F. 1092, 1102
- Latham, G. P. 29, 251, 1455, 1482
- Lavy, V. 483, 554, 556, 560
- Layard, R. 1275, 1281, 1600
- Lazear, E. P. 1527
- Leahy, J. 1245, 1246–7, 1250
- LeBoeuf, R. A. 1370, 1373–4
- Ledyard, J. O. 373, 377, 448, 550, 855
- Lee, C. M. C. 56, 1492–3, 1520
- Legg, S. 56, 263, 1494
- Lehrer, E. 291, 1097, 1156, 1157
- Leibbrandt, A. 349, 390–1
- Leider, S. 479, 490, 495
- Lemieux, T. 478, 494
- Lepper, M. R. 1616
- Levin, D. 992–6, 1000–1
- Levine, D. K. 40, 401, 445, 447, 448, 450, 535, 627,
 1085, 1087, 1091, 1108, 1110–1, 1117, 1130, 1154, 1156,
 1259, 1288, 1298–9, 1300, 1303, 1647, 1689
- Levitt, S. D. 10–11, 13, 51, 58, 194, 348, 352, 386–7, 394,
 483, 560, 714, 732–3, 754–7
- Levy, O. 216, 229
- Lewis, K. P. 1481, 1683
- Li, S. X. 454, 468, 473
- Lichtenstein, S. 20, 45, 84, 107, 108, 1344, 1365–6,
 1369, 1373
- Lin, A. 1646, 1669–71
- Lindgren, B. 1289, 1311
- Linnainma, J. T. 1501
- Lipman, B. L. 6, 1260, 1263–4
- Lipsey, R. G. 8
- List, J. A. 10–11, 13, 15, 29, 50–2, 97, 172, 214, 215, 222,
 225–7, 232, 240, 346, 348, 350–2, 359, 361–2, 386–7,
 394, 439, 756–7, 835, 871, 994
- Locke, E. A. 29, 251, 1455, 1482
- Loewenstein, G. F. 2, 3, 33, 37–8, 45–6, 217, 221, 264,
 340, 446, 544–5, 583, 587, 595–9, 602–3, 605–6,
 607–9, 612, 616, 632, 637, 720, 813, 975, 1223, 1226–7,
 1229–30, 1233–5, 1242–3, 125–2, 1254–5, 1265, 1288,
 1290–1, 1293, 1295–7, 1344, 1368, 1454–5, 1463–6,
 1468, 1473, 1475, 1481, 1579–82, 1595, 1597, 1599–600,
 1601, 1612, 1617, 1663, 1686, 1688–9
- Loomes, G. 97, 101, 108, 113, 154, 175, 178, 181–2,
 225, 1245
- Lopes, L. L. 1349, 1427
- Lord, C. G. 47, 1346, 1389, 1391
- Lott, J. R., Jr. 993
- Loughran, T. 1498, 1533, 1534

- Lovallo, D. 207, 1528
 Lucas, R. E., Jr. 858, 862
 Luce, R. D. 148, 150–1, 189, 285, 1092, 1098, 1667
 Lundborg, P. 1289, 1311
 Lusk, J. L. 206, 619

 Maccheroni, F. 292–3, 298
 MacCrimmon, K. R. 84, 97, 101
 MacCulloch, R. J. 1274, 1275
 Machina, M. J. 92, 284, 288, 290, 296, 299
 Maciejovsky, B. 56, 1449
 Madrian, B. C. 1594, 1597
 Majluf, N. S. 721, 831, 1488, 1530, 1540
 Mak, V. 40, 1131
 Malkiel, B. 1500, 1534
 Malmendier, U. 39, 56, 646, 684, 688, 1256, 1357, 1488, 1528
 Malouf, M. W. K. 718, 788–90, 795
 Maniar, S. D. 1386
 Mankiw, N. G. 241, 863
 Manzini, P. 601, 606, 628, 633, 638, 1583, 1618, 1623
 Mareno, M. C. 238
 Margolis, H. 646, 1427, 1459
 Mariotti, M. 601, 606, 628, 633, 638, 1583, 1618, 1623
 Mariotti, T. 1289, 1308, 1309, 1320
 Markle, A. B. 113, 175
 Markowitz, H. 3, 175
 Marley, A. A. J. 148, 189
 Mas, A. 15, 213, 359
 Masatlioglu, Y. 601, 602, 633, 635, 641–3, 1582, 1618, 1619, 1621–3
 Maschler, M. 769, 932, 1703
 Masclot, D. 347, 375, 558
 Mas–Collel, A. 29, 83, 85, 89, 139, 413, 863, 1097, 1155–56, 1579
 Maskin, E. S. 627, 644, 963, 1073, 1077, 1348
 Matsui, A. 1348, 1412–13, 1415
 Matsusaka, J. G. 417, 1534
 Matsushima, H. 734, 735, 736
 Mayer, C. 213, 226
 Mayer, T. 587, 603
 Maynard Smith, J. 1039, 1041, 1043, 1046, 1073
 Mazar, N. 459, 545
 McAllister, P. H. 1093, 1103
 McAlvanah, P. 599, 601
 McCabe, K. A. 351, 358, 534–5, 549, 744, 1647, 1676, 1690
 McClure, S. M. 1647, 1663, 1687, 1688, 1689
 McConnell, K. E. 214, 220
 McCord, M. 83, 93
 McFadden, D. L. 1193, 1348
 McKelvey, R. D. 713, 729–1, 733, 875, 882, 886–8, 1116–7, 1132
 McNeil, B. J. 57, 1449, 1453, 1456
 Meehl, P. E. 56, 1375, 1449
 Mehra, R. 241, 1499
 Mehta, J. 716, 759, 760
 Meier, S. 349, 394, 556, 1286
 Mellström, C. 482, 553
 Meng, J. 216, 260, 262
 Mengel, F. 1068, 1097, 1148, 1150–2, 1406
 Metzger, M. 1343, 1353
 Meulbroek, L. K. 1528
 Miao, B. 108, 202, 205
 Michaely, R. 56, 1486, 1498
 Milgrom, P. 481, 525, 721, 733, 822, 825, 827–8, 852, 988, 990, 1350, 1401, 1431
 Milkman, K. L. 1460, 1463
 Miller, J. B. 1343, 1355
 Miller, J. H. 345, 351, 356, 399, 431, 433, 1067, 1130
 Miller, R. M. 873, 874
 Mischel, W. 597, 1291, 1293, 1295
 Mislin, A. 346, 367
 Mitchell, J. P. 1684, 1690, 1691
 Mitchell, O. S. 1481, 1483
 Mitchell, S. H. 604
 Miyasawa, K. 1093, 1109
 Moinas, S. 971, 972
 Monderer, D. 780, 1093, 1109
 Mookherjee, D. 58, 714, 741–3, 1093–4, 1106, 1115–6, 1120, 1123, 1127, 1128
 Moore, J. 216, 264, 265, 268, 270, 272, 278
 Moreno, D. 722, 845
 Morgan, J. 1125, 1350, 1431
 Morgenstern, O. 26, 83, 86, 296, 297
 Moskowitz, T. J. 1450, 1498
 Mosteller, F. 1092, 1098
 Mullainathan, S. 39, 683, 1347, 1350, 1401, 1407, 1409, 1437, 1488, 1527, 1583, 1602, 1632, 1634
 Munro, G. D. 1346, 1391
 Murnighan, J. K. 718, 792
 Murphy, K. M. 1226, 1230, 1236–7, 1241–2
 Mussweiler, T. 46, 57, 1345, 1370–2, 1373, 1449
 Myatt, D. P. 1159, 1180, 1181
 Myers, S. C. 721, 831, 1488, 1530, 1540
 Myerson, R. B. 720, 804–6, 843, 990, 998, 1703

 Nagel, R. 41, 43, 731, 734, 876, 897–900, 924, 952, 1120, 1647–8, 1690, 1691
 Nakamura, N. 291
 Nash, J. F. 57, 709, 714, 718, 738, 786–7, 1708
 Neale, M. A. 46, 57, 1345, 1371, 1373, 1449
 Neelin, J. 718, 719, 795, 798, 809
 Neilson, W. S. 35, 104, 144, 250, 401, 451, 452
 Nickerson, R. S. 1346, 1389
 Noor, J. 592, 1258
 Nordhaus, W. 8–9
 North, D. C. 1347, 1407
 Northcraft, G. B. 46, 57, 1345, 1371, 1373, 1449
 Novemsky, N. 29, 214, 222, 224, 232
 Nowak, M. A. 1039, 1040, 1073–4, 1076–9, 1084
 Nunn, N. 1347, 1409
 Nyarko, Y. 1093, 1112–3
 Nydegger, R. V. 718, 788

 Oberholzer–Gee, F. 551
 Ochs, J. 715, 718, 749, 779, 786, 795–99, 831, 886
 Ockenfels, A. 35, 36, 339, 352–3, 400–2, 410–1, 416–7, 430, 433–4, 436–7, 440, 441–2, 443, 446, 466, 799, 1136, 1278, 1670
 Odean, T. 29, 45, 56, 213, 1343, 1354, 1361, 1455, 1477, 1486, 1493, 1496, 1501, 1517, 1535, 1540
 O'Doherty, J. P. 1668, 1671–2

- O'Donoghue, T. 39, 645–6, 650–1, 655, 670–1, 675–7, 679, 683, 1265, 1288, 1289–90, 1291, 1296–7, 1633
- Oechssler, J. 1151, 1347, 1400–1
- Oettinger, G. S. 258, 259
- Offerman, T. 435, 438, 995
- Ok, E. A. 601–2, 633, 635, 641–3, 1081, 1084
- Öncüler, A. 178, 206, 207, 221
- O'Neill, B. 58, 714, 738, 739–43, 745, 749–51, 754, 758, 895
- Oppenheimer, D. M. 1349, 1423
- Orme, C. 30–1, 179, 180, 183
- Ortmann, A. 16, 17, 18, 784
- Osborne, M. J. 1703
- Ostrom, E. 378, 1080
- Oswald, A. J. 1273, 1275–8, 1284–5
- Owen, G. 718, 788
- Pachur, T. 1344, 1366, 1367
- Padoa-Schioppa, C. 1662
- Page, T. 54, 347, 370, 375
- Paichayontvijit, T. 347, 375
- Palacios-Huerta, I. 57–8, 249, 714–5, 731, 732, 753, 754, 755, 756, 757
- Palfrey, T. R. 713, 729, 730, 731, 733, 843, 875, 882, 886–8, 1116–7, 1132, 1519
- Pape, A. D. 114, 188
- Paulus, M. P. 1646, 1667
- Pauly, M. 192, 193
- Payne, J. W. 816, 1423, 1425
- Pearce, D. G. 914, 926, 931, 986, 1706
- Penczynski, 878, 900, 924
- Perry, M. 734, 735, 736–7, 826
- Persson, T. 417, 418
- Pesendorfer, W. 40, 298, 816, 1228, 1258, 1260–1, 1263–4, 1581–2, 1599–1600, 1603–6, 1610, 1645
- Petrie, R. 346, 368
- Pfeiffer, F. 618–19
- Phelps, E. S. 3, 38, 606, 612, 615, 632, 637, 1264, 1348
- Pinker, S. 1064, 1345, 1377
- Pinto, J. L. 112, 145–7
- Plassmann, H. 1662, 1663, 1665
- Plott, C. R. 10, 20–1, 214, 226, 229–31, 864, 871, 993, 1617
- Plous, S. 1373, 1391
- Pogarsky, G. 46, 1345, 1372–3
- Pogrebna, G. 217, 240
- Polak, B. 976, 1707
- Pollak, R. A. 3, 38, 606, 612, 615, 632, 637, 654, 1237, 1264
- Pollet, J. M. 1433, 1517
- Pope, D. 29, 214, 227–8
- Popper, K. 5, 6
- Porter, M. E. 1534
- Porteus, E. L. 202, 207, 1246–7, 1258, 1605, 1608
- Posch, M. 1092, 1102, 1103
- Poterba, J. M. 301, 1481
- Potters, J. 207, 213, 215, 239, 240, 722–3, 835, 836, 838–9, 853–6
- Pouget, S. 971, 972
- Poulsen, A. 814, 815
- Powdthavee, N. 1285, 1343, 1356, 1361, 1449
- Prelec, D. 27, 29, 37–8, 103, 110, 118–19, 121, 134, 146, 148, 150–1, 184, 587, 595–7, 603, 605–6, 607–9, 612, 616, 624, 632, 637, 1289, 1454, 1463–4, 1466, 1468, 1473, 1475, 1667
- Prescott, E. C. 241, 1499
- Price, G. R. 1039, 1043, 1046, 1075
- Prowse, V. 43, 183
- Pruckner, G. J. 453, 458, 555
- Pycia, M. 1228, 1261, 1264
- Quattrone, G. A. 975, 1373
- Quiggin, J. 28, 110, 124, 128, 148
- Rabin, M. 30, 36, 39, 45, 47–8, 84, 104–6, 112, 131, 153–9, 161–6, 168–9, 170–2, 183, 214–5, 220, 233, 248, 250, 257, 260, 263, 273, 277–80, 401, 411, 431–3, 444–8, 450, 469, 502, 645–6, 650–1, 655, 670–1, 675–7, 679, 683, 840, 842, 878, 907, 937, 941, 946, 956, 960, 972, 974, 995–7, 1001–2, 1289, 1343, 1346, 1352–3, 1358, 1389, 1391–2, 1430, 1471, 1582, 1605, 1609–10, 1633, 1670
- Rachlinski, J. J. 57, 1346, 1386–7
- Radner, R. 720, 806–11
- Raiffa, H. 285, 791
- Rajagopal, P. 1454, 1476
- Rangel, A. 1227, 1242, 1257, 1292, 1354, 1582, 1595, 1601, 1604, 1611–12, 1614–8, 1621, 1645, 1662, 1671–2
- Rapoport, A. 45, 715, 718–9, 730–1, 734, 742–4, 747–8, 788, 803, 807–9, 895, 915, 976–7, 984–5, 1093, 1106, 1125, 1127, 1131–2, 1343, 1352–3
- Raviv, A. 1488, 1516, 1540
- Read, D. 17, 587, 600–1, 604, 606, 620–1, 622, 629, 631–3, 636–7, 1227, 1251, 1454–5, 1469, 1471, 1481
- Redelmeier, D. A. 57, 302, 1271, 1449
- Reeson, A. F. 483, 556–7
- Reich, P. 482, 554
- Reiley, D. H. 756–7
- Rey-Biel, P. 356, 432, 490
- Rha, J.-Y. 1454, 1476
- Richerson, P. J. 1065, 1066
- Riedl, A. 50, 375, 447
- Riyanto, Y. E. 1343, 1356, 1361, 1449
- Robbins, J. M. 974, 1396
- Roberts, J. 525, 721, 733, 822, 825, 827–8, 852
- Robinson, J. 1093, 1107, 1109, 1452
- Robson, A. J. 1040, 1073, 1080–1, 1167
- Rockenbach, B. 359, 379, 460, 482, 541, 542, 1480
- Rode, M. 1286–7
- Roelofsma, P. H. M. P. 587, 600–1, 619, 629, 632
- Rogers, B. W. 880, 969, 971
- Roiser, J. P. 29, 1646, 1665
- Roll, R. 1488, 1491, 1528, 1535
- Romer, D. 856, 860
- Rosen, A. 56, 1375, 1449
- Rosenthal, R. W. 58, 714, 724, 726, 729, 734, 739–42, 744, 747–8, 750–1, 754
- Ross, L. 974, 1396
- Ross, S. 1488, 1540
- Rotemberg, J. J. 36, 448
- Roth, A. E. 10, 14, 16, 25, 49, 339, 344, 349, 350–1, 405, 521, 718, 786, 788–90, 792–3, 795–801, 813, 864, 868, 874, 1092–3, 1098, 1100–1, 1103–5, 1120, 1127, 1426

- Roth, B. 57, 1397
 Roth, J. 1488, 1528
 Rothschild, M. 400, 412
 Rottenstreich, Y. 178, 302
 Rubinstein, A. 6, 107, 250, 586, 587, 593, 601, 606, 629–30, 718–9, 721, 786, 794–5, 816, 923–4, 1416, 1582, 1618, 1703
 Russell, T. 587, 603
 Rustagi, D. 349, 392–3
 Rustichini, A. 17, 482, 550, 551, 556, 1102
 Ryan, R. M. 482, 542
- Safra, Z. 108, 250
 Sagi, J. S. 92, 288, 298, 304
 Saito, K. 606, 625, 626–7
 Salanié, B. 216, 263
 Salant, Y. 1582, 1618
 Samuelson, L. 250, 826, 1040–1, 1048, 1073, 1080–1, 1090–1, 1097, 1118, 1143, 1148
 Samuelson, P. A. 3, 8–9, 37, 96, 586, 637, 1470, 1579
 Samuelson, W. F. 219, 719, 807, 811, 990–1, 993, 1095, 1134, 1596
 Sandholm, W. H. 1041, 1167
 Sandmo, A. 215, 243
 Sanfey, A. G. 1647, 1676, 1678, 1680, 1684, 1685
 Sanjurjo, A. 1343, 1355
 Sarin, R. 92, 291, 1089, 1092, 1101–2, 1123, 1348, 1415
 Satterthwaite, M. A. 720, 843
 Sausgruber, R. 36, 416–7, 453, 458, 555
 Savage, L. J. 83, 91, 96, 122
 Scale, D. A. 807, 809
 Schelling, T. C. 646, 716, 720, 758, 759, 813–4, 877, 904–5, 1459
 Schmeidler, D. 92, 114, 128–9, 148, 187–8, 285, 288–92, 297–9, 925
 Schmidt, K. M. 32, 34–6, 216, 279, 281, 283, 339, 352, 373, 378, 399, 401–11, 412, 416, 419, 422, 426, 428–9, 432–4, 436–7, 443–4, 446–7, 450–2, 462–3, 465, 478–81, 484, 489, 498, 507, 512, 524–6, 535, 541, 542, 800, 836–7, 1182, 1670, 1672
 Schmidt, U. 108, 112, 131, 145, 148, 150, 152–3, 161, 199
 Schneider, F. 352, 387, 1682
 Schneller, K. 46, 1345, 1371–2
 Scholten, M. 601, 606, 620–1, 631–3, 636–7
 Scholz, J. K. 1350, 1439
 Schotter, A. 8, 207, 482, 548–9, 713, 717, 720, 727–8, 781–2, 806–8, 809–11, 1093, 1112–3, 1582, 1604–6, 1608, 1713
 Schrag, J. 47, 1346, 1389, 1391–2, 1430
 Schram, A. 349, 374, 395–6
 Schubert, R. 177–8
 Schunk, D. 142, 224, 1682
 Schwarz, N. 1376, 1398
 Schweitzer, M. 29, 214, 227–8
 Seasholes, M. S. 1354, 1517
 Sefton, M. 347, 375, 714, 734, 737, 1125, 1130
 Segal, U. 108, 148, 250, 285, 294–6, 401, 450
 Seinen, I. 349, 395–6
 Seki, E. 13, 348–9, 389
 Selten, R. 3, 10, 42, 47, 352, 767, 830, 877, 895, 899, 915, 984, 992, 1038, 1043, 1044, 1049, 1063, 1077, 1095, 1096, 1114, 1132–8, 1348, 1378, 1409–12, 1416, 1711, 1716
 Sen, S. 219, 817
 Senik, C. 1272, 1274, 1278–9
 Seymour, B. 1664, 1678
 Shachat, J. M. 58, 714–5, 744–5, 749, 750–1
 Shafrir, E. 57, 859, 1350, 1370, 1373–4, 1437, 1449, 1454, 1475
 Shah, A. K. 1350, 1437
 Shaked, A. 746, 795, 802
 Shapira, Z. 1354, 1528
 Shapiro, C. 513, 1436
 Shapley, L. S. 1093, 1109–10
 Sharma, T. 722, 843–8
 Shearer, B. S. 346, 362–3
 Shefrin, H. M. 40, 216–7, 644, 646–7, 649, 1264, 1298, 1354, 1453, 1459–63, 1477, 1481, 1647, 1689
 Shehata, M. 175–6
 Sherstyuck, K. 715, 745–6
 Shiffman, S. 1291
 Shiffrin, R. M. 1290, 1399
 Shiller, R. J. 55, 972, 1486, 1496–7, 1500
 Shiv, B. 1291, 1295
 Shleifer, A. 55, 184, 1347, 1401–2, 1405, 1486–8, 1491, 1493, 1495–6, 1500, 1512, 1514–6, 1529, 1533
 Sicherman, N. 587, 602
 Siegel, S. 10, 718, 789
 Sigmund, K. 1038, 1040–1, 1048, 1077–8
 Signorino, C. S. 888–9
 Silverman, D. 39, 617, 689
 Simon, C. P. 1186
 Simon, H. A. 3, 10, 47, 521, 1101, 1339, 1347–8, 1409–12, 1416, 1419, 1423, 1425
 Simonson, I. 217, 1455, 1481
 Sims, C. A. 1350, 1448
 Singer, T. 1675–6
 Siniscalchi, M. 25, 284, 290, 293, 296–9, 730, 936
 Skinner, J. 243, 1462
 Slemrod, J. 242–3
 Sloman, S. A. 1345, 1349, 1377–8, 1399
 Slonim, R. L. 13–4, 58, 344, 350, 741
 Slovic, P. 20, 45, 46, 56, 84, 107–8, 1234, 1344, 1367–9, 1373, 1449, 1454, 1472, 1483
 Smale, S. 1186, 1190–1
 Smith, A. 3, 172, 1290
 Smith, D. V. 1227, 1250, 1646, 1669
 Smith, V. L. 10, 13, 16, 48–9, 360, 549, 723, 724, 864, 868–71, 873, 890, 972
 Snyder, M. 1374, 1399
 Sobel, J. 401, 431, 450, 822, 826, 907
 Soetevent, A. R. 354, 835
 Sokol–Hessner, P. 1666–7
 Solan, L. 57, 1397, 1449
 Soman, D. 1454–5, 1461, 1474–5, 1482
 Sonnemans, J. 377, 481
 Sopher, B. 58, 603, 714, 717, 741–3, 781–2, 1093–4, 1106, 1115–6, 1120, 1123, 1127–8
 Spiegel, R. 277, 1602
 Sprenger, C. 115, 199–202, 204–5
 Stacchetti, E. 926, 931
 Stähl, D. O. 41, 718, 786, 794–5, 876, 921, 1096, 1124–5, 1142–3, 1145, 1147–8

- Stalmeier, P. F. 93
 Stanovich, K. E. 44, 47–8, 1345–7, 1349, 1376, 1389, 1398–9, 1401, 1427–8, 1430–1
 Starmer, C. 18, 26, 28, 95, 101, 103, 114, 116, 130, 133, 173, 182
 Statman, M. 216–7, 1354, 1469, 1477, 1481
 Stein, J. C. 1487, 1493, 1495, 1517, 1534
 Stevens, J. R. 1295, 1686
 Stevenson, B. 1229, 1282–3
 Stigler, G. J. 1401, 1411
 Stiglitz, J. E. 400, 412–3, 1579
 Stöcker, R. 1095, 1132–4
 Stolarz–Fantino, S. 1344, 1364
 Stott, H. P. 122, 134, 1667
 Strack, F. 46, 1345, 1347, 1370–1, 1373, 1376, 1398
 Strahilevitz, M. A. 217, 221
 Straub, P. G. 767, 770
 Strobel, M. 401, 432–3, 447
 Strotz, R. H. 37, 586, 590, 1263, 1264
 Stutzer, A. 482, 554, 1272, 1275–6, 1278, 1281, 1284, 1286–7
 Suetens, S. 723, 853–6
 Sugden, R. 3, 28, 101, 103, 108, 112–3, 130, 133, 153–4, 173, 178, 181, 182, 763, 814, 815, 962, 988, 1178, 1245, 1581, 1583, 1586, 1591–3, 1599, 1612, 1617, 1623, 1624–5
 Sundali, J. A. 877, 915, 1343, 1353–4, 1356
 Sunder, S. 49, 724, 864, 872, 872
 Sunstein, C. R. 1366, 1368, 1409, 1437, 1580–1, 1591–3, 1596, 1599, 1611, 1633
 Suppes, P. 10, 1092, 1098
 Sutter, M. 215, 240–1, 1128
 Sutton, J. 795, 802
 Swarthout, J. T. 744–5
 Szech, N. 551–2

 Tabellini, G. E. 417–8
 Tabibnia, G. 1676–7, 1680
 Tajfel, H. 454, 465–6
 Talmi, D. 1646, 1665
 Tanaka, T. 39, 618, 623, 625
 Tang, F. F. 731, 734, 1120, 1122
 Tate, G. 56, 1357, 1488, 1528
 Taylor, H. M. 1196–7
 Taylor, P. 1038, 1048, 1052
 Tentori, K. 1344, 1362, 1364
 Terrell, D. 45, 1343, 1353
 Tesfatsion, L. 25, 872
 Tetlock, P. E. 56, 1342, 1350, 1386, 1449
 Thakor, A. V. 1488, 1527
 Thaler, R. H. 2, 4, 6, 14, 20, 25–6, 29–30, 33, 33, 37–40, 45, 49, 55–6, 107–8, 144, 199, 207, 213–5, 217, 222, 230, 233, 240–2, 257, 349, 373, 587, 595–6, 597–8, 616, 622, 625, 644, 646–7, 649, 658, 689, 993, 1264, 1271, 1298, 1340–1, 1437, 1451, 1453–7, 1459–63, 1468, 1470, 1473–7, 1479–84, 1486, 1491–3, 1495–500, 1520, 1580–1, 1591–3, 1596, 1599, 1611, 1633, 1647, 1689
 Thomas, J. K. 1486, 1498, 1520
 Thompson, L. 810, 812
 Thöni, C. 347, 374, 376, 384–5, 975
 Tietz, R. 349, 718–9, 795, 799

 Tirole, J. 17, 341, 345, 353, 477, 481, 507, 528, 531–3, 539, 552, 769, 822, 844, 990, 998, 1289, 1311–2, 1316, 1319–21, 1703, 1707–8, 1716
 Tisdell, J. G. T. 483, 556–7
 Titman, S. 55, 1486, 1498, 1502, 1520
 Titmuss, R. M. 341, 482, 533, 553
 Todd, P. M. 1418, 1424
 Tom, S. M. 29, 1646, 1662, 1665–6
 Tooby, J. 47, 1064, 1288, 1290, 1345, 1377–8
 Tranel, D. 1647, 1682–3
 Traub, S. 145
 Trautmann, S. T. 108, 285, 298, 301
 Tricomi, E. 1646, 1671, 1672–3
 Trivers, R. L. 342, 385, 1039, 1066, 1076
 Trope, Y. 1290, 1374
 Tsutsui, K. 814–5
 Turner, J. C. 454, 465, 466
 Tversky, A. 4, 10, 19, 25–32, 44–49, 54–6, 57, 80, 84, 95–6, 98, 107–8, 111, 115–6, 122, 129–33, 136, 142–6, 148, 150–1, 153, 173, 176, 179, 185, 190–3, 196, 199, 213, 222, 247, 298, 300–4, 601, 607, 624, 632, 634, 776, 881, 975–6, 1296, 1339, 1342–7, 1349, 1351, 1355, 1357, 1361–3, 1365, 1370–1, 1373–5, 1379–82, 1396, 1399, 1401, 1416–7, 1425–30, 1449, 1453–7, 1471–2, 1474, 1494, 1596, 1601, 1609, 1644, 1667
 Tyran, J.–R. 36, 369, 411, 416–7, 483, 559, 723, 856–7, 859–62

 Ubel, P. A. 1580, 1601
 Utkus, S. P. 1481, 1483

 Vahid, F. 1092, 1101, 1123
 Valley, K. 720, 810, 812
 van Assen, M. 142
 Van Boening, M. V. 874
 Van Damme, E. 881, 1046
 van de Kuilen, G. 122, 145
 van den Bos, W. 1685–6
 van Dijk, E. 219, 221, 480, 521
 Van Huyck, J. B. 716–7, 763–7, 769, 772–3, 775, 778–81, 783, 845, 902, 1093, 1096, 1103, 1106, 1127, 1139, 1140–1
 van Knippenberg, D. 219, 221
 van Leeuwen, B. 1227, 1251
 Van Praag, B. M. S. 1274, 1280, 1283
 Vanberg, C. 454–5, 1684
 Varey, C. 603, 1271
 Varian, H. 29, 83, 85, 593, 863
 Vayanos, D. 1352–3, 1430, 1499
 Veblen, T. 32, 341, 1277
 Veenhoven, R. 1275, 1283, 1287
 Vega–Redondo, F. 1041, 1048, 1081, 1084, 1087, 1114, 1167
 Vesterlund, L. 15, 431–2, 835
 Viscusi, K. W. 39, 116, 191, 617, 1289
 Vishny, R. 55, 1487–8, 1495, 1512, 1514–16, 1529, 1533
 Vohra, R. 1154–5
 Vohs, K. D. 1291, 1294–5
 Volij, O. 58, 714–5, 731–2, 754–7
 von Gaudecker, H.–M. 179, 207
 von Neumann, J. 26, 83, 86, 296–7, 714, 738, 1711
 von Siemens, F. A. 490

- Voskort, A. 57, 1397
 Vriend, N. J. 716, 761–2, 1093, 1115
 Wagenaar, W. A. 45, 1343, 1352
 Wakker, P. P. 26, 29, 31, 83, 90, 92–5, 101, 103–4, 111–2, 114, 116, 122, 127, 129, 140, 142, 144–6, 148, 150, 182, 191, 198, 213, 215, 249–50, 284, 285–6, 288, 291–2, 294, 297–8, 302, 306
 Walker, M. 57, 715, 752–3
 Wallace, C. 1159, 1161, 1180–1
 Waller, W. S. 97, 103
 Walther, H. 178, 199, 208
 Wason, P. C. 1347, 1397, 1427
 Watson, J. 1130, 1703
 WDR 1342, 1347, 1350, 1407–9, 1437, 1449–50, 1456
 Webb, D. C. 29, 216, 272, 275–7
 Weber, B. J. 175, 207, 619
 Weber, M. 46, 54, 56, 129, 207, 240, 301, 1346, 1387–8, 1390, 1430, 1449, 1493, 1501
 Weber, R. A. 717, 778–9, 886–7
 Weibull, J. W. 1040–1, 1081, 1083–4, 1089, 1091
 Weigelt, K. 366, 727–8, 830–1, 1120
 Weinstein, N. D. 195, 1491, 1528
 Weitz, B. A. 1488, 1527
 Weizsäcker, G. 878, 890, 922, 1471
 Werner, P. 353, 466
 Wertenbroch, K. 221, 689–90, 1477, 1598
 West, R. F. 44, 48, 1345, 1346, 1347, 1349, 1376, 1389, 1398–9, 1401, 1427–8, 1430
 Wiederholt, M. 1448
 Wilcox, N. T. 874, 1125
 Wilson, E. O. 1039, 1064–7, 1075–6, 1288, 1290
 Wilson, P. W. 41, 876, 921, 1145, 1147
 Wilson, R. 525, 721, 733, 798, 828, 830–1, 1716
 Winter, E. 509, 1223, 1232–3
 Wolfers, J. 263, 1229, 1282–4
 Wooders, J. 57–8, 715, 722, 752–4, 758, 845
 Woodford, M. 1448, 1662
 Wout, M. van't 1680
 Wu, G. 113, 122–3, 146, 174–5, 1517
 Wurgler, J. 1462, 1486, 1488, 1500, 1526, 1533, 1541–2
 Xiao, M. 43, 916
 Yaari, M. E. 148, 1040, 1080
 Yacubian, J. 1663, 1666
 Yamagishi, T. 378, 473
 Yang, B. 43, 877, 916
 Yavas, A. 714, 734, 737
 Yellen, J. L. 357, 863
 Yitzhaki, S. 215, 242–3, 245
 Young, H. P. 1097, 1153–4, 1156–60, 1162, 1164, 1167, 1172, 1174–5, 1177–80, 1184–5, 1200–3
 Young, P. H. 1124, 1161, 1180
 Zajonc, R. B. 1344, 1368
 Zak, P. J. 1650, 1684
 Zamir, S. 1124, 1161
 Zank, H. 145, 148, 150
 Zauberman, G. 599, 601, 616
 Zeckhauser, R. J. 219, 222, 1596
 Zeiler, K. 214, 226, 229–31
 Zhang, J. 298, 1090–1
 Zhong, S. 108, 202, 205
 Zhu, N. 1354, 1501
 Zin, S. E. 183, 202
 Zizzo, D. J. 10, 386, 455, 473, 952
 Zuravskaya, K. 1486, 1500
 Zwick, R. 387, 441, 719, 801–2
 Zwiebel, J. 1343, 1355

SUBJECT INDEX

- Abreu–Matsushima mechanism 714, 734, 735–7
- absolute magnitude effect, exponential discounted utility model 598
- acquisition utility 1454–5, 1474, 1476–7
- action-sample learning 1136, 1137, 1138
- adaptive learning 827, 1094
- addiction
 - animal experiments 1649
 - cue-conditioned behavior 1236, 1237, 1240, 1241–2
 - present-biased preferences 646, 647
 - projection bias 1257
 - visceral influences 1230
- additive discount functions 634–5
- advertising 1401, 1402, 1423
- affect heuristic 45–6, 1344, 1366, 1367–9
 - experts 1449
 - neuroeconomics 1665
 - stake size effect 178
- ffective response of experts 56
- age factors
 - educational incentives 560
 - investor skills 1501
 - lying 461
 - neuroeconomics of time preferences 1689
 - overconfidence 1528
 - patience 664
- agent-based computational economics (ACE) 25, 872–3
- aggregate monotonicity condition, evolutionary game theory 1091
- agrarian contracts, social networks applied to 1184–5
- all causes model 14–15
- Allais paradox 18, 26, 95–7
 - as common consequence violation 98
 - dual selves model with costly commitment 1303
 - hyperbolic discounting 627
 - probability weighting functions 116
 - prospect theory 28
 - rank dependent utility theory 124, 129
 - regret theory 182
 - salience and context in decision theory 185, 186
- allocative efficiency, competitive equilibrium experiments 871, 872, 874
- α -maximin expected utility (α -MEU) model 293, 298, 299, 309
- alternating offers bargaining game 801
- altruism 341–2, 343, 398–9
 - axiomatic foundations 450, 451
 - centipede games 733
 - dictator game 352, 353, 354, 356
 - evolution 385
 - evolutionary game theory 1039–40, 1073–84
 - cultural evolution, formal model of 1068–73
 - gene-culture coevolution 1066–7
 - fairness and stochastic dominance 414–15
 - Fehr–Schmidt model 417
 - genuine 345, 354
 - incentives 531, 533–40, 552–3, 554, 555
 - interdependent preferences models 447–9
 - neuroeconomics 1647, 1677–8, 1683
 - parochial 1233, 1647, 1677–8
 - public goods games 377
 - reciprocal 342, 1039–40, 1076–8
 - reluctant 345, 353, 354
 - social identity 469, 471
 - trust game 368
 - white lies 455, 461
- ambiguity 79, 284–6
- aversion 293, 298
 - empirical evidence 298–301
 - support theory 304
 - see also* smooth ambiguity aversion models
- behavioral models 81–2, 301–9
- and compound risk premium, relation between 108–9
- Ellsberg paradox 286–8
- neoclassical models 81–2, 284, 285, 286, 288–301
- preference reversals 108
- prospect theory 31, 80
- amygdala 1653
- human virtues 1685
- punishment and trust 1676, 1683
- risky decisions 1664–5, 1666, 1667
- strategic interaction 1692
- analogy-based expectations equilibrium (ABEE) 880, 966–9, 972
- anchoring 46, 1344–5, 1370–5
 - cognitive ability 1401
 - efficient markets hypothesis 1493
 - experts 57, 1449
 - frequency vs. probability format 1429
 - index 1370
 - mental accounting 1472–3
 - mental models 1408
 - soft paternalism 1599
 - strategic complements and strategic substitutes 856–63
 - type, level- k models 895
- anger 1232–3
- animals
 - cooperation 342
 - expected utility theory, violations of 104
 - neurology 1649
 - techniques 1655, 1656, 1657
 - time preferences 1686
 - see also* primates

- Anscombe–Aumann acts 290
- antecedents of behavioral economics 3–4
- anterior cingulate cortex (ACC)
 - guilt 1686
 - human virtues 1685
 - punishment and trust 1675, 1676, 1680, 1684
 - risky decisions 1664, 165
 - strategic interaction 1690
- anterior insula (AI)
 - human virtues 1685
 - punishment and trust 1675, 1676, 1680, 1684
 - time preferences 1688
- anticipated emotions 1233
- anticipation under certainty 1227, 1242–5
- anticipatory emotions 1223, 1226, 1233–4, 1235
- anticipatory learning models 1130
- antisocial punishments 53
 - crowding-in effect of incentives 559
 - equity, reciprocity, and competition model 410
 - public goods games 382–5
- anxiety
 - broad choice bracketing 1471
 - and risk aversion, distinction between 1250
 - under uncertainty 1227, 1245–50
- arbitrage limits 1487, 1509–16
- arbitrageurs 54–5, 1485
 - efficient markets hypothesis 1492, 1495, 1500
 - noise trader risk 1502–9
 - optimization problem 1512–13
- Arrow–Hurwicz criterion, multiple priors models 293
- artificial social identity 466
- aspect listing 234–5
- aspiration adaptation theory 1277, 1348, 1411–16
- aspiration-based experimentation 1148
- aspiration-based reinforcement learning 1101–2
- aspiration level learning 1155
- asset prices under anxiety 1247–50
- association area of the brain 1653
- assortativity index 1082
- assurance game 1710
- asymmetric information
 - bargaining games 789, 790, 791, 803–11
 - between emotions and cognition 1289, 1304–8
 - indirect paternalism 1588–9
 - psychological game theory 936–7
 - signaling games 819–43
 - temptation and conflicts between commitment and flexibility 1267–9
 - two-sided 843
- asymmetric paternalism 1580, 1593–5
 - cooling-off periods 1598
 - defaults 1596
- asymptotic stability
 - deterministic situations 1159–60
 - non-existent 1059–60
 - ordinary differential equations 1191, 1192
- attention grabbing stocks 1517
- attribute-based models 606, 627–32
- attribute substitution 1347, 1397–8
- attribution bias 1527
- attribution hypothesis, intention-based
 - reciprocity 439, 440, 442, 443, 444
- auctions
 - behavioral economics of 881, 988–1002
 - double
 - competitive equilibrium 723–4, 863, 864, 868–73, 874
 - induced value methodology 48–9
 - rationality of emotions 1231
- audience effect 354, 387
- augmented double auction institution 874
- automatic choices 354–5
- automaticity of the brain 1648
- autonomy 542
- availability cascades 1366
- availability heuristic 45, 1344, 1365–7
 - efficient markets hypothesis 1491, 1494
 - mental models 1408
- average-payoff reinforcement model 1118
- awareness
 - management 1316–20
 - present-biased preferences 647, 650–1, 653–5
- backward induction
 - dynamic games of perfect information 1714–15
 - finitely repeated asymmetric games 828, 830
 - search and lookups, inferring cognition through 818–19
- banking *see* financial markets
- bargaining games 718–21, 786–819
- bargaining impasse 340, 806, 813
- base rate neglect 46, 1345–6, 1351, 1375–8
 - and conservatism, reconciliation of 1379–83
 - experts 1449
 - frequency vs. probability format 1429
- baseball stadium vendors, labor supply decision 258
- battle of the sexes (BOS) game 1710
 - coordination 758
 - and communication 877, 907–11
 - extensive form games 1713, 1714, 1715
 - fictitious play 1108, 1112
 - forward induction 769–70
 - intergenerational advice 781–2
 - iterated deletion of dominated strategies 728
 - market entry game 915
 - mixed strategy Nash equilibrium 1708–9
 - preplay communication 771
 - social identity 469–70
- Bayes' rule
 - awareness management 1317, 1318
 - base rate neglect 1375, 1377, 1378, 1381
 - conservatism 1381
 - correlated equilibria 844
 - efficient markets hypothesis 1491, 1493
 - heuristics and biases program 46, 47
 - managerial overconfidence with rational investors 1535
 - perfect Bayesian equilibrium 1716, 1717
 - personal evaluation biases 1320–1
 - scope of economics 1604
 - signaling games 821, 834, 836, 882
 - stock market underreaction and overreaction 1522
- Bayesian approach to uncertainty 91

- Bayesian Nash equilibrium (BNE)
 - defined 973, 1082, 1719
 - incomplete information games 1719, 1721
 - indirect evolutionary approach 1081–2
- Becker–DeGroot–Marshak mechanism (BDM) 18
- Becker paradox 194
- behavioral corporate finance 1487–8, 1526–43
- behavioral corporate signaling 1488, 1540–3
- behavioral decision theory 79–82
 - ambiguity 284–309
 - applications 213–83
 - evidence on human choice under risk and uncertainty 83–109
 - models 110–212
- behavioral evidence, examples of 48–59
- behavioral life-cycle (BLC) models
 - mental accounting 1453, 1459–63
 - present-biased preferences 647, 655–8
- behavioral macroeconomics 651
- behavioral political economy 36
- behavioral public finance 1578, 1583, 1625–32
 - reduced-form approaches 1632–4
- behavioral rules 1144
- behavioral welfare criteria 1600–1
- behavioral welfare economics 1577–83
- belief-based learning 1093–4, 1106–12
 - empirical evidence 1093, 1112–15
 - experience weighted attraction model 1094, 1120–1, 1123–4
 - reinforcement learning vs. fictitious play 1094, 1115–20
- betrayal aversion 368
- better than average effect 1527–8
- betweenness axiom 102, 104
 - disappointment aversion theory 183
- biased assimilation *see* confirmation bias
- biased numerical attention 1440–3
- bicycle messengers, labor supply decision 258–60
- bimodal perception of risks model 116
- blame the theory argument 13
- blank slate view of the human mind 1065
- blood donations incentives 553–4
- blood-oxygen-level-dependent (BOLD) contrast 1658
 - risky decisions 1666, 1668–9
- bonus contracts 479–80, 495, 497
 - evidence on contract choice under contractual incompleteness 506–7, 508, 509
 - other-regarding preferences 498, 504–5
 - self-regarding preferences 498
- bonuses 1462–3
- bounded rationality 1339–41
 - asymmetric paternalism 1593–4
 - financial markets 1485–543
 - judgment heuristics 1342–450
 - level-*k* models 907–11
 - mental accounting 1451–84
 - soft paternalism 1597, 1599
 - stochastic social dynamics 1159, 1160, 1167, 1178
- brain 1645, 1648–55
 - gene–culture coevolution 1065
 - lesions 1656, 1660–1
 - loss aversion 1666
 - lobes 1649–50
 - regions 1650–1
 - two-modules view 1288, 1290–8
- brainstem 1651, 1676
- branch independence, prospect theory 193
- breast cancer screening 195–6
- broad choice bracketing 1469–71, 1480
- Brodman areas 1653, 1654
- Brownian motion 9–10
- bubbles, financial market 54–5, 1516, 1534
 - efficient markets hypothesis 1500
- buyer–seller game 1112–13
- bygones 52–4, 1673
- C-independence 292, 293, 298
- cab drivers *see* taxi drivers
- cancellation heuristic 193
- cancellation principle 1568
- canonical split form, prospect theory 174
- capital asset pricing model (CAPM) 1500
- case-based decision theory 7, 81, 114, 187–8
- categorization, learning models 1097, 1148–53
- catering hypothesis, merger/acquisition activity 1533, 1534
- caudate 1651, 1669, 1675
- causal base rates 1376
- causal reasoning 974
 - defined 979
 - and evidential reasoning 980
 - public goods games 975
 - rationality 976
- centipede game 1714
 - experience weighted attraction learning 1132
 - iterated deletion of dominated strategies 713–14, 729–34
 - reinforcement learning 1132
 - subjective heterogeneous quantal response equilibrium 972
- cerebral cortex 1651–3
- certainty effect
 - and common difference effect 625, 626
 - expected utility theory 93
 - preferences 626
 - Prelec function 122
 - support theory 303
- charitable donations 722, 835–9
 - functional magnetic resonance imaging 1658
 - incentives 554–5
- Charness–Rabin hybrid model 444–7
- cheap talk 456, 771, 840–2, 907
- χ -cursed equilibrium 880, 972–4
 - winner's curse 995–7, 1000–1
 - multiple signals 840–1, 842
 - polymorphic equilibria 1057–9
 - public signals and correlated equilibria 848–50
- choice acclimating personal equilibrium (CPE)
 - newsvendor problem 161–4
 - stochastic reference points under rational expectations 158, 159–64, 168, 171–2
- choice architect/architecture 1437–8, 1581, 1599
- libertarian paternalism 1591–3

- choice-based behavioral welfare economics 1577, 1582, 1611–18
- choice bracketing 1454, 1469–71, 1480
- choice data 1581, 1603–10
- choice function 1619
- choice uncertainty 235
- choice with limited attention (CLA) 1619–23
- choking under high incentives 544–8
- Choquet expected utility 82
 - ambiguity 284, 285, 289–91, 299, 308
- Choquet integral 289
- cingulate gyrus 1653, 1668, 1669
- civic cooperation, norms of 382–3
- classical decision makers 1584, 1585
- classical decision theory 84–9
- classical game theory 52, 53, 342
- classical preferences 1260–1
- closed-end fund puzzles 1492–3
- closing an account 1455, 1477
- coalescing lotteries 173
- coalition formation 1067
- coarse correlated equilibria 1155, 1723
- coarse thinking 1347, 1401–7
- Coase theorem 231
- cognition
 - and emotions, interaction between 1225, 1288–321
 - inferring through search and lookups in a bargaining game 720–1, 815–16
- cognitive ability
 - dual process models 1400
 - judgment heuristics 1400–1
- cognitive dissonance 1399–400
- cognitive hierarchy (CH) model 41, 43–4, 876–7, 895, 900–4
 - applications 877, 916
 - focal points 763
 - functional experience weighted attraction learning 1127
 - iterated deletion of dominated strategies 733–4
 - market entry game 916
 - prisoner's dilemma game 978
 - subjective heterogeneous quantal response equilibrium 969, 971
 - truncated heterogeneous quantal response equilibrium 971
 - winner's curse 1001
- cognitive miser model 1399
- cognitive niche hypothesis 1064–5
- cognitive psychology 4
- cognitive reflection test (CRT) 1400–1
- coherent arbitrariness 1454, 1471–3
- commitment
 - and flexibility, conflicts between 1228, 1266–70
 - investment goods 687
 - optimal consumption under 655–6
- common consequence effect 122
- common consequence violation
 - of EU independence axiom 97, 98–9, 103, 104
 - of EU rationality axioms 101
- common difference effect 583, 584, 587, 605, 611
 - exponential discounted utility model 37, 590, 594–6
- hyperbolic discounting 612–19
 - Loewenstein–Prelec explanation 607, 608
 - reference time theory 638, 641
 - subadditive discounting 619–22
 - uncertainty 623–7
 - vague time preferences model 628–9
- common ratio effect
 - preferences 625–6
 - probability weighting functions 122
 - risk and time preferences, integration of 201
- common ratio violation
 - of EU independence axiom 97–8, 103, 104
 - of EU rationality axioms 99–101
 - of EU reduction axiom 97–8
- common values (CV) auctions 989, 991–2, 993, 994, 995
 - winner's curse 997
- communication
 - battle of the sexes game 877, 907–11
 - coordination games 783–5
 - gene–culture coevolution 1064, 1067
 - level-*k* models 877, 907–11
 - Markov chains 1198
 - and morality 552
 - preplay *see* preplay communication
 - psychological game theory 953–4
- comonotonic independence, Choquet expected utility 290
- comonotonicity, Choquet expected utility 290
- compensating variation method, exponential discounted utility model 597
- compensatory processes, cue-conditioned behavior 1236–7
- competence hypothesis 303
- competition
 - and contracts 264, 266, 267, 283
 - Fehr–Schmidt model 403–7
 - preferences 356
 - present-biased time preferences 645, 677–9
 - trust game 370
- competitive equilibrium (CE) 723–4, 863–74
 - fairness considerations 48–52
 - gift exchange game 360
- completely mixed strategy, evolutionary game theory 1041, 1059
- complexity
 - learning models 1097, 1148–53
 - theory 25
- composite Prelec probability weighting function (CPF) 196–8
- composite prospect theory (CPT) 81, 115, 140
 - and configural weights models, comparison of 1909
 - defined 198
 - extreme probability events 191, 196–9
 - Prelec function 122
- composite rank dependent utility theory (CRDU) 191, 198–9
- compound invariance, Prelec function 151
- compound lotteries
 - ambiguity 300
 - classical decision theory 85
 - failure of compounding 285, 294–6
 - process dependence 211, 212

- compound risk premium 108–9
- compounding rule of probabilities 1043
- conditional cooperation 342
 - evolutionary game theory 1076–7
 - public goods games 356, 374, 375, 377
 - external validity of lab evidence 390, 391, 393
- Condorcet winner 418–22, 423
- confidence
 - defined 1315
 - effect 1608
 - emotions and cognition, interaction
 - between 1311–21
 - frequency effect 1422
 - psychological game theory 931
- configural weights models 31, 81, 114, 188–91
 - probability weight function 32
 - prospect theory 173–4
- confirmation bias 47, 1346, 1389–95
- experts 1450
- information flows in financial markets 1518
- mental models 1408
- neuroeconomics 1648
- overconfidence 1389–95, 1422
- conglomerate wave 1529, 1534
- conjunction fallacy 1343–4, 1361–4
- cognitive ability 1400–1
 - frequency vs. probability format 1429
- conservatism 1346, 1378–9
 - and base rate neglect, reconciliation of 1379–83
 - cognitive ability 1401
 - efficient markets hypothesis 1491
 - noise trader risk 1504
 - stock market underreaction and overreaction 1520
- consistent evidential equilibria 981–2
- constrained sparse max operator 1447
- consumer durables, and mental accounting 1464
- consummate performance, contracts 264–5, 266, 267
- consumption
 - independence, violations of 603
 - optimal, under hyperbolic discounting 651–64
 - utility 154–5, 260
- context in decision theory 81, 114, 184–7
- continental divide game
 - functional experience weighted attraction
 - learning 1127, 1128
 - historical accidents and coordination 772, 773
- continuity axiom of expected utility theory 85
- contract theory 36, 81
- contractarian approach 1578, 1583, 1586, 1623–5
- contracts
 - baseline 277
 - employment 480, 521–7
 - explicit 507–8
 - implicit 507–8
 - independent 489
 - legal environment 556
 - long-term 278–83, 480–1, 510–27
 - perturbed 277–8
 - as reference points 216, 264–72
 - relative performance 488, 489, 490
 - renegotiation 216, 264, 278–83
 - rigid vs. flexible 265, 270–2
 - team 488, 489, 490
 - see also* bonus contracts; incentive contracts; incomplete contracts; optimal contracts; trust contracts
- control
 - aversion 538–9, 542–8
 - illusion of 1528
- conventions 1170, 1176, 1178
- convex monotonicity, evolutionary game theory 1091
- cooling-off laws/periods 1257, 1598
- cooperation 342–3
 - conditional *see* conditional cooperation
 - correlated equilibrium and social norms 963
 - evolutionary game theory 1040, 1073, 1074, 1084
 - gene–culture coevolution 1066–7
 - group selection 1075
 - indirect reciprocity 1078
 - reciprocal altruism 1077
 - repeated games 1079
 - Fehr–Schmidt model 407–10
 - learning direction theories 1133
 - prisoner's dilemma game 977–8
 - in production 350
 - rule-based learning 1148
 - strategic complements and strategic substitutes, implications of 853–6
 - ultimatum game 350
- coordination
 - games 715–18, 758–86, 1710
 - aspiration adaptation 1413–15
 - battle of the sexes 877, 907–11
 - belief-based learning 1108–9
 - cognitive hierarchy model 902
 - correlated equilibrium and social norms 963
 - evolutionary game theory 1060–3
 - fictitious play 1112
 - level-*k* models 904–6
 - rational learning 1154
 - rule-based learning 1142
 - stochastic social dynamics 1158, 1161–4, 1170–5, 1178
 - level-*k* models 877, 904–11, 914
- corporate finance, signaling in 721–2, 831–5
- corporate turnaround game 774–5, 783–5
- correlated equilibria 722–3, 843–51, 1723
 - coarse 1155, 1723
 - defined 1723
 - and social norms 879–80, 962–6
- counterfactual emotions 1223
- Cournot adjustment dynamics 1112, 1707
- Cournot games 1708
- Cox–Friedman–Gjerstad model 449
- credit card debt, illiquid assets held with 658
- critical cost-efficiency index (CCEI) 356
- crowding-in effect of incentives 558–61
- crowding-out effect of incentives 553, 558
- cue-conditioned behavior and habit
 - formation 1226–7, 1235–40
- cues
 - fast and frugal heuristics 1419–24, 1426
 - model 1240–2
 - scope of economics 1606

- cultural evolution, formal model of 1068–73
- cultural factors
 - gift exchange game 360
 - public goods games 384
 - ultimatum game 1072
- cultural intelligence hypothesis 1064
- cultural niche hypothesis 1065
- culture
 - defined 1064
 - gene–culture coevolution *see* gene–culture coevolution
 - see also* cultural factors
- cumulative payoff matching rule (CPM) 1102
- cumulative prospect theory
 - common difference effect 624
 - and original prospect theory, difference between 111
 - see also* prospect theory
- current asset mental account 1460
- current income mental account 1460
- curse of debiasing 1436
- cursed equilibrium 880, 972–4
 - winner's curse 995–7, 1000–1
- cyclical choices
 - exponential discounted utility model 601–2
 - theories compatible with 601
- deadlines
 - hyperbolic discounting 689–90
 - normative bargaining solutions 792
 - soft paternalism 1598–9
- debiasing, curse of 1436
- debt aversion 1463, 1464, 1468
- deception, neural basis of 1684
- decision field theory (DFR) 299
- decision theory
 - classical 84–9
 - see also* behavioral decision theory
- decision utility 680
 - behavioral welfare economics 1600–1, 1633
 - and experienced utility, distinction between 1600, 1662
 - happiness economics 1228, 1270, 1274
 - neuroeconomics 1662, 1665
- decision weights, rank dependent utility 28, 124
- declining sensitivity property, reference time theory 633
- dedicated intelligence 1064
- defaults
 - behavioral welfare economics 1577
 - hyperbolic discounting 689
 - soft paternalism 1596–7
- delay
 - under certainty 1227, 1242–5
 - dependence 206–10
 - function 638–9, 640, 641
 - premium 642
 - reference time theory 638–9, 640, 641, 642
- delay–speedup asymmetry 598–600
- deliberative choices 354–5
- democracy
 - direct 417–26
 - and happiness 1286–7
- description invariance, violations of 106–7
- descriptive methodology 5
- deterministic evolutionary dynamics 1140
- dictator game (DG) 344, 345, 352–5
 - altruism 398–9
 - Charness–Rabin hybrid model 445, 447
 - crowding-in effect of incentives 559
 - envy 399
 - equity, reciprocity, and competition model 410
 - gene–culture coevolution 1067
 - guilt 954–5, 956–60
 - human sociality 33, 35
 - human virtues 461
 - incentives 549, 550, 552, 559
 - intention-based reciprocity 439
 - interdependent preferences models 448
 - internal validity 12, 13
 - mental accounting 1453
 - neuroeconomics 1683
 - positive bargaining solutions 806
 - rationality of emotions 1232
 - rationality of other-regarding preferences 355–7
 - social identity 466
 - social preferences models, evidence on 431–2, 433
 - student subject pool 387, 388
 - third party punishments 396, 397
 - and ultimatum game, comparison between 352
- direct paternalism 1580, 1585, 1587–8
- direct reciprocity, evolutionary game theory 1076–8
- disappointment aversion theory 81, 101–2, 114, 154, 155, 182–4
 - and Köszegi–Rabin framework, distinguishing between 161
- discount factors
 - exponential discounted utility model 588, 589–90, 594, 595, 615–16
 - hyperbolic discounting 615–16
 - quasi-hyperbolic discounting 615–16
- discount function
 - additive 634–5
 - discounted utility model 588
 - exponential discounted utility model 588
 - interval 619–20
 - reference time theory 637
- discount rate
 - discounted utility model 588
 - exponential discounted utility model 588, 589–90, 595–6, 603–4
- discounted utility (DU) model 583, 586, 587–8, 632
 - anticipation and delay under uncertainty 1243, 1245
 - mental accounting 1463–4, 1466
- disposition effect
 - efficient markets hypothesis 1493
 - financial markets 54, 56
 - gambler's fallacy 1354–5
 - mental accounting 1477
- distribution hypothesis, intention-based reciprocity 439, 441, 443–4
- diversification heuristic 1455, 481–2, 1483
- dogfooding 1450

- domain of gains/losses 131, 133, 300
 - choice bracketing 1471
 - goals 251–3, 254, 255
 - neuroeconomics 1663–4
 - risk 135–6
 - stake size effect 177–8
 - tax evasion puzzles 244, 245, 246
 - utility function 141–2
 - domain satisfaction 1283
 - domain-specific intelligence 1064
 - dominance
 - defined 675
 - heuristic 193
 - seeking 1067
 - dopaminergic reward prediction error (DRPE)
 - hypothesis 1663
 - dorsolateral prefrontal cortex (dlPFC) 1653
 - guilt 1686
 - punishment and trust 1679, 1680–2
 - risky decisions 1663, 1664, 1667
 - strategic interaction 1690, 1692
 - time preferences 1688, 1689
 - dorsomedial prefrontal cortex 1678, 1679, 1686
 - double auctions
 - competitive equilibrium 723–4, 863, 864, 868–73, 874
 - induced value methodology 48–9
 - drift-diffusion model 1662
 - dual brain modules *see* two-modules view of the brain
 - dual selves model
 - with costly commitment 1288–9, 1298–304
 - neuroeconomics of time preferences 1689
 - Duhem–Quine thesis (DQT) 5
 - duration effects, gift exchange game 361–3
 - Dutch price auctions 989–90, 992, 1231
 - dynamic game
 - of imperfect information 1713, 1714, 1715
 - of perfect information 1713, 1714–15
 - earned income tax credit (EITC) program, USA 1439
 - Easterlin paradox 1229, 1275–83
 - ecological validity 751
 - defined 1420
 - fast and frugal heuristics 1420, 1421
 - mixed strategy Nash equilibria 751
 - economic freedom and happiness 1286–7
 - economist-therapists 1599–600
 - educational incentives
 - crowding-in effect 560
 - extrinsic and intrinsic motivation 554, 556, 560
 - efficiency
 - and control, trade-off between 546
 - gain, multiple signals 841, 842
 - preferences 432–3
 - efficient markets hypothesis (EMH) 55–6, 1485–6
 - anchoring 1371
 - behavioral finance 1526
 - defined 1490
 - empirical evidence 1495–502
 - implications 1490–1
 - information flows 1517
 - semi-strong form 1491, 1498
 - short selling 1488–90
 - stock market underreaction and overreaction 1520, 1521
 - strong form 1491
 - theoretical case 1491–5
 - weak form 1491, 1497–8
 - elasticity of intertemporal substitution of consumption (EIS) 663–4
 - elimination by aspects (EBA) strategy 1425
 - Ellsberg paradox 79, 284–5, 286–8
 - behavioral models 286
 - vs. neoclassical models 82
 - prospect theory 304, 305, 306, 307
 - support theory 303
 - neoclassical models 285, 289, 297
 - vs. behavioral models 82
 - Choquet expected utility 289, 291
 - empirical evidence 298–9
 - multiple priors models 292
 - smooth ambiguity aversion model 296
 - two-stage recursive models 294–5, 296
 - probabilistic sophistication violation 92
 - rank dependent utility theory 129
 - reduction axiom of EU theory 109
 - suspicion 298–9
- emotional repair 1482
- emotions 1223–5
 - and cognition, interaction between 1225, 1288–321
 - and human behavior 1224, 1226–87
 - rationality 1226, 1231–5
 - stake size effect 178–9
 - taxonomy 1223
- empathy 1675–6
- empirically progressive research programs 6
- employment contracts 480, 521–7
- emulation dynamics 1040, 1085–7
- end-game effects 855
- endogenous reference points 30
- endowment effect 213–14, 217–35
 - indirect evolutionary approach 1080
 - loss aversion 214, 217, 220–9, 232, 233
 - primates 237–8
 - public goods games 376
 - risk aversion 222–4
 - visceral influences 1230
 - see also* exchange asymmetries
- English price auctions 989–90, 992, 1231, 1432
- Enlightenment view 1416
- entitlement norm, ultimatum game 351
- entrepreneurs, overconfidence 1528
- entropy, limited attention 1448
- entry deterrence model 828–31
- envy 342, 399
 - incentives 500
 - political economy 417, 421
 - social identity 469
- epigenetic changes 1065, 1066
- epistemic foundations of equilibrium concepts 25
- epistemic games 964
- equal weighting strategy (EQW) 1425

- equilibrium
 - in beliefs interpretation of mixed strategies 59, 738
 - dominance, in signaling games 1721, 1722
- equity premium puzzle
 - anxiety 1250
 - efficient markets hypothesis 1493, 1494, 1499
 - financial markets 55, 56
 - myopic loss aversion 81, 215, 241
- equity, reciprocity, and competition (ERC)
 - model 35–6, 400, 410–12
 - evidence on 434
 - positive bargaining solutions 799
 - and quantal response equilibrium 800
- equivalent variation method 597
- ergodicity state 1158, 1160, 1198, 1199
- esteem incentives 533–40
- ethnicity factors
 - social identity 467, 468
 - trust game 368
- Euler equation
 - exponential discounting 659–60
 - hyperbolic discounting 660–2, 663, 664
- eusociality 1064
- event contingent lotteries 91, 92
- event splitting 190
- event studies, efficient markets hypothesis 1495
- event weighting function 128–9
- evidence based rules 1145
- evidential equilibrium 880–1, 974–84
- evidential games 979–84
- evidential reasoning (ER) 974–7, 1396
 - defined 979
 - evidential games 980, 981, 982–4
 - heuristics and biases approach 977
 - prisoner's dilemma game 978, 982–4
 - rationality 976
 - about states of other humans 975
- evolution 1037
 - brain 1649, 1651
 - emotions and cognition, interaction between 1290
 - features 1290–1
 - low probability events, neglect of 104
 - mental models 1408
 - primates 235–9
 - public goods games 385–6
 - punishment 1673
 - rationality 1431
 - of emotions 1233
 - regular evolutionary system
 - visceral influences 1230
- evolutionary game theory 1033, 1037–43
 - altruism and reciprocity 1073–84
 - applications 1038, 1053–63
 - evolutionary stable strategies
 - relation with other equilibrium concepts 1047–8
 - in two-player games 1043–7
 - gene–culture coevolution 1063–73
 - playing the field 1063
 - prisoner's dilemma game 978
 - replicator dynamics 1048–53
 - social learning and replicator-like dynamics 1085–91
- evolutionary political economy 1067–8
- evolutionary stability 1082
- evolutionary stable strategies (ESS) 1037, 1038, 1043–7
 - asymptotically stable equilibrium, non-existence of 1059
 - chicken game 1057, 1058–9
 - defined 1045, 1046
 - indirect evolutionary approach 1080, 1081–2
 - indirect reciprocity 1079
 - population game 1063
 - primary criterion 1045, 1046, 1058–9
 - prisoner's dilemma game 1054–6
 - reciprocal altruism 1076–7
 - relation with other equilibrium concepts 1038, 1047–8, 1051
 - and replicator dynamics, relation between 1051–2, 1053
 - secondary criterion 1046, 1059
- evolutionary system 1089
- excess burden of taxation under limited attention 1583, 1628–32
- exchange asymmetries 213–14, 217
 - marketing implications 232–3
 - property rights 231–2
 - subject misinterpretation 229–31
 - uncertainty 235
 - see also* endowment effect
- exchangeable events 305, 306
- expected–ideal (EI) gap, self-control 1270
- expected utility theory (EU) 79–80
 - ambiguity 284, 288, 297, 299
 - Becker–DeGroot–Marshak mechanism 18
 - cognitive system 1292
 - common difference effect 623, 624, 625
 - defences 30
 - efficient markets hypothesis 56, 1494
 - and exponential discounted model, combining 606, 624, 625
 - goals 215, 250–1, 254
 - implicit 101
 - incentive compatibility 18, 19
 - linear weighting of probabilities 110
 - low probability events 195, 196
 - market experience and loss aversion 228–9
 - order axiom of 85
 - parameters 30–1
 - number of 138, 139
 - preference foundations 85–7
 - primates 237
 - and prospect theory 26, 27, 28, 29, 30, 81
 - Rabin's paradox 250
 - random lottery incentive mechanism 18
 - rank dependent utility 28, 29
 - rationality axioms 86–7
 - violations 99–101, 108
 - risk, attitudes to 31
 - second price auction 18
 - tax evasion puzzles 243, 244, 245–6, 247–8
 - violations 95–109, 137–8

- expected utility theory (EU) (*cont.*)
 - see also* independence axiom of expected utility theory; reduction axiom of expected utility theory; subjective expected utility theory; utility function: expected utility theory
- expected value maximization (EV)
 - classical decision theory 87, 88
 - parameters, number of 138, 139
 - subjective expected utility theory 90
- expected value theory 30
- experience weighted attraction (EWA)
 - learning 1094–5, 1120–30, 1138
 - empirical evidence 1094, 1125
 - EWA cube 1124
 - EWA-lite 1126
 - functional 1095, 1125–8, 1137, 1138
 - mixed strategy Nash equilibria 748
 - sophisticated 1095, 1128–30, 1131
 - viewpoints 1131–2
- experienced utility 680
 - behavioral welfare economics 1600–1, 1633
 - and decision utility, distinction between 1600, 1662
 - happiness economics 1228, 1270–1, 1272, 1274
- experimental economics 4, 9, 10–21
 - and behavioral economics, relation between 2
- experimental scrutiny 387, 388
- experimenter demand effect 10, 11, 12, 386
 - other-regarding preferences 387
- experts
 - biases 1350, 1449–50
 - affect heuristic 1369
 - anchoring 1372
 - hindsight bias 1386
 - neoclassical economics 56–7
- exponential discount function 637
- exponential discounted utility (EDU) model 37–8, 39, 583, 586–7, 588–9
 - anomalies 594–607
 - reference time theory 638–41
 - axiomatic foundations 593
 - Euler equation 659–60
 - job search 689
 - payday effect 689
 - precautionary savings 662–3
 - salient properties 598–90
 - time consistency of consumption plans 590–1
 - time consistent preferences 652
 - time preferences, elicitation of 591–2
- exponential discounting 1292
- external arbiter recommendations 717, 783
- external validity
 - of experiments 11, 14–15
 - of other-regarding preferences 386, 388–95
- extreme probability events 191–9
 - see also* low probability events
- extreme value distribution 883
 - Type I 1193–4, 1195
- extrinsic motivation
 - asymmetric information in principal–agent framework 507
 - empirical evidence 482–3
 - in experiments 16–18, 19
 - theoretical framework 481
- fairness
 - axiomatic foundations 450
 - Fehr–Schmidt model 401, 403–11
 - political economy 417, 418
 - general equilibrium and welfare 426–31
 - human sociality 33–4, 36
 - intention-based reciprocity 438
 - interdependent preferences models 448
 - of lies 457
 - motivations 19
 - norm
 - dictator game 354, 355
 - human sociality 36
 - third party punishment game 396
 - ultimatum game 351
 - preferences, acquired nature of 355
 - psychological game theory 879, 937–49
 - reference points 439–44
 - stochastic dominance 400, 412–17
- false consensus effect 1396–7
- evidential equilibrium 974–5
- experts 57, 1449
- fanning-in hypothesis 103–4, 183
- fanning-out hypothesis 101–4, 182, 183
- fast and frugal heuristics 1348–9, 1416–26
 - great rationality debate 1426–31
- fear
 - of negative evaluation (FNE) 301
 - under uncertainty 1227, 1245–50
- feedback market 369
- Fehr–Schmidt model 34–5, 36, 399–400, 401–11, 419, 420, 422
 - and equity, reciprocity, and competition model, comparison between 410, 411
 - evidence on 434
 - fairness and stochastic dominance 412–15, 417
 - fairness, general equilibrium, and welfare 428, 429
 - general form 412
 - incentives
 - incomplete contracts 498, 499
 - moral hazard 484
 - reciprocity and long-term contracts 512, 524
 - leader contributions to charitable giving 836–9
 - political economy 417–18
 - promises and threats model 462–3, 464
 - and quantal response equilibrium model, comparison between 800
- Fehr–Schmidt utility 419, 420, 422
- fictitious play 1106–12
 - empirical evidence 1112–13
 - historical accidents and coordination 774
 - quantal response equilibrium 893
 - vs. reinforcement learning 1094, 1115–20
 - smooth 1108
 - traveller's dilemma game 893
 - weighted 1111–13, 1127–8
- financial incentives
 - conjunction fallacy 1363–4
 - coordination games 717, 774–5, 784

- financial markets
 - bounded rationality in 1340–1, 1485–8
 - arbitrage limits 1509–16
 - behavioral corporate finance 1526–43
 - efficient markets hypothesis 1488–502
 - information, gradual flow of 1516–20
 - noise trader risk 1502–9
 - stock market underreaction and overreaction 1520–6
 - efficiency 54–6
 - fast and frugal heuristics 1424, 1425
 - recognition heuristic 1424
- financial volatility, underestimation of 1387–9
- firefighter volunteers, incentives 555
- firm
 - nature of the 521–7
 - theory of the 264
- first order Fehr–Schmidt dominance 416
- first order stochastic dominance
 - classical decision theory 88, 89
 - defined 413
 - and fairness 412, 416
 - incentives 518
 - prospect theory 173–4
 - rank dependent utility theory 128
- first price sealed-bid auctions 989–90, 992, 993, 994
- winner's curse 998–9, 1002
- fitness 1037
 - defined 1043
 - evolutionary stable strategies 1043
 - inclusive 1039, 1073–4, 1076
- fixed fee tariffs 1468
- flexibility and commitment, conflicts between 1228, 1266–70
- focal points
 - in bargaining 811–15
 - and coordination 716, 759–63
 - level-*k* models 904–7
 - see also* salience
- follow the herd theory 1146
- Food Stamp Program, USA 1478–9
- forward induction, coordination games 716, 769–71, 776–7
- fourth generation prospect theory 199
- framing effects 107
 - absence of 106–7
 - behavioral welfare economics 1601, 1625
 - contractarian approach 1625
 - coordination games 771
 - dictator game 352
 - efficient markets hypothesis 1493
 - experts 57, 1449, 1450
 - gift exchange game 360
 - guilt 951
 - incentives 560
 - information flows in financial markets 1517
 - iterated deletion of dominated strategies 727–9
 - learning models 1151
 - mental accounting 1453, 1455–6
 - neuroeconomics 1664–5
 - scope of economics 1608–9
 - survey data 20
- free-riding 340
 - incentives
 - crowding-in effect 559
 - extrinsic and intrinsic motivation 550, 557, 559
 - incomplete contracts 509
 - public goods games 356, 377, 378, 379
 - causal reasoning 975
 - empirical evidence 372, 373, 375
 - external validity of lab evidence 390, 392, 393
 - Fehr–Schmidt model 408
 - unproductive expenditures 382
 - with punishments 373
 - without punishments 371
- freedom and happiness 1286–7
- front-end delay dependence 600–1
- full information
 - bargaining games
 - Nash bargaining approach 788, 790, 791–2
 - positive bargaining solutions 793–803
 - leader contributions to charitable giving 839
- fully rational outcomes, stochastic social dominance 1178
- functional experience weighted attraction (fEWA)
 - learning 1095, 1125–8, 1137, 1138
- functional magnetic resonance imaging (fMRI) 1658
 - risky decisions 1662, 1663, 1664, 1666, 1668
 - social preferences
 - guilt 1685, 1686
 - human virtues 1685
 - inequity aversion 1670, 1672
 - punishment and trust 1676, 1677, 1678, 1679, 1680, 1684
 - social rewards 1669
 - strategic interaction 1690
 - time preferences 1687, 1689
- fundamental risk
 - arbitrageurs 1502–3
 - efficient markets hypothesis 1489–90
 - noise traders 1503
- Fundamental Theorem of Sociology 1072
- fundamental transformation, contracts 265
- future wealth mental account 1460
- gain–loss asymmetry
 - exponential discounted utility model 597
 - Loewenstein–Prelec explanation 605, 607, 608
- gain–loss separability, prospect theory 174–5
- gain–loss utility, prospect theory
 - choice acclimating personal equilibrium 161–2
 - Kőszegi–Rabin framework 169, 172
 - labor supply decisions 260
 - renegotiation, long-term contracts, and loss aversion 279, 280
 - stochastic reference points under rational expectations 154–6, 157
 - unacclimating personal equilibrium 165–6
- gains decomposition utility (GDU) model 189–90
- gamble, classical decision theory 84
- gambler's fallacy 45, 1343, 1352–5, 1356, 1357–8, 1359–60
- cognitive ability 1401

- gambler's fallacy (*cont.*)
 - dual process model of the brain 1431
 - scope of economics 1609–10
- game of chicken
 - learning models 1153
 - test (CAT) 1401
- game theory 709–12, 1703–23
 - evidence 709–11, 713–874
 - evolutionary *see* evolutionary game theory
 - models 711–12, 875–1002
 - p-beauty context 43
- gender factors
 - ambiguity aversion 300
 - crowding-out effect of incentives 553
 - educational incentives 560
 - empathy 1676
 - goals, effect on performance 255
 - incentives 553, 554, 560
 - investor skills 1501
 - lying 461
 - mental models 1409
 - present-biased preferences 691
 - risk attitudes 137, 181, 1235
 - social identity 467–8
 - trust game 368
- gene–culture coevolution 1039, 1063–73
 - reciprocal altruism 1077
 - social norms 962
- general equilibrium, fairness and welfare 400–1, 426–31
- generalized axiom of revealed preference (GARP)
 - altruism 399
 - defined 356
 - dictator game 355, 356
 - primates 237
- generalized Lorenz dominance 412, 417
- generous tit-for-tat strategy 1077
- genetics 1278
- genotype 1037, 1066
- genuine altruism 345, 354
- gift exchange game 344, 345–6, 357–66
 - envy 399
 - equity, reciprocity, and competition model 410
 - experiments and internal validity 11
 - external validity of lab evidence 388–9
 - human sociality 33
 - incentives 540–1, 543
 - intention-based reciprocity 440
 - psychological game theory 928
 - social preferences models, evidence on 433
 - student subject pool 386
 - and trust game, relation between 366–70
- gift giving 1468–9, 1477–9
- glamour stocks 1497
- Glazer–Perry mechanism 714, 734, 735, 736–7
- goals
 - mental accounting of 1455, 1482
 - as reference points 215–16, 250–5
 - systems 1416
- Goldstein–Einhorn probability weighting
 - function 122
 - prospect theory 146
 - stake size effect 176, 177, 179
 - rank dependent utility theory 127
- great rationality debate 47–8, 1349, 1426–31
- green-beard effect 1080, 1081
- gross domestic product (GDP) and happiness, relation
 - between 1274–83, 1285–6
- group selection 1037, 1039, 1075–6
- group size, gradual growth in 717, 778–80
- groupthink 1362–3
- growth factors 1044
- guilt
 - aversion 465
 - defined 1685
 - neural basis 1647, 1685–6
 - psychological game theory 879, 925, 927–8, 949–60
- habit formation
 - cue-conditioned behavior 1226–7, 1235–40
 - projection bias 1253–6, 1257
- happiness economics 1228–9, 1270–87
 - altruism 414
 - behavioral welfare criteria 1600–1
 - scope of economics 1607–8
 - survey data 20
- hard core of a research program 5–6
- hawk–dove game 395, 1112–13
- health satisfaction and happiness 1285–6
- heavy-handed paternalism 1577, 1579
- hedonic benefits/costs 1464–5, 1466–8
- hedonic treadmill 1277
- herding behavior 55, 890, 1495
- heterogeneous quantal response equilibrium (HQRE) 880, 969–72
 - normal form games 886, 888
 - subjective heterogeneous 880, 969–72
- heuristics and biases program 44–8, 976–7, 1339–40, 1342–7
 - affect heuristic 1367–9
 - anchoring and adjustment 1370–5
 - attribute substitution 1397–8
 - availability heuristic 1365–7
 - base rate neglect and conservatism 1375–83
 - coarse thinking and persuasion 1401–7
 - confirmation bias 1389–95
 - conjunction fallacy 1361–4
 - dual process models 1398–401
 - false consensus effect 1396–7
 - financial markets, efficiency of 54
 - great rationality debate 1426–31
 - hindsight bias 1383–9
 - mental models 1407–9
 - necessary and sufficient conditions, confusion
 - between 1397
 - regression towards the mean 1395–6
 - representativeness heuristic 1351–2
- hindbrain 1651
- hindsight bias 46, 1346, 1383–9, 1390
 - anchoring 1373
 - dual process model of the brain 1431
 - efficient markets hypothesis 1493
 - experts 56, 1449
 - financial markets 54

- historical accidents and coordination 717, 772–4
- history of behavioral economics 3–4
- Homo hamiltonensis 1083
- Homo kantiensis 1083
- Homo moralis 1082
- Homo oeconomicus 1083, 1084
- honesty, evidence on 455, 457–8, 459–60
- hostile intent 482, 540–2
- hostile takeover wave 1529
- hot hands fallacy 45, 1343, 1355–8
 - dual process model of the brain 1431
- hot hands phenomenon 1355
- house money effect 19, 1455, 1479–81
 - exchange asymmetries 230
 - exponential discounted utility model 594
- human sociality *see* sociality, human
- human virtues *see* virtues, human
- hybrid models of other-regarding preferences 398, 401, 444–52
- hyperbolic discounting 38, 39–40, 583–4, 587
 - common difference effect 605–6, 612–19, 627
 - emotional system 1292
 - empirical evidence 616–19
 - Euler equation 660–2
 - exponential discounted utility model 594–6
 - general model 612–14
 - neuroeconomics of time preferences 1687
 - present-biased preferences 644
 - applications 646, 688–91
 - modeling multiple selves 646–51
 - optimal consumption 651–64
 - procrastination and preoperation 677
 - quasi-hyperbolic model 37–40
 - reference time theory 637
 - similarity relation 630, 631, 632
 - simple increasing elasticity class of utility functions 611
 - soft paternalism 1596, 1597, 1598
- hypothesis-based filtering 1389
- hypothesis testing 1153, 1156–7
- identity
 - social *see* social identity
 - soldier vs. civilian 341
 - within firms 341
- identity economics 473–7
- illiquid assets, optimal consumption under 656–7, 658
- illusion of control, managers 1528
- imagination bound 970
- imitation
 - gene–culture coevolution 1064, 1065
 - imitate the best strategy 1167, 1178
 - social 1040, 1085–7
- imitation learning 1087–9, 1114, 1115
- immediate costs 671–3, 676
- immediate rewards 673–5, 676
- impatience and patience
 - covariation with wealth and age 664
 - discounted utility function 588
 - exponential discounted utility model 595–6
 - hyperbolic discounting 612, 617, 619
 - reference time theory 641
 - sign and magnitude effects in time discounting 607
 - social identity 468
 - subadditive discounting 620–1, 622
 - see also* present-biased preferences
- imperfect self-information, regulation under 1581, 1601–3
- improvisational intelligence 1064
- impulse balance equilibrium 1095, 1133, 1134, 1135–6, 1137, 1138
- impulse matching learning 1134–5, 1137, 1138
- impulse matching theories 1095–6, 1134–8
- in-kind transfers 1455, 1477–9
- incentive compatibility constraint (IC) 274, 275
- incentive contracts 479–80, 495, 496–7
 - evidence on contract choice under contractual incompleteness 506, 508
 - extrinsic and intrinsic motivation 540, 542–3
 - other-regarding preferences 498, 503–4, 505
 - self-regarding preferences 498
- incentives 478–83
 - extrinsic and intrinsic motivation
 - empirical evidence 540–61
 - theoretical framework 527–40
 - financial *see* financial incentives
 - identity economics 473–7
 - incomplete contracts 495–510
 - for lying/truth-telling 455, 456, 457, 460, 461
 - moral hazard 483–95
 - moral hazard, loss aversion, and optimal contracts 276, 277
 - reciprocity and long-term contracts 510–27
 - role in economics 16–21
- inclusive fitness 1039, 1073–4, 1076
- income
 - and consumption, hyperbolic discounting 657
 - comparisons, and happiness 1278–80
 - inequality, and happiness 1284–5
- incomplete contracts 479–80, 511, 519
- increasing sequences, preferences for 602–3
- independence
 - control aversion 542
 - of irrelevant alternatives 675, 676
- independence axiom of expected utility theory 26, 85
 - random lottery incentive mechanism 18, 59
 - relaxation attempts 101–4
 - violations 95–9
- independent private values (IPV) auctions 988–90, 995
- indirect evolutionary approach 1040, 1080–4
- indirect paternalism 1580, 1586, 1588–9
- indirect reciprocity 342, 349, 395–7
 - equity, reciprocity, and competition model 411
 - evolutionary game theory 1040, 1078–9
- induced value methodology 48–9
- inequity aversion 34–6, 342
 - axiomatic foundations 451, 452
 - Charness–Rabin hybrid model 444, 446, 447
 - dictator game 352, 353, 356
 - equity, reciprocity, and competition model 434
 - Fehr–Schmidt model 402, 403, 408, 434
 - gift exchange game 364

- inequity aversion (*cont.*)
 - intention-based reciprocity 435–6, 439–44
 - moral hazard 478–9, 483–90
 - neuroeconomics 1646, 1670–3, 1675
 - promises and threats model 463, 464
 - psychological game theory 945–9
 - social preferences models, evidence on 432–3
- inertia, macroeconomics of nominal 723, 856–63
- inflation 1283–4
- information
 - cascades 890
 - disclosure value to self 1313–14
 - gradual flow of 1487, 1516–20
 - sets, extensive form games 1714
- informed desire criterion 1591, 1599
- informed utility approach 1601
- instrumental position 6–7, 8
- instrumental rationality 963
- instrumental reciprocity 342, 510
- insula
 - risky decisions 1665
 - social preferences 1669
 - guilt 1686
 - human virtues 1685
 - punishment and trust 1675, 1676, 1683–4
 - time preferences 1688
- insurance
 - ambiguity aversion 301
 - deductibles 263
 - for low probability events 193–4
 - probabilistic 310–11
- intention–action divide 690, 1437
- intention-based reciprocity, models of 398, 433
 - evidence on 434–44
- intentions
 - axiomatic foundations 451
 - Charness–Rabin hybrid model 444, 445, 446
 - extrinsic and intrinsic motivation 533–40
 - psychological game theory 946
 - reciprocity theories 233
- interdependence, present-biased time
 - preferences 677–9
- interdependent preferences, models of 398, 447–9
- intergenerational advice and coordination 717, 780–2
- internal validity of experiments 11–13
- intertemporal choice 1295
- intertemporal risk aversion 115
- intertemporal substitution of consumption 663–4
- interval discount function 619–20
- intransitive preferences 635–7, 641–3
- intrinsic motivation
 - asymmetric information in principal–agent framework 507
 - empirical evidence 482–3
 - employment contracts 522
 - in experiments 16–18, 19
 - gift exchange game 364
 - theoretical framework 481
- intrinsic reciprocity 342
 - gift exchange game 357–8
 - long-term contracts 510
 - public goods games 378
- intuitive equilibrium, in signaling games 820, 821, 822, 827
- invasiveness, brain measurement techniques 1656, 1657, 1658
- inverse S-shaped weighting functions, implications of 216, 263–4
- investment goods 585, 646, 683–7, 688
- investor sentiment 1485, 1502
 - efficient markets hypothesis 1491–2, 1493
 - see also* noise traders
- ion channels 1655
- iterated deletion
 - of dominated strategies 713–14, 724–37, 1706
 - of strictly dominated strategies 1168
- iterated regret minimization 988
- January effect, and efficient markets hypothesis 1498
- job search 689
- judgment heuristics 1339–40, 1342–50
 - affect heuristic 1367–9
 - anchoring and adjustment 1370–5
 - attribute substitution 1397–8
 - availability heuristic 1365–7
 - base rate neglect and conservatism 1375–83
 - behavioral welfare economics 1579–80
 - biases in experts 1449–50
 - coarse thinking and persuasion 1401–7
 - confirmation bias 1389–91
 - conjunction fallacy 1361–4
 - dual process models 1398–401
 - efficient markets hypothesis 1491–4
 - false consensus effect 1396–7
 - hindsight bias 1383–9
 - law of small numbers 1350–61
 - limited attention 1463–48
 - mental models 1407–9
 - necessary and sufficient conditions, confusion between 1397
 - rationality debate 1426–31
 - regression towards the mean 1395–6
 - shrouded attributes 1431–6
 - Simon's approach 1409–11
- justice sensitivity (JS) 354
- k-rationalizability 914
- Karmakar probability weighting function 122
- kin selection 1039, 1073–4, 1076
- kindness
 - defined 939
 - fairness and reciprocity 939–44
 - fairness, reciprocity, and inequity aversion 945–7
- Köszegi–Rabin framework 80, 112–13
 - moral hazard, loss aversion, and optimal contracts 273, 277–8
- prospect theory
 - choice acclimating personal equilibrium 162, 163
 - stochastic reference points under rational expectations 153–9, 164–5, 167–72
 - third generation 161
- psychological game theory
 - guilt 956

- plans, dependence of utility on 960
- renegotiation, long-term contracts, and loss aversion 279, 280
- lab evidence, representativeness 348–9, 386–95
- label salience 877, 904–7, 978
- labor economics 689
- laboratory evidence, representativeness 348–9, 386–95
- language
 - evolutionary game theory 1078
 - gene–culture coevolution 1064
 - see also* communication
- large stakes
 - Rabin's paradox 104–6
 - loss aversion 156
- law, rule of 382–3, 384
- law of effect 1092
 - reinforcement learning 1098, 1099, 1121, 1122
 - rule-based learning 1144
- law of large numbers 45, 1350
- law of small numbers 1343, 1350–61
 - cognitive miser model 1399
 - dual process model of the brain 1431
 - efficient markets hypothesis 1491
 - experts 56, 1449
 - formal model 1343, 1358–61
 - heuristics and biases program 45
- learning 1033–6
 - adaptive 827, 1094
 - anticipatory 1130
 - dynamics 821–2
 - evolutionary game theory 1037–91
 - heterogeneity in 1131
 - and hindsight bias, distinction between 1383
 - imitation 1087–9, 1114, 1115
 - incentives 556–8
 - models 1034–5, 1092–157
 - ultimatum game 351
 - viewpoints 1095, 1131–2
 - naïve reinforcement 1104
 - to predict 1156
 - prepared 1065
 - quantal response 1116
 - rational 1097, 1153–7
 - social 1040, 1085–91
 - soft paternalism 1600
 - stimulus–response 1098
 - stochastic social dynamics 1158–204
 - transference of 1146
 - trial and error 1114, 1115
 - see also* belief-based learning; experience weighted attraction (EWA) learning; reinforcement learning; rule-based learning
- learning direction theories 1095–6, 1114, 1115, 1132–42
- left-digit bias 1440–3
- legal situations
 - anchoring 1372–3
 - hindsight bias 1386–7
- leisure goods 646, 683–4, 687–8
- Leontief preferences 356
- less is more effect 1422
- level-*k* models 41–4, 876, 894–6
 - applications 877–8, 904
 - coordination and communication with boundedly rational players 907–11
 - focal points 904–7
 - market entry game 914–16
 - risk dominance and payoff dominance in sender–receiver games 911–14
 - and cognitive hierarchy model, comparison between 900–1
 - comparison with team reasoning and Stackelberg reasoning 987–8
 - empirical evidence 877–8, 916–21
 - evidential equilibrium 978, 980
 - heuristics and biases approach 977
 - non-strategic players 921–5
 - p-beauty contest 896–900, 902
 - prisoner's dilemma game 978
 - subjective heterogeneous quantal response equilibrium 969, 971
 - truncated quantal response equilibrium 971
 - winner's curse 998–1002
- LEX heuristic 1423, 1425
- LEXSEMI heuristic 1423
- libertarian paternalism 1580, 1581, 1586, 1589–90, 1591–3, 1599
- libertarianism 1577
- life-cycle consumption
 - under anxiety 1247–50
 - projection bias 1253–6
- life satisfaction
 - and GDP, relation between 1274–83
 - and happiness, distinction between 1280–1
 - see also* happiness economics
- light paternalism 1580, 1581, 1595
- Likert scale 1674
- limbic system 1649, 1653, 1687–9
- limit entry pricing game 721, 822–7
- limited attention 1350, 1433
 - asymmetric paternalism 1594
 - behavioral welfare economics 1577, 1582–3, 1618–23
 - biased numerical attention 1440–3
 - excess burden of taxation under 1583, 1628–32
 - experts 1449
 - information flows in financial markets 1518
 - and the poor 1436–8
 - and taxes 1438–40, 1583, 1625–8
 - theoretical framework 1443–8
- Linda problem 1361–4, 1400–1
- Loewenstein–Prelec explanation of magnitude effect 605, 607–9
- logit quantal response equilibrium 970
 - extensive form games 888–9
 - normal form games 883–5, 886
 - traveller's dilemma game 893–4
- Long Term Capital Management (LTCM) fund 1516
- long-term contracts 480–1, 510–27
- lookups, inferring cognition through 720–1, 815–19
- loss aversion
 - coefficient of 132, 151
 - endowment effect 220
 - moral hazard and optimal contracts 273

- loss aversion (*cont.*)
 - third generation 153
 - and coordination 717, 775–8
 - disposition effect 54
 - emotions and cognition, interaction between 1292, 1296–7
 - happiness economics 1274
 - learning direction models 1136–7, 1138
 - mental accounting 1465
 - behavioral economics of retirement savings 1483–4
 - prospect theory 1457–8
 - myopic 81, 215, 239–42
 - neuroeconomics 1646, 1665–7
 - prospect theory 26–7, 29, 80, 131–2, 142, 213, 264
 - choice acclimating personal equilibrium 163–4
 - contract renegotiation 216, 278–83
 - elicitation 143–5
 - endowment effect 214, 217, 220–9, 232, 233
 - goals 255
 - Kőszegi–Rabin framework 169, 170
 - labor supply decisions 260
 - market experience 224–9
 - mental accounting 1457–8
 - moral hazard and optimal contracts 216, 272–8
 - power form of utility 151, 152
 - primates 237
 - public goods games 376
 - Rabin's paradox 248
 - stochastic reference points under rational expectations 156
 - third generation 153
 - trade-off attribute model of intertemporal choice 631
 - rank dependent utility theory 130
 - reference time theory 633
 - and risk aversion 81
 - simple increasing elasticity class of utility functions 610
 - soft paternalism 1596, 1597
 - status-quo bias 1596
 - survey data 19
 - willingness to accept and willingness to pay 81
- lotteries
 - binary lottery games 790, 795
 - classical decision theory 84
 - coalescing 173
 - compound *see* compound lotteries
 - event contingent 91, 92
 - in incremental form/prospects 131, 132, 134
 - probability contingent 91, 92
 - random lottery incentive mechanism (RLI) 18, 59
 - set of 133
 - simple 85, 211
 - split 173, 193
 - temporal 1247, 1248
- lottery winners, Easterlin paradox 1276–7, 1278
- low bandwidth problem 1437
- low probability events 81, 114–15, 193–6
 - insurance 193–4
 - neglect by animals 104
 - see also* extreme probability events
- lying
 - evidence on 455, 456–8, 459–60
 - neoclassical economics 453, 455
- macroeconomics 723, 856–63
- magnetic resonance imaging (MRI) 1658, 1683
- magnitude effect
 - exponential discounted utility model 37
 - time discounting 605, 607–11
 - trade-off attribute model of intertemporal choice 631
- managers
 - biases 1488, 1527–9
 - overconfidence 1488, 1527–9, 1534–40
 - rationality 1488, 1529–34
- marginal propensity to consume (MPC)
 - asset specificity 658
 - behavioral life-cycle model 1460–1, 1462, 1463
 - covariation with wealth and age 664
- market entry game
 - asymmetric information 721, 822–7
 - experience weighted attraction learning 1131–2
 - level-*k* models 914–16
 - Nash equilibrium 41
- market improvement rule 868–9, 874
- market integration 350
- market timing hypothesis, merger/acquisition activity 1533–4
- market to book value ratios (MBVR) 1497–8
- Marshall–Machina probability triangle *see* probability triangle
- matching pennies game
 - learning direction models 1137
 - mixed strategy Nash equilibria 743–5
 - quantal response equilibrium 884–5
 - scope of economics 1607
- material emotions 1223
- maximin expected utility model (MEU) 285, 292–3
 - empirical evidence 299
 - and smooth ambiguity aversion models, unified account of 298
- maximin play 764
- maximin preferences
 - Charness–Rabin hybrid model 447
 - social preferences models, evidence on 433
- mean quadratic distance, learning direction models 1137–8
- mean-variance utility analysis 1663
- medial prefrontal cortex 1687, 1689, 1690, 1691
- median action game 716, 763–5
 - communication and experience 784
 - functional experience weighted attraction learning 1127, 1128
 - group size, gradual growth in 779, 780
 - historical accidents 772
 - learning direction theories 1141, 1142
 - loss aversion 775–6
- median enhancing emulation dynamics 1087
- memory
 - availability heuristic 1365
 - emotions and cognition, interaction between 1311–12, 1316–17

- hippocampus 1653
- stochastic social dynamics 1166–7, 1169, 1175–6, 1178
- mental accounting 1340, 1451–5
 - account closure 1477
 - behavioral economics of retirement savings 1482–4
 - behavioral life-cycle model 1459–63
 - choice bracketing 1469–71
 - coherent arbitrariness 1471–3
 - diversification heuristic 1481–2
 - examples 1451–3
 - framing 1455–6
 - of goals 1482
 - house money effect 1479–81
 - independence axiom of EU, violations 97
 - myopic loss aversion 239
 - prospect theory 1456–9
 - red and black of 1454, 1463–9
 - self-control and gift giving 1477–9
 - sunk costs 1473–6
 - transaction utility & acquisition utility 1476–7
- mental models 1347, 1407–9
 - cognitive dissonance 1399
- mentalizing *see* mind, theory of
- merger/acquisition activity 1528, 1529–34
- methodological individualism 962, 963
- methodology in economics 4, 5–10
- microcredit 691
- mind, theory of 1647, 1678–9
 - rule-based learning 1145–6
 - strategic interaction 1690, 1691
- mindful economics 1577
- mindless economics 1577, 1605
- minimal groups, social identity 466, 469, 470–3
- minimalist heuristic 1420, 1424
- minimax play, mixed strategy Nash equilibria 738–43, 744, 749
 - professional sports players 752–4, 755, 756, 757, 758
- multilevel selection 1037, 1039, 1075–6
- minimum effort games 473, 779–80
- minimum effort weak link games *see* weak link minimum effort games
- minimum savings rules 1269
- minimum wage 363
- mixed strategy Nash equilibrium (MSE) 57–9, 714–15, 738, 748–51
 - battle of the sexes game 1708–9
 - in economic context 745–8
 - in general context 738–45
 - matching pennies game 884, 885
 - mixed strategy devices 748–51
 - professionals 751–8
 - two-person zero-sum games 1710, 1711
- mixture model approach, prospect theory 179–81
- mixture symmetry 103
- mobile phone use while driving 195
- modified public goods (MPG) game 389–90
- modularity of the brain 1648
- molecular genetics 1033–4
- momentum
 - defined 1519
 - efficient markets hypothesis 1498
 - financial markets 55
 - information flows in financial markets 1519–20
- monetary incentives *see* financial incentives
- money illusion 233, 856–63
- monomorphic population 1037, 1047
- monotone likelihood ratio property (MLRP)
 - contracts 273, 274, 276
 - incentives 491, 493, 528
- monotonicity
 - reference time theory 633, 642
 - subjective expected utility theory 90
- Monty Hall game 922, 923
- moonlighting game 436–8
- moral disengagement 482, 548–52
- moral hazard
 - employer 523
 - evolutionary game theory 1084
 - guilt 952–3
 - incentives 478–9, 483
 - inequity aversion 483–90
 - reciprocity 490–5
 - loss aversion and optimal contracts 216, 272–8
 - present-biased preferences 677
 - soft paternalism 1600
 - trust game 369
- morality
 - evolutionary game theory 1078, 1081–4
 - signaling 482, 552–5
- more is less effect 1423
- mortgage payments 1468
- motivation
 - emotions and cognition, interaction between 1311–21
 - narrow choice bracketing 1471
- MOUSELAB
 - level-*k* models 916–18, 920
 - non-choice data 1607
 - search and lookups, inferring cognition through 816, 817–19
- multilevel selection 1037, 1039, 1075–6
- multiple games, learning models 1097, 1148–53
- multiple priors models 82
 - ambiguity 285, 291–3, 299
- multiple regression and heuristics, performance comparison 1424–5
- multiple selves models 1299
 - choice-based behavioral welfare economics 1617
 - emotions and cognition, interaction between 1292, 1308–21
 - Euler equation 660–2
 - present-biased preferences 644–5, 646–7, 650–1
 - procrastination 1598
 - temptation and conflicts between commitment and flexibility 1266–7
- multiple signals 722, 840–3
- multiself Pareto dominance 1617
- multitask agency models 478
- mutations
 - evolutionary stable strategies 1043, 1044–5, 1046
 - stochastic social dynamics 1158, 1160
 - coordination game 1162, 1163
 - perturbed Markov dynamics 1168–70
 - technology adoption model 1165
- mutually consistent strategies 981

- myopic best response models 772–4, 1142
- myopic loss aversion (MLA) 81, 215, 239–42
- naive hyperbolic decision makers 617–18
- naive reinforcement learning 1104
- narcissism, and merger activities 1529
- narrow choice bracketing 1469–71, 1480, 1494
- Nash bargaining approach 718, 786–93
- natural frequency format 1362, 1378
- natural selection *see* evolution
- necessary and sufficient conditions, confusion between 1347, 1397
- negative autocorrelation, heuristics and biases program 45
- negotiations research 718
- neoclassical economics 1
 - ambiguity 81–2, 284, 285, 286, 288–301
 - amorality 33
 - and behavioral economics, relation between 2
 - bygones 52
 - cash vs. in-kind transfers 1478
 - contracts 266–7
 - experiments 13, 14
 - and expert behavior 56–7
 - frame-invariance assumption 20
 - framing effects 1456
 - and heuristics and biases program 44, 45
 - incentives 539, 560
 - liberal tradition 1579
 - marginalist principle 1473
 - other-regarding preferences 32, 341
 - paternalism 1579, 1585, 1590
 - preferences 1033
 - promises and threats 454, 461
 - prosocial behavior 1064
 - rational inattention models 1443
 - rationality 30, 1431
 - scope 1604, 1605, 1609, 1610
 - selfish black liars 453, 455, 458, 1084
 - welfare analysis 1298, 1611
- neuroeconomics 22, 1643–8
 - brain 1648–55
 - choice-based welfare economics 1612
 - defined 1644
 - pharmacoeconomics 1643, 1692–4
 - prospect theory 138
 - of risky decisions 1661–9
 - scope of economics 1604, 1606
 - of social preferences 1646–7, 1669–86
 - of strategic interaction 1689–92
 - techniques 1643, 1645, 1655–61
 - manipulation 1655–6, 1659–61
 - measurement 1655, 1656–8
 - of time preferences 1686–9
 - ultimatum game 351
- never-a-weak-best response (NWBR) criterion, signaling games 822
- newsvendor problem, impulse balance equilibrium 1136
- niche construction 1066
- no-cues model 1238–40
- no-trade theorems 843
- noise traders 54–5, 1485
 - arbitrage limits 1510, 1511, 1512, 1513, 1514, 1515
 - efficient markets hypothesis 1491–2, 1493, 1495
 - risk 1486–7, 1502–9
 - see also* investor sentiment
- noisy introspection (NI) 881, 984–6
 - p-beauty contest 900
- noisy Nash model (NNM) 887–8
- nominal inertia, macroeconomics of 723, 856–63
- non-choice data 1581–2, 1603–10
- non-linear probability weighting
 - anchoring 1373
 - neuroeconomics 1646, 1667–9
 - prospect theory 27
- non-strategic players 878, 921–5
- non-transitive choices 601–2
- norm theory, professional sports players 754
- normative bargaining solutions 718, 786–93
- normative preferences 1579
- norms
 - and mental models, distinction between 1409
 - social *see* social norms
 - stochastic social dominance 1178–80
- NP-complete problems 47
- nucleus accumbens 1663, 1686, 1687
- nudges 1438, 1577, 1581, 1586, 1589–90, 1597–8
 - behavioral public finance 1633
 - contractarian approach 1624, 1625
 - critique 1599, 1600
 - libertarian paternalism 1592–3
- numerical attention, biased 1440–3
- obesity 688–9, 1597–8
- optimal contracts 268–70
 - incentives
 - incomplete contracts 498–505
 - moral hazard 485, 486, 487, 492–3, 495
 - reciprocity and long-term contracts 515
 - moral hazard and loss aversion 216, 272–8
- optimism of decision maker 125–7, 129
- orbitofrontal cortex (OFC) 1653
 - punishment and trust 1675, 1684
 - risky decision 1665
 - time preferences 1687
 - value 1663
- order statistic games
 - learning direction theories 1096, 1138–42
 - reinforcement learning 1106
- original prospect theory (OPT) 28, 80
 - ambiguity 299
 - and cumulative prospect theory, difference between 111
 - extreme probability events 191, 192–3
 - limitations 173
 - loss aversion 145
 - see also* prospect theory
- other-regarding preferences 339–43
 - competitive market equilibrium 49
 - cultural evolution model 1073
 - human sociality 32, 34, 35, 36, 344–97
 - human virtues 453–4, 455–61
 - incentives 478–561

- iterated deletion of dominated strategies 727, 731
- leader contributions to charitable giving 835
- models 398–452
 - neoclassical economics 32
 - social identity 454–5, 465–6
 - see also* social preferences
- out-of-pocket costs vs. opportunity costs 232–3
- outcome sensitivity, reference time theory 642
- overconfidence
 - confirmation bias 1389–95, 1422
 - consumer 684, 685, 687
 - efficient markets hypothesis 1493, 1494
 - experts 56, 1449
 - formal model 1392–5
 - hindsight bias 1431
 - information flows in financial markets 1519
 - investor skills 1501
 - law of small numbers 1360–1
 - managers 1488, 1527–9, 1534–40
 - noise trader risk 1504
 - take-the-best algorithm 1422
- overinvestment 1539–40
- own-payoff effect 886
- oxytocin, social effects of 1648, 1692–4
- p-beauty contest
 - cognitive hierarchy model 902
 - experience weighted attraction learning 1125, 1128–30
 - functional 1127, 1128
 - level- k models 41, 42–3, 876, 896–900, 902
 - non-strategic players 922–4
 - neuroeconomics of strategic interaction 1690–1
- Pareto dominance, multiset 1617
- Pareto white lies 455, 461
- parochial altruism 1233, 1647, 1677–8
- partial-naivety about self-control 39
- patent race games 1127, 1128
- paternalism 1577, 1590
 - contractarian approach 1624–5
 - direct 1580, 1585, 1587–8
 - heavy-handed 1577, 1579
 - indirect 1580, 1586, 1588–9
 - libertarian 1580, 1581, 1586, 1589–90, 1591–3, 1599
 - light 1580, 1581, 1595
 - see also* asymmetric paternalism; soft paternalism
- path independence, reference time theory 642
- path of least resistance, stochastic social dynamics
 - coordination game 1163
 - memory, conventions, and risk dominance in 3×3 games 1176, 1181
 - perturbed Markov dynamics in 2×2 games 1172
 - technology adoption model 1165–6
- patience *see* impatience and patience
- pay what you want (PWYW) pricing 555
- payday effect 689
- payment depreciation 1474
- payoff dominance
 - belief-based learning 1118
 - Nash equilibria 1711–12
 - normal form games 1710
 - stag-hunt game 767, 768, 877, 911–14
- payoff efficiency index 841
- payoff monotone and regular (PMR)
 - dynamics 1090–1
- payoff monotonicity 1040, 1089–91
- payoff salience 904–7
- payoff-sampling learning 1136
- peak-end evaluation, happiness economics 1271
- pensions
 - automatic enrollment in 401(k) plans 1594
 - mental accounting
 - behavioral economics of retirement savings 1455, 1482–4
 - behavioral life-cycle model 1462
 - diversification heuristic 1481
 - see also* retirement
- perception-perfect strategies 670–1, 676
- performance based arbitrage (PBA) 1512, 1513, 1514, 1515
- perfunctionary performance 264–5
- persistent randomness, sources of 1160
- personal commitment, managers 1528
- personal evaluation biases 1320–1
- personal rules and willpower 1321
- persuasion 1347, 1401–7
- perturbed Markov dynamics 1158–9, 1166–74
- pessimism of decision maker 125–6, 129
- pharmacoeconomics 1643, 1648, 1692–4
- phenotypes 1037
 - evolutionary stable strategies 1043, 1044–5
 - gene–culture coevolution 1066
 - resident 1037
- phobias 1230, 1234
- Pigouvian taxes 1579
 - indirect paternalism 1588, 1589
 - present-biased preferences 646, 679, 681
- planner–doer models 644, 646, 647–9
 - neuroeconomics of time preferences 1689
 - temptation preferences 1261, 1264–6
- plans, dependence of utility on 960–2
- plasticity of the brain 1648
- playing the field 1038, 1063
- policy view of experiments 14
- political economy
 - behavioral 400, 417–26
 - evolutionary 1067–8
- polymorphic population 1037, 1047
- poor people, and limited attention 1436–8
- population game 1038, 1063
- population genetics 1049
- portability, level- k models 895
- portfolio choice under anxiety 1247–50
- positive autocorrelation, heuristics and biases
 - program 45
- positive bargaining solutions 718–20, 793–811
- positive feedback traders 1492
- post-earnings-news drift, and efficient markets
 - hypothesis 1498
- power form of utility 112, 142, 144
 - axiomatic foundations 151–2
- power invariance, Prelec function 150–1
- precautionary savings, hyperbolic discounting 662–3
- predatory pricing 830

- predicted utility 1271
- prediction error (PE)
 - neuroeconomics 1663, 1665, 1669
 - rational learning 1156
- preference reversals 107–8
 - anchoring 1373
 - exponential discounted utility model 38, 594–5
 - salience and context in decision theory 186
 - survey data 20–1
 - temporal 37
 - third generation prospect theory 153
- preferred personal equilibrium (PPE) 159, 167
- prefrontal cortex (PFC) 1653, 1661
 - evolution 1290, 1649
 - risky decisions 1664–5, 1667, 1668
 - social preferences 1669, 1670, 1672, 1673
 - guilt 1686
 - punishment and trust 1676, 1678, 1679, 1680–3
 - strategic interaction 1689–90, 1691, 1692
 - time preferences 1686, 1687, 1688, 1689
 - value 1662
- Prelec probability weighting function 118–23
 - composite 196–8
 - neuroeconomics 1667–9
 - prospect theory 27–8, 29, 32, 133, 134, 146
 - ambiguity 307, 308
 - axiomatic foundations 150–1
 - risk 135, 137
 - tax evasion puzzles 247
 - rank dependent utility theory 126–7
 - risk and time preferences, integration of 204
 - delay dependence 209–10
 - intertemporal risk aversion 206
 - process dependence 211–12
 - salience and context in decision theory 184
- prepared learning 1065
- prepayment preference 1464, 1466–8
- preplay communication 716–18, 771–2
 - level-*k* models 907, 912–14
 - trust game 928
- preproperation 669–77
- prescriptive methodology 5
- present-biased preferences
 - applications 644–6
 - endogenous retirement decisions 664–9
 - hyperbolic discounting 651–64, 688–91
 - interdependence, multiple activities, and competition 677–9
 - investment and leisure goods 683–4
 - modeling frameworks 646–51
 - optimal sin taxes 679–83
 - procrastination and preproperation 669–77
 - exponential discounted utility model 37
 - hyperbolic discounting 618
 - quasi-hyperbolic discounting 39, 40, 605
 - trade-off attribute model of intertemporal choice 631, 632
- Price equation 1075
- pride 414–15
- primates
 - brain 1649, 1666, 1686
 - cooperation 342
 - increasing sequences, preferences for 603
 - loss aversion 1666
 - prospect theory 215, 235–9, 1666
 - time preferences 1686
- priming
 - anchoring 1374
 - social identity 466
 - evidence on 467, 468, 469, 473
- prisoner's dilemma (PD) game
 - antisocial punishments 385
 - aspiration adaptation 1415
 - democracy 375
 - evidential equilibrium 977–8
 - evidential reasoning 982–4
 - evolutionary game theory 1042
 - group selection 1075–6
 - inclusive fitness 1074
 - indirect reciprocity 1078–9
 - monomorphic equilibrium 1053–7
 - network reciprocity 1084
 - reciprocal altruism 1076, 1077–8
 - replicator dynamics 1050
 - experience weighted attraction learning 1131
 - fairness and reciprocity 937, 943–5
 - future interaction with same partner 521
 - group size, gradual growth in 779
 - incentives
 - crowding-in effect 559, 560
 - extrinsic and intrinsic motivation 558, 559, 560
 - learning direction theories 1133–4
 - multiple signals 840–1, 842
 - neuroeconomics 1694
 - punishment and trust 1676, 1677–8
 - reciprocity 342
 - social identity 469–70, 473
- probabilistic insurance 310–11
- probabilistic sophistication
 - Ellsberg paradox 285, 288
 - rank dependent utility theory 129
 - subjective expected utility theory 91, 92
- probability contingent lotteries 91, 92
- probability triangle
 - disappointment aversion theory 183
 - fanning-in 103–4
 - fanning-out 101–4
 - independence axiom of EU, relaxation of 101–4
 - prospect theory 136–7, 190
 - rationality axioms of EU, violations 99–101
 - regret theory 182
 - transfer of attention in exchange model 190
- probability weighting, non-linear *see* non-linear probability weighting
 - probability weighting
- probability weighting function 115–24
 - configural weights models 32
 - expected utility theory 80
 - prospect theory 27, 31, 32, 133, 134
 - elicitation 146–7
 - risk 135, 136
 - stake size effect 175–9
- rank dependent utility 32, 125
 - curvature 125–6
 - elevation 126–7

- rationality of emotions 1235
- risk and time preferences, integration of 199–205
 - delay dependence 207–10
 - process dependence 210–12
- procedural rationality 1409, 1410, 1416
- procedural utility 1272
- procedure invariance, violations of 108
- process dependence 207, 210–12
- procrastination 645, 669–77, 1598
- productivity incentives 510
- professional sports players, mixed strategy Nash equilibria 752–8
- projection bias 1227–8, 1250–1, 1396
 - applications 1257
 - consumer durables 1256–7
 - habit formation and life-cycle consumption 1253–6
 - preferences under 1251–3
- promises and threats, model of 454, 461–5
- property rights 231–2
- prosocial behavior 341
 - accounting for 1063–4
 - equity, reciprocity, and competition model 411
 - evolution 385
 - evolutionary game theory 1075
 - gene–culture coevolution 1066–7
 - guilt 952, 1685, 1686
 - incentives 554, 555
 - oxytocin 1693, 1694
 - public goods games 374
 - evolution and reciprocity 385
 - external validity of lab evidence 391–2, 393–4, 395
 - punishments 383, 384
 - social identity 470
 - student subject pool 386–7
- prospect theory (PT) 26–32, 48, 80, 111–13, 130–4
 - ambiguity 286, 304–9
 - applications 213–83
 - axiomatic foundations 148–52
 - Becker paradox 194
 - common difference effect 624
 - composite *see* composite prospect theory
 - and configural weights models, comparison of 190
 - disposition effect 1355
 - efficient markets hypothesis 1493
 - erroneous criticisms 138–40
 - and expected utility theory, comparison between 81
 - extreme probability events 193
 - financial markets 56
 - fourth generation 199
 - framing effects 107
 - goals 251–4
 - limitations 172–81
 - loss aversion, elicitation of 143–5
 - low probability events 194, 195, 196
 - and mental accounting 1453, 1456–9
 - choice bracketing 1470–1
 - examples 1452
 - neural evidence 1661–9
 - Prelec function 121
 - probability weighting function, elicitation of 146–7
 - public goods games 376
 - and rank dependent utility theory, difference between 131
 - rationality of emotions 1233
 - and reference time theory 643
 - risk 106, 134–7
 - stochastic reference points under rational expectations 153–72
 - survey data 19–20
 - third generation 80, 112, 152–3, 199
 - uncertainty 134
 - utility function, elicitation of 140–3, 144
- prospective accounting 1464, 1466, 1467–8, 1475
- psychological expected utility model 1245–7
- psychological game
 - with reciprocity incentives 940
- theory 36, 878–9, 925–62
- psychological Nash equilibrium
 - defined 930
 - multi-stage games 931–7
 - single-stage games 928–31
- psychology
 - scope of economics 1604
 - of tail events 56, 81, 1494
- psychophysical activity, measurement of 1656
- public goods (PG) games 344, 346–8, 370
 - contributions 356–7, 975
 - empirical evidence 372–6
 - evolution and reciprocity 385–6
 - evolutionary game theory 1077
 - external validity of lab evidence 389–93
- Fehr–Schmidt model 407
- guilt 950–2
- heterogeneous preferences and the pattern of contributions 376–9
- inequity aversion 434
- modified public goods game 389–90
- psychological game theory 941–2, 943, 950–2
- stochastic social dynamics 1159, 1180–2
- threshold 376
- unproductive expenditures 379–82
- with punishments 371–2
 - altruism 399
 - antisocial punishments 382–5
 - bygones 52–3
 - Charness–Rabin hybrid model 444, 445
 - crowding-in effect of incentives 558–9
 - empirical evidence 372, 373–6
 - equity, reciprocity, and competition model 411
 - evolutionary game theory 1080, 1084
 - Fehr–Schmidt model 408–10
 - heterogeneous preferences and the pattern of contributions 377–9
 - human sociality 33, 35, 36
 - intention-based reciprocity 436
 - neuroeconomics 1673
 - unproductive expenditures on public goods 379–81, 382
- without punishments 371
 - bygones 52–3
 - empirical evidence 372, 373, 374, 375, 376
 - equity, reciprocity, and competition model 411
 - external validity of lab evidence 390

- public goods (PG) games (*cont.*)
 - Fehr–Schmidt model 407–8
 - heterogeneous preferences and the pattern of contributions 377
 - human sociality 33, 36
 - incentives 556–7
 - unproductive expenditures on public goods 381
- public signals and correlated equilibria 722–3, 843–51
- punctuated equilibrium 782, 1160
- punishments 52–4
 - evolutionary game theory 1084
 - functional magnetic resonance imaging 1658
 - human sociality 342
 - incentives 500, 558
 - neural basis 1646–7, 1673–84
 - social identity 471, 472–3
 - see also* public goods (PG) games: with punishments; third party punishment (TPG) game
- purification agenda 57
- pursuer–evader game 745, 754, 755, 756
- putamen 1651, 1669, 1683
- quantal response equilibrium (QRE) 41, 875–6, 881–2
 - belief-based learning 1117
 - defined 883
 - extensive form games 888–90
 - falsifiable restrictions 890–2
 - heterogeneous (HRQE) 880, 969–72
 - normal form games 886, 888
 - hypothesis testing 1157
 - and impulse balance equilibrium, comparison between 1137
 - iterated deletion of dominated strategies 733
 - label salience 906
 - learning models 1127–8
 - noisy introspection 984, 985–6
 - normal form games 882–8
 - p-beauty contest 900
 - positive bargaining solutions 799–800
 - prisoner's dilemma game 978
 - regular 891–2
 - smooth fictitious play 1108
 - sophisticated experience weighted attraction learning 1129
 - subjective heterogeneous 880, 969–72
 - traveller's dilemma game 892–4
 - truncated 971
 - winner's curse 1001, 1002
- quantal response learning 1116
- quasi-hyperbolic discounting 37–40
 - common difference effect 605, 614–16, 626–7
 - emotions and cognition, interaction between 1295
 - empirical evidence 618
 - neuroeconomics of time preferences 1687
 - reference time theory 637
- query theory of value construction 214, 234–5
- quotient dynamics 1050
- Rabin's paradox 81, 104–6, 215, 248–50
- race factors *see* ethnicity factors
- random lottery incentive mechanism (RLI) 18, 59
- random numbers, inability to produce a sequence of 1352
- rank affected multiplicative weights model (RAM) 188–9, 191
- rank changing coalitions 1067
- rank dependent utility (RDU) 28, 29, 30, 31, 80, 110–11, 124–30
 - ambiguity 286, 299
 - two-stage recursive models 294, 295
 - Becker paradox 194
 - binary rank dependent utility 305
 - common difference effect 623, 624
 - drawbacks 129–30
 - goals 254
 - low probability events 194, 195, 196
 - market experience and loss aversion 228–9
 - Prelec function 121
 - probability weight function 32
 - and prospect theory, difference between 131
 - risk 125–8
 - risk and time preferences, integration of 202, 203, 205
 - delay dependence 207–8
 - intertemporal risk aversion 206
 - process dependence 211
 - uncertainty 128–9
 - violation 137–8
- rational expectations 153–72
- rational inattention models 1443
- rational learning 1097, 1153–7
- rationality
 - bounded *see* bounded rationality
 - efficient markets hypothesis 1490
 - emotions 1226, 1231–5, 1290
 - great rationality debate 47–8, 1349, 1426–31
 - interpretations 1409–10
 - investors 1529–40
 - managers 1488, 1529–34, 1537, 1539
 - procedural 1409, 1410, 1416
 - substantive 1409–10
- rationality axioms for time discounting 593
- rationality axioms of expected utility theory 86–7, 99–101, 108
- Read–Scholten (RS) discount function 584
 - common difference effect 606, 621
 - reference time theory 637
- realized emotions 1223
- reasoning-based expected utility procedure 988
- rebound effect 1482
- RECAP regulation 1593, 1602, 1603
- reciprocal altruism 342
 - evolutionary game theory 1039–40, 1076–8
- reciprocal determinism, behavioral welfare economics 1584
- reciprocal fairness equilibrium (RFE) 446
- reciprocal inclination 949
- reciprocation term 947, 948
- reciprocity 342, 343
 - axiomatic foundations 450, 451
 - Charness–Rabin hybrid model 446
 - Cox–Friedman–Gjerstad model 449
 - defined 940

- direct 1076–8
- equilibrium, psychological game theory 947, 948
- evolutionary game theory 1039–40, 1073–84
- external validity of lab evidence 394–5
- gift exchange game 362–3, 364, 365–6
- incentives
 - extrinsic and intrinsic motivation 533–4, 535, 536, 540–1, 546, 550
 - incomplete contracts 497, 498, 508, 509, 510
- instrumental 342, 510
- intention-based 398, 433
 - evidence on 434–44
- leader contributions to charitable giving 835–9
- long-term contracts 480–1, 510–27
- models 401, 434–44
- moral hazard 479, 490–5
- psychological game theory
 - extensive form games 879, 937–49
 - gift exchange game 928
 - guilt 952
- public goods games 377–8, 385–6
- rule-based learning 1148
- sequential reciprocity equilibrium 879, 941–5
- social identity 469
- strategic complements and strategic substitutes 855
- strong 342–3
- student subject pool 386
- and tenure, relation between 362–3
- theories 233
 - see also* indirect reciprocity; intrinsic reciprocity
- recognition heuristic 1419–26
- red teaming 1450
- red traffic lights, jumping 195
- reduction axiom of expected utility theory 86
 - evidence 108–9
 - violations 97–8
- reduction invariance, Prelec function 150, 151
- reference dependence 26
 - expected utility theory 250
 - gift exchange game 363–4
 - mental accounting examples 1452
 - neuroeconomics 1646, 1663–5
 - prospect theory 130, 131, 213
 - contract renegotiation 216
 - contracts as reference points 264
 - labor supply decisions 257, 259–60, 262
 - primates 237
 - Rabin's paradox 248, 250
 - rank dependent utility theory 130
 - reference time theory 633
- reference dependent subjective expected utility theory (RDSEU) 153
- reference points
 - composite prospect theory 199
 - contracts as 216, 264–72
 - fairness 439–44
 - gift exchange game 362
 - goals as 215–16, 250–5
 - prospect theory 130–1, 137–8
 - criticisms 139–40
 - framing effects 107
 - primates 237
 - renegotiation, long-term contracts, and loss aversion 283
 - trade-off consistency method 148, 150
 - rank dependent utility theory 130
 - stochastic *see* stochastic reference points
 - strategic complements and strategic substitutes 858–9
- reference time theory (RT) 584, 606–7, 632–3
 - additivity of discount functions, implications of 634–5
 - basic description 633–4
 - common discount functions recast in terms of 637
 - exponential discounted utility model anomalies explained 638–41
 - intransitive preferences 635–7
 - and relative discounting 641–3
- regions of interest (ROI) analysis 1678, 1679
- regression towards the mean 1346–7, 1395–6
- regret
 - iterated regret minimization 988
 - learning 1153, 1154
 - matching 1153–6
 - theory 81, 113–14, 181–2
- regular quantal response equilibrium 891–2
- regulations
 - anchoring 1372
 - imperfect self-information 1581, 1601–3
 - soft paternalism 1599
- reinforcement learning 1092–3, 1097–8, 1132, 1137, 1138, 1426
 - empirical evidence 1093, 1103–6
 - experience weighted attraction model 1094, 1120–1, 1122–3
 - vs. fictitious play 1094, 1115–20
 - and learning direction theories, comparison between 1133
 - naive 1104
 - one-parameter model 1092, 1098–100
 - optimality in stationary environment 1092, 1102–3
 - with payoff variability 1127–8
 - soft learning 1600
 - three-parameter model 1092, 1100
 - uncoupled learning rule 1156
 - variants 1092, 1101–2
- relative discounting 601, 641–3
- relative performance contracts 488, 489, 490
- relative risk premium (RRP) 116–17
- relativity, theory of 7–8
- reluctant altruism 345, 353, 354
- remembered utility 1271
- rent dissipation 381, 382, 383
- repeated games
 - coordination in 785–6
 - evolutionary game theory 1040, 1079
- repetitive transcranial magnetic stimulation (rTMS) 1659, 1679–83
- replicator dynamics (RD) 1038, 1048–53
 - asymptotically stable equilibrium, non-existence of 1060
 - chicken game 1057
 - cultural evolution 1071
 - evolutionary game theory 1040

- replicator dynamics (RD) (*cont.*)
 - microfoundations 1040, 1087–9
 - prisoner's dilemma game 1054
 - rule-based learning 1145
 - stochastic social dynamics 1161, 1162
 - and weighted fictitious play, comparison
 - between 1112–13
- representativeness heuristic 45, 1342–3, 1351–2, 1357
 - conjunction fallacy 364
 - stock market underreaction and overreaction 1520
- reputation
 - evolutionary game theory 1078–9
 - human virtues 1684
 - incentives 531–3, 546, 547, 555
 - indirect reciprocity 395, 396, 1078–9
 - neuroeconomics 1669
 - other-regarding preferences 392
 - reciprocity and long-term contracts 510, 512, 513, 521, 523, 525, 527
 - social preferences 1681
 - sophisticated experience weighted attraction
 - learning 1130
 - ultimatum game 351
- resource constraint, sin taxes 680
- restricted branch independence (RBI) 173, 174
- retirement 584–5
 - consumption at 658
 - decisions 645, 664–9
 - savings
 - oversaving 669
 - soft paternalism 1596–7
 - undersaving 657–8, 669
 - see also* pensions
- revealed preferences
 - behavioral welfare economics 1577, 1582–3, 1618–23
 - dictator game 355
 - generalized axiom *see* generalized axiom of revealed preference
 - under limited attention 1577, 1582–3, 1618–23
 - neoclassical economics 1579
 - self-control 1270
 - temptation 1258–66
 - weak axiom of revealed preference with limited
 - attention 1622, 1623
- revenge 384
- revenue equivalence theorem 990
- reverse common ratio violation of EU independence
 - axiom 97
- rewards
 - anticipation 1667, 1668
 - human sociality 342
 - neuroeconomics 1646, 1667, 1668, 1669–70
 - punishment and trust 1674, 1675
 - social 1646, 1669–70
- risk 83
 - and ambiguity attitudes 300–1
 - behavior under 1295–7
 - behavioral decision theory 79
 - bimodal perception of 191
 - classical decision theory 85
 - dominance 877, 911–14
 - emotions and cognition, interaction
 - between 1295–7
 - expected utility theory 31, 80, 87–9
 - Rabin's paradox 104–6
 - as feelings 1226, 1234–5, 1368
 - fundamental *see* fundamental risk
 - house money effect 1480
 - idiosyncratic 1489, 1490
 - and intertemporal substitution, relation
 - between 663–4
 - noise traders 1486–7, 1502–9
 - normative bargaining solutions 792
 - original prospect theory 80
 - probability weighting functions 116–17
 - prospect theory 31, 80, 131, 134–7, 176
 - trade-off consistency method 148–50
 - rank dependent utility 80, 125–8
 - subjective expected utility theory 91
 - temporal behavior under 199–205
- risk aversion
 - and anxiety, distinction between 1250
 - arbitrageurs 1505, 1508
 - classical decision theory 87–8
 - and cognitive load, relation between 1303
 - expected utility theory 104–5
 - gender factors 137, 181, 1235
 - incentives 509–10
 - intertemporal 205–6
 - and loss aversion 81
 - managers 1536, 1539
 - mental accounting 1469–70
 - neuroeconomics 1666
 - noise traders 1505, 1507, 1508
 - probability weighting functions 117
 - prospect theory 29, 131, 135, 136, 137
 - endowment effect 214, 222–4
 - inverse S-shaped weighting functions 263
 - stake size effect 175
 - tax evasion puzzles 243, 245
 - rank dependent utility theory 125, 126
 - stock market participation 1494
 - underlying mechanisms 212
- risk dominance
 - coordination games 767, 768
 - Nash equilibria 1711–12
 - stochastic social dynamics 1174–8, 1179–80
 - perturbed Markov dynamics 1170–2
- risk loving/seeking
 - classical decision theory 87–8
 - probability weighting functions 117
 - prospect theory 131, 135, 136, 137
- risk preferences
 - mental accounting 1472
 - and time preferences, relation between 115, 199–212, 619, 626
- risky decisions, neuroeconomics of 1645–6, 1661–9
- risky willingness to accept (RWTA) 222, 223, 224
- risky willingness to pay (RWTP) 223, 224
- Roberts–Romer–Meltzer–Richards (RRMR)
 - model 418
- rock–scissors–paper game 1059–60, 1061
- rote learning *see* reinforcement learning

- rule-based learning 1096, 1142–3
 - evidence 1096, 1147–8
 - model 1096, 1143–6
- rule-based reasoning 7
- rule of law 382–3, 384
- sales contracts 480, 521–7
- salience 81
 - availability heuristic 1365
 - base rates 1375, 1376
 - decision theory 114, 184–7
 - information flows in financial markets 1517
 - label 877, 904–7
 - prisoner's dilemma game 978
 - payoff 904–7
 - primary vs. secondary 759–60
 - taxes 1438–40
 - see also* focal points
- satisficing strategy (SAT) 1348, 1412, 1425
 - anchoring 1373
 - and fast and frugal heuristics, distinction between 1419
 - heuristics and biases program 47
- save more tomorrow (SMarT) plan 1484, 1596–7
- savings
 - declining national rates 658
 - inadequacy 649
 - mental accounting
 - behavioral economics of retirement savings 1455, 1482–4
 - behavioral life-cycle model 1462
 - precautionary 662–3
 - and retirement decisions 657–8, 664–9
- scientific view of experiments 14
- scope of economics 1581–2, 1603–10
- sealed-bid auctions
 - first price 989–90, 992, 993, 994
 - winner's curse 998–9, 1002
 - positive bargaining solutions 806, 810–11
 - second price 989–90, 992–3
 - winner's curse 997, 1001, 1002
- searches, inferring cognition through 720–1, 815–19
- seat belt usage 195
- second order beliefs, models of 82
- second order Fehr–Schmidt dominance 416
- second order stochastic dominance
 - classical decision theory 88
 - defined 413
 - and fairness 412, 416, 417
- second price auctions 989–90, 992–3
 - incentive compatibility 18
 - winner's curse 997, 1001, 1002
- self-awareness *see* awareness
- self-confidence *see* confidence
- self-control 1261–3, 1270
 - mental accounting 1477–9
 - problems 37–40, 655–8
- self-deception 1312, 1317
- self-esteem 1311–21
- self-handicapping 1316
- self-information, regulation under imperfect 1581, 1601–3
- self-regarding preferences 339–42
 - bygones 52, 53
 - children 354
 - competitive market equilibrium 49
 - contracts 271, 272
 - dictator game 352, 356
 - evolution 385
 - fairness and stochastic dominance 412
 - fairness, general equilibrium, and welfare 429
 - Fehr–Schmidt model 405, 406, 407
 - gift exchange game 357, 358, 360
 - human sociality 32, 33, 34, 35
 - incentives 478
 - extrinsic and intrinsic motivation 539, 541, 546, 555
 - incomplete contracts 507, 508
 - moral hazard 484–5, 487, 489–90, 492, 494
 - optimal contracts 498, 499, 501, 502–3, 504, 505
 - reciprocity and long-term contracts 511–18, 524–7
 - intention-based reciprocity 437
 - political economy 417, 418, 419, 422–6
 - promises and threats 461–2
 - public goods games 377
 - empirical evidence 373
 - unproductive expenditures 381
 - with punishments 372, 373
 - without punishments 371, 373
 - trust game 366
 - ultimatum game 349, 351
- self-serving bias 33
 - in bargaining 720, 811–15
 - contracts 264
 - managers 1527, 1528
- self-signaling 528–33, 553
- self-tuning experience weighted attraction learning (functional EWA learning) 1095, 1125–8, 1137, 1138
- selfish black lies
 - evidence on 456, 459, 461
 - evolutionary game theory 1084
 - neoclassical economics 453, 455, 458
- sensation seeking 1502
- sense making 1648
- sensory areas of the brain 1653
- separability
 - axiom 451–2
 - preferences 428
 - reference time theory 642
- sequential bargaining game
 - iterated deletion of dominated strategies 726
 - positive bargaining solutions 793, 801
 - search and lookups, inferring cognition through 816
- sequential equilibrium
 - reciprocity 879, 941–5
 - signaling games 821, 822, 826
 - finitely repeated asymmetric games 830–1
 - leader contributions to charitable giving 836
- set point theory of happiness 1278
- shading of performance 265, 266, 267, 271, 272
- short selling 1485, 1488–90

- shrouded attributes 1350, 1431–6, 1437, 1629
- Siamese-twin companies 55, 1498–9
- sign effect
 - exponential discounted utility model 37
 - time discounting 605, 607–11
 - trade-off attribute model of intertemporal choice 631, 632
- signaling
 - behavioral corporate 1488, 1540–3
 - games 721–2, 819–43, 1719–22
 - correlated equilibria 722–3, 843–51
 - quantal response equilibrium 882
 - incentives
 - extrinsic and intrinsic motivation 553, 554
 - self-signaling 528–33, 553
 - social 553
 - silver lining principle 1568–9
- similarity relation
 - cyclical choices 601
 - time preference 606, 629–31
- simple dominance, in signaling games 1722
- simple increasing elasticity (SIE) class of utility functions 605, 609–11
- simple lotteries 85, 211
- sin taxes 645–6, 679–83, 690
- small stakes
 - expected utility theory 104–6
 - loss aversion 156
- small worlds theory 298
- smooth ambiguity aversion models 82, 285, 296–7
 - empirical evidence 299
 - and maximin expected utility model, unified account of 298
- smooth fictitious play 1108
- social capital 1286
- social comparisons 360–1
- social conformity models 935
- social dilemma game 440–1
- social emotions 1223
- social evolution 1178–80
- social identity 454–5, 465–6
 - artificial 466
 - and economic incentives 473–7
 - experiments 1677–9
 - mental models 1408
 - and other-regarding preferences, evidence on 467–73
- social image, management of 1691–2
- social imitation 1040, 1085–7
- social learning 1040, 1085–91
- social monotonicity 430–1
- social networks
 - and amygdala size 1683
 - stochastic social dynamics 1182–5
- social norms
 - and correlated equilibrium 879–80, 962–6
 - dictator game 354
 - evolutionary game theory
 - cultural evolution model 1073
 - indirect reciprocity 1078
 - stochastic social dominance 1178–9
 - ultimatum game 351
- social preferences
 - emotions and cognition, interaction between 1297–8
 - leader contributions to charitable giving 836–9
 - models 398, 431–4
 - axiomatic foundations 451
 - Charness–Rabin hybrid model 444–5, 446
 - equity, reciprocity, and competition model 410–12
 - fairness and stochastic dominance 412–17
 - fairness, general equilibrium, and welfare 426–31
 - Fehr–Schmidt model 401–10
 - political economy 417–26
 - neuroeconomics of 1646–7, 1669–86
 - see also* other-regarding preferences
 - social psychology 466
 - social rewards 1646, 1669–70
 - social signaling 528–33, 553
 - social status 395–6
 - social ties 521
 - social welfare maximization 356
 - sociality, human 32–6
 - competitive market equilibrium 50
 - evidence on 344–97
 - other-regarding preferences 342
 - see also* prosocial behavior
 - societal redistribution 417
- soft paternalism 1577, 1580–1, 1586, 1590–5
 - applications 1595–9
 - behavioral welfare criteria 1600–1
 - choice-based behavioral welfare economics 1616
 - critique 1599–600
- somatic markers hypothesis 1231–2, 1234
- sophisticated experience weighted attraction
 - learning 1095, 1128–30, 1131
- sophisticated thinking 1402
- source-based uncertainty 82
- source-dependent uncertainty 286
 - see also* ambiguity
- source functions 306, 307, 308
- sources
 - ambiguity 305–6
 - definition 305
 - uniform 305–6
- sparse max operator model 1443–7, 1448
- spatial resolution, brain measurement techniques 1656
 - electroencephalography 1657
 - functional magnetic resonance imaging 1658
 - positron emission tomography 1657
 - single unit recording 1657
- special status demand 9, 10
- special transfer of attention in exchange (TAX) model 189
- speedup function, reference time theory 638, 639–40, 641
- spite 341, 448
- sports players, mixed strategy Nash equilibria 752–8
- St. Petersburg paradox
 - classical decision theory 88
 - expected utility theory vs. expected value maximization 138, 139

- Stackelberg game
 - comparison with level- k models and team reasoning 987–8
 - Cox–Friedman–Gjerstad model 449
 - psychological game theory 932
- stag-hunt game 716, 767–9, 1710
 - cognitive hierarchy model 902
 - fictitious play 1112, 1116–17
 - loss aversion 777
 - multiple signals 840–1, 842
 - reinforcement learning 1112, 1116–17
 - risk dominance and payoff dominance 877, 911–14
- state effects 14, 113
 - gift exchange game 359
 - prospect theory 175–9
 - ultimatum game 350
- standard weighting functions 119
- stationarity
 - exponential discounted utility model 590, 595, 601
 - reference time theory 641
- status quo
 - bias 1596, 1597, 1599
 - prospect theory 30, 80, 130, 138, 139, 153–4, 156
 - endowment effect 219, 235
 - original 80
- stigma, in tax evasion puzzles 244, 245, 246, 247–8
- stimulus learning *see* reinforcement learning
- stimulus–response learning 1098
- stochastic dominance
 - and fairness 400, 412–17
 - non-linear probability weighting 123–4
 - prospect theory 148, 160–1
 - see also* first order stochastic dominance; second order stochastic dominance
- stochastic dynamics, learning direction theories 1140
- stochastic potential
 - coordination game 1164
 - defined 1203
 - of a directed graph 1201
 - technology adoption model 1166
- stochastic process 1196
- stochastic reference points
 - endowment effect 233
 - prospect theory
 - rational expectations 153–72
 - third generation 152–3
- stochastic social dynamics 1036, 1158–9
 - fixing ideas in 1159–65
 - memory, conventions, and risk dominance in 3×3 games 1174–82
 - perturbed Markov dynamics in 2×2 games 1166–74
 - social networks 1182–5
- stochastic utility
 - choice under 1193–6
 - payoff shocks 882
- stochastically stable states 1046, 1158, 1160
 - coordination game 1164
 - defined 1200
 - dominated strategies as 1159, 1167–70
 - Markov chains 1200–1, 1203, 1204
 - memory, conventions, and risk dominance in 3×3 games 1175, 1176–7, 1179, 1180
 - perturbed Markov dynamics in 2×2 games 1172, 1173
- stock markets
 - limited participation 1494
 - underreaction and overreaction 1487, 1520–6
- stock variables, cue-conditioned behavior 1236–7
- strategic complements (SC) 723, 852–63
- strategic ignorance 1289, 1308–11
- strategic interaction, neuroeconomics of 1647, 1689–92
- strategic substitutes (SS) 723, 852–63
- strategic teaching models 1130, 1131
- striatum 1651
 - risky decisions 1663, 1665, 1666, 1667, 1668
 - social preferences 1670, 1672, 1673, 1676
 - strategic interaction 1690
 - time preferences 1687
- strict best response 1704
- strict dominance 1705
- strict Nash equilibrium 1047, 1705
- strong reciprocity 342–3
- Stroop task 355, 1400, 1606
- Strotz preferences 1263–4
- student subject pool 386–8
- subadditive discounting
 - common difference effect 606, 619–22
 - cyclical choices 601
 - exponential discounted utility model 37, 601
- subadditivity 584
 - defined 620
 - reference time theory 641
 - support theory 302–3
- subject pools in experiments 13
- subjective expected utility theory (SEU) 26, 79, 89–92
 - ambiguity 305
 - Ellsberg paradox 285, 286–8, 299
 - horse-roulette acts 290
 - prospect theory 308
 - support theory 301
 - limitations 1416
- subjective heterogeneous quantal response
 - equilibrium (SQRE) 880, 969–72
- subjective probability distribution 79
- subjective probability weighting function 128–9
- subjective well-being (SWB) 1271–4
- subproportionality effect, Prelec function 122
- substantive rationality 1409–10
- sufficient and necessary conditions, confusion
 - between 1347, 1397
- sunk costs 1454, 1473–6
- support theory 82
 - ambiguity 285–6, 301–4
- sure thing principle 91
- survey data 19–21
 - scope of economics 1607–8
- synergy, merger/acquisition activity 1530
- T period game 794–5
- tacit bargaining 813
- tail events, psychology of 56, 81, 1494

- take-the-best heuristic 1419–26
- tallying and take-the-best algorithm, comparison between 1422
- taxation
 - evasion puzzles 215, 242–8
 - limited attention 1438–40
 - excess burden under 1583, 1628–32
 - incidence under 1583, 1625–8
 - optimal sin taxes 645–6, 679–83, 690
- taxi drivers
 - base rate neglect 1375–6
 - prospect theory 164–7
 - rainy days 216, 256–62
- team contracts 488, 489, 490
- team reasoning
 - comparison with level- k models and Stackelberg reasoning 987–8
 - focal points 763
 - models 44
- technology adoption model 1158, 1164–5
- temporal preference reversal 37
- temporal resolution, brain measurement
 - techniques 1656
 - electroencephalography 1657
 - functional magnetic resonance imaging 1658
 - magnetoencephalography 1657
 - positron emission tomography 1657
 - single unit recording 1656, 1657
- temporoparietal junction (TPJ)
 - guilt 1686
 - punishment and trust 1678, 1679
 - strategic interaction 1690–1
- temptation
 - and conflicts between commitment and flexibility 1228, 1266–70
 - emotions and cognition, interaction between 1292
 - preferences 1228, 1258–66
- theoretically progressive research programs 6
- theory of mind *see* mind, theory of
- theory of relativity 7–8
- third generation prospect theory (PT³) 80, 112, 152–3, 199
 - and Köszegi–Rabin framework, relation between 161
 - preference reversals 108
- third party punishment (TPG) game
 - equity, reciprocity, and competition model 411
 - incentives 550
 - indirect reciprocity 396, 397
- threats and promises, model of 454, 461–5
- time consistent preferences 652
- time discounting 583–5
 - evidence 586–604
 - models 605–43
 - present-biased preferences, applications of 644–91
 - and risk 207
- time preferences
 - neuroeconomics of 1647, 1686–9
 - reference time theory 641–2
 - and risk preferences, relation between 115, 199–212, 619, 626
- time sensitivity, reference time theory 642
- tipping point, stochastic social dynamics 1163, 1168–70
- tit-for-tat strategy
 - evolutionary game theory 1055, 1057
 - reciprocal altruism 1076, 1077
 - rule-based learning 1148
- trade-ins, exchange asymmetries 232
- trade-off attribute model of intertemporal choice 606, 631–2, 633
- trade-off consistency method 148–50
- trade-off method, eliciting the utility function under EU 94
- trading volume 1519–20
- transaction costs
 - arbitrage 1495
 - contracts 264
 - efficient markets hypothesis 1495, 1501
- transaction utility 1454–5, 1474, 1476–7
- transcranial direct current stimulation (tDCS) 1659, 1660, 1663
- transcranial magnetic stimulation (TMS) 1659, 1679–83, 1689
- transfer of attention in exchange (TAX) models 189, 190, 191
- transference of learning 1146
- traveller's dilemma
 - functional experience weighted attraction learning 1127, 1128
 - quantal response equilibrium 892–4
- travelling salesman problem 47
- trial and error learning 1114, 1115
- trial offers, exchange asymmetries 232
- truncated quantal response equilibrium model (TQRE) 971
- trust 366–70
 - incentives
 - extrinsic and intrinsic motivation 528, 539–41, 546
 - incomplete contracts 510
 - legal environment for contracts 556
 - neural basis 1646–7, 1673–84
 - trust contracts 479–80, 495, 497
 - evidence on contract choice under contractual incompleteness 506, 509
 - extrinsic and intrinsic motivation 542–3
 - other-regarding preferences 498, 500–3
 - self-regarding preferences 498
- trust game 346, 366–70
 - altruism 399
 - Charness–Rabin hybrid model 445
 - competitive market equilibrium 50
 - envy 399
 - external validity of lab evidence 393, 394
 - guilt 954, 955, 1686
 - human sociality 33, 35
 - incentives
 - crowding-in effect 559
 - extrinsic and intrinsic motivation 535, 536, 541–2
 - indirect reciprocity 396–7
 - intention-based reciprocity 438
 - neuroeconomics

- human virtues 1684–5
- oxytocin 1692–3, 1694
- social preferences 1686
- social punishment and trust 1673–4, 1676, 1681, 1683
- psychological game theory 928
 - guilt 954, 955
 - plans, dependence of utility on 960
- social identity 471, 473
- student subject pool 386, 387–8
- trustworthiness 366–70
 - legal environment for contracts 556
- truthfulness, evidence on 455, 457–8, 459–60
- Tversky–Kahneman probability weighting function 122, 146
- two-modules view of the brain 1288, 1290–8
- judgment heuristics 1347, 1398–401, 1430–1
- two-stage recursive models 82
 - ambiguity 285, 294–6
- ultimatum game 344–5, 349–51
 - Charness–Rabin hybrid model 444, 447
 - competitive market equilibrium 50
 - Cox–Friedman–Gjerstad model 449
 - cultural factors 1072
 - deliberative vs. automatic choices 354
 - and dictator game, comparison between 352
 - envy 399
 - equity, reciprocity, and competition model 410, 411
 - Fehr–Schmidt model 403, 405, 406
 - human sociality 33–4, 35, 36
 - incentives
 - extrinsic and intrinsic motivation 549, 550
 - moral hazard 486
 - intention-based reciprocity 434–5, 439, 440, 441–4
 - interdependent preferences models 448, 449
 - internal validity 12
 - positive bargaining solutions 794, 799, 800, 802
 - promises and threats model 463
 - psychological game theory 948
 - punishment and trust 1676–7, 1680, 1683, 1684
 - reinforcement learning 1101
 - search and lookups, inferring cognition through 818
 - student subject pool 387, 388
 - willpower, depletion of 355
- unacclimating personal equilibrium (UPE)
 - stochastic reference points under rational expectations 158–9, 160, 161, 164–7
 - taxi drivers' problem 164–7
- unbiased Bayesian statistician 1393
- uncertainty 79, 83
 - aversion, multiple priors models 292, 293
 - Bayesian approach to 91
 - choice 235
 - choice data 1608
 - classical decision theory 85
 - common difference effect 606, 623–7
 - exchange asymmetries 235
 - hyperbolic discounting 619
 - limited attention 1448
 - prospect theory 80, 134
 - stake size effect 178
 - trade-off consistency method 150
 - rank dependent utility theory 128–9
 - resolution 207
 - source-dependent 286
 - see also* ambiguity
 - subjective expected utility theory 89, 91
 - trade 235
 - weighting function for 305
- uncoupled learning rules 1156
- underconfidence
 - confirmation bias 1394, 1395
 - conservatism 1379, 1383
- unemployment and happiness 1283–4
- uniform sources 305–6
- unsophisticated investors *see* noise traders
- vacations, and mental accounting 1464, 1467–8
- vague time preferences model 601, 606, 628–9
- value
 - function 131, 133, 134, 245
 - neural basis of 1645–6, 1661–3
- variational representation of preferences 298, 299
- ventromedial prefrontal cortex (vmPFC) 1653
 - risky decisions 1662–3, 1665
 - social preferences 1669, 1670, 1672, 1673
 - guilt 1686
 - punishment and trust 1676, 1681, 1682–3
- Vickrey auctions 18–19
- virtues, human 453–4
 - evidence on 455–61
 - neural basis 1647, 1684–5
 - promises and threats model 461–5
- visceral influences 1226, 1229–30
 - and cues 1235–6
- Wald criterion, multiple priors models 293
- Walrasian auctioneer 863, 868, 873
- Walrasian competitive equilibrium 870
 - fairness and welfare 400–1, 429, 430–1
- weak axiom of revealed preference with limited attention (WARP(LA)) 1622, 1623
- weak certainty independence axiom, variational representation of preferences 298
- weak law of large numbers 1350
- weak link games 716, 765–7
 - reinforcement learning 1103
 - see also* weak link minimum effort games
 - weak link minimum effort games
 - communication and experience 783–5
 - financial incentives 774–5
 - group size, gradual growth in 778–9
- weak WARP axiom 1623
- wealth factors, and patience 664
- weighted fictitious play 1111–13, 1127–8
- weighting function for uncertainty 305
- welfare
 - behavioral welfare economics 1577–83
 - comparison 1298
 - evolutionary stable strategies 1046
 - fairness and general equilibrium 400–1, 426–31
 - improvement 1615

- welfare (*cont.*)
 - payments, long-term dependence on 689
 - present-biased preferences 675
 - social identity 473
- well-being
 - externalities 429–30
 - see also* happiness economics
- Westermarck effect 1067
- white lies
 - altruistic 455, 461
 - evidence on 455
 - Pareto 455, 461
- white matter 1651
- willingness to accept (WTA) 81, 217
 - defined 220
 - mental accounting 1472–3
 - neuroeconomics 1665
 - risky (RWTA) 222, 223, 224
 - third generation prospect theory 153
 - and willingness to pay, discrepancy between 213
 - see also* endowment effect; exchange asymmetries
- willingness to pay (WTP) 81, 217
 - defined 220
 - neuroeconomics 1665, 1666
 - risky (RWTP) 223, 224
 - third generation prospect theory 153
 - and willingness to accept, discrepancy between 213
 - see also* endowment effect; exchange asymmetries
- willpower
 - asymmetric information about 1308
 - confidence, motivation, and self-esteem 1311
 - dual selves model with costly commitment 1299, 1300, 1302, 1303
 - limited attention 1437
 - and personal rules 1321
 - two-modules view of the brain 1291, 1293–4, 1295, 1296, 1297, 1298
 - ultimatum game 355
- win–stay, lose–shift strategy, evolutionary game theory 1077
- windfalls, and mental accounting 1460, 1461, 1463
- winner's curse 881, 991–2
 - cursed equilibrium 995–7, 1000–1
 - empirical evidence 993–5
 - level-*k* explanation 998–1002
 - managers 1528
- wishful thinking 988
- Yerkes–Dodson law 543