FORGETTING WE FORGET: OVERCONFIDENCE AND MEMORY

Keith M. Marzilli Ericson
Harvard University

Abstract
Do individuals have unbiased beliefs, or are they over- or underconfident? Overconfident individuals may fail to prepare optimally for the future, and economists who infer preferences from behavior under the assumption of unbiased beliefs will make mistaken inferences. This paper documents overconfidence in a new domain, prospective memory, using an experimental design that is more robust to potential confounds than previous research. Subjects chose between smaller automatic payments and larger payments they had to remember to claim at a six-month delay. In a large sample of college and MBA students at two different universities, subjects make choices that imply a forecast of a 76% claim rate, but only 53% of subjects actually claimed the payment. (JEL: C91, D81, D83, D84)

1. Introduction

Are the beliefs of individuals well-calibrated, or are individuals under- or overconfident? Accurate beliefs matter for the welfare of individuals. Expectations and beliefs are key determinants of behavior, and behavior will be suboptimal if beliefs are biased. Individuals who underestimate the probability that they will get sick or die will hold too little health insurance and life insurance. Individuals who are overconfident about their ability to hold a job will save too little for spells of unemployment. Moreover, economists may make mistaken inferences if we incorrectly assume that beliefs are unbiased and then attempt to infer preferences such as risk aversion or time discounting from observed behavior.

Influenced by evidence from the psychological literature, economists have begun to model the consequences of overconfidence.[1] However, existing evidence on

This paper was edited by Orazio Attanasio

Acknowledgments: I thank John Beshears, James Choi, Michael Grubb, David Laibson, Brigitte Madrian, and Al Roth for their thoughtful comments, along with participants at the Yale Whitebox Behavioral Science Conference, the Harvard seminar on Decision Making in Economics and Psychology, and this paper's thoughtful referees. This research was supported by the National Institute on Aging (P01-AG05842 to the NBER) and by a Williams College Tyng Grant. I thank David Borden, Ananya Chakravarti, Jonathan Hall, Shih En Lu, Dina Mishra, and Kelly Shue for assistance with data collection and for helpful discussion.

E-mail: kericson@fas.harvard.edu

1 For instance, Sandroni and Squintani (2007) analyze insurance markets in the presence of overconfidence, and Scheinkman and Xiong (2003) model speculative bubbles as resulting from trader overconfidence.

Journal of the European Economic Association February 2011 9(1):43–60
 DOI: 10.1111/j.1365-2966.2010.01005.x

overconfidence is surprisingly weak and subject to a number of critiques. There has been limited work in the economics literature measuring overconfidence via choices. Psychology studies typically rely on statements of confidence, and classic studies use ambiguous questions. Further, re-examination of the finding of overconfidence in the often-studied domain of trivia questions has shown patterns consistent with Bayesian reasoning (Gigerenzer, Hoffage, and Kleinbölting 1991).

This paper uses data on choices to document overconfidence in the domain of prospective memory—the ability to remember to undertake actions. Prospective memory is a domain that is economically important, and beliefs in this domain have not been studied. Prospective memory is necessary to translate present intentions into future actions, and is implicated in nearly all situations in which we make plans for future behavior. In particular, prospective memory affects individuals' ability to both manage their finances and undertake health behaviors. Individuals must remember to enroll in retirement savings plans, change their savings rates in response to changes in their financial situations, rebalance their investment portfolios, and simply pay bills on time. Individuals must remember to purchase or cancel insurance appropriately: when providing for dependents, they must remember to purchase life insurance, and if they wish to reduce their insurance when their children become financially independent, they must remember to do so. Similarly, in the domain of health, individuals must remember to schedule preventative tests and checkups with their doctors. Overconfidence on prospective memory tasks can partially explain why there is inertia in retirement savings plan participation (Madrian and Shea 2001) and why life insurance is not attuned to financial vulnerabilities (Bernheim et al. 2003).

While individuals clearly have access to memory aids, they may fail to invest enough in such aids if they overestimate their ability to remember. Moreover, while the cost of limited prospective memory is bounded by the cost of memory aids for individuals with accurate beliefs, the cost to overconfident individuals is unbounded (Holman and Zaidi 2004).

When consumers are biased, profit-maximizing firms will design contracts to exploit such biases (DellaVigna and Malmendier 2004). Firms should then respond to overconfidence. Overconfidence in prospective memory can explain the existence of free-trial/automatic renewal offers and rebate markets (Holman and Zaidi 2004). Overconfident consumers do not claim rebates or cancel free trials as often as they expect. Thus, even apart from any price discrimination rationale, firms may offer rebates since the cost to the firms is less than the value that consumers incorrectly expect to receive. It may even be the case that no firm has an incentive to debias consumers, since once consumers are educated they may invest in memory aids and claim the rebate more often than overconfident consumers.

Differences in sophistication about prospective memory also have economic implications. For instance, overconfident consumers think they will remember to pay credit card bills on time more often than they actually do; when they fail to do so, they face costly penalty fees. Sophisticated consumers with accurate beliefs can avoid such fees by taking costly steps such as setting more reminders or enrolling in automatic payment programs. Sophisticated consumers would then receive a subsidy from overconfident consumers in the form of credit cards with lower annual fees and other benefits. This result is akin to the "shrouded equilibrium" of Gabaix and Laibson (2006).[2] The current paper provides evidence that overconfident consumers exist in the population. The design does not allow us to identify whether a subset of subjects have accurate beliefs, but at least some such agents are likely to exist.

This paper provides experimental evidence that individuals are overconfident in the domain of prospective memory using an experimental design that is more robust to potential confounds than previous research. In large samples at two different universities, subjects make choices that imply they believe they will claim a delayed payment much more often than they actually do: their choices imply a forecast of a 76% claim rate, but only 53% of subjects actually claimed the payment. Overconfidence is measured at the population level, not at the level of individual subjects. Overconfidence on any prospective memory task may result from overconfidence in memory per se, or overconfidence regarding future effort. I examine the potential sources of overconfidence, and conclude that due to the low cost of claiming the payment, overconfidence in memory must play a significant role in the explanation of the results.

2. Related Literature

Economists have tested for overconfidence using inference from incentivized choices. Incentives may motivate subjects to be more accurate, as they will bear the consequences of their decisions.[3] Camerer and Lovallo (1999) experimentally test for overconfidence in a strategic entry setting for a trivia contest and find that subjects enter the competition too often. Hoelzl and Rustichini (2005) examine votes between pay-for-performance and pay-by-lottery compensation for experimental test performance and find mixed evidence indicating both underconfidence and overconfidence. However, subjects may seek status in such competitions (e.g. a trivia contest) as well as monetary payoffs. Moreover, when voting on a compensation scheme for a group, subjects may be guided by norms of fairness as well as profit-maximization. The present study demonstrates overconfidence in an environment where status seeking, game playing, and norms of fairness are unlikely to be involved.

Overconfidence with regard to relative ability can result from biased assessments of either a subject's own ability, or others' abilities, or both. Although many economic models assume only relative overconfidence, others may rely on the more specific case of absolute overconfidence. Previous tests for overconfidence in the economics

2 The shrouded equilibrium of Gabaix and Laibson (2006) can occur not only with add-ons that entail hidden costs, but also with costly add-ons that consumers underestimate their probability of remembering to avoid or benefits that consumers overestimate their probability of remembering to claim (e.g. rebates). For example, car renters face added fees if they forget to return the car with a full tank of gasoline. If consumers are overconfident about their probability of remembering, a shrouded equilibrium can exist.

3 However, see Camerer and Hogarth (1999) for a review of the effect of incentives in laboratory experiments. They find that hypothetical choices are often very similar to incentivized choices. However, because self-image may be a consumption good, incentives may be relatively more important for self-evaluation than for other types of preference elicitation.

literature (Camerer and Lovallo 1999, Hoelzl and Rustichini 2005) have focused on subjects' assessment of their abilities relative to other subjects. In contrast, this study provides evidence that subjects have absolute overconfidence.

The existing literature on overconfidence in psychology has largely focused on statements of confidence: subjects are asked to state a confidence interval (Soll and Klayman 2004), to judge the probability they will answer a trivia question correctly (Gigerenzer, Hoffage, and Kleinbölting 1991), or indicate the probability they will face a life event (Weinstein 1980). Although overconfidence is often shown in these statements, it is less clear whether overconfidence translates into choice behavior. The present study shows that overconfidence is demonstrated in actual choices.

Classic psychology studies also relied on ambiguous questions: the well known fact that 70% to 80% of subjects think they are "safer" than the median driver (Svenson 1981) can be explained by disagreement about what constitutes a safe driver: Should one keep up with traffic or follow the posted speed limit? The present study uses choice between clearly defined options to minimize potential disagreements in interpretation.

Observed overconfidence in the frequently studied domain of trivia questions may also be a result of choosing unfamiliar or harder-than-average questions (Gigerenzer, Hoffage, and Kleinbölting 1991). The so-called hard–easy effect—overconfidence on hard questions and underconfidence on easy questions—is consistent with subjects making unbiased Bayesian forecasts of their ability. Consider subjects with an accurate prior distribution over the difficulty level of questions an experimenter would present to them in the context of an experiment. Subjects observe the questions and produce a Bayesian posterior distribution $F(d)$ of the questions' difficulty. Subjects then use this distribution to produce an unbiased forecast of the probability they will answer correctly. The experimenter chooses the difficulty d^*. If $d^* > E(F(d))$ overconfidence on hard items is observed, and if $d^* < E(F(d))$ underconfidence on easy items is observed. In the context of relative overconfidence, Moore and Cain (2007) present evidence that subjects use similar Bayesian reasoning when faced with tasks about which they have imperfect information. The present study addresses this critique, as subjects know the full structure of the task and have experience with similar tasks.

Failures of prospective memory have also been studied in the psychology literature (Guynn, McDaniel, and Einstein 1998), but this work has generally focused on laboratory-based tasks (e.g. remembering to press a button when presented with a cue) and has not included tests of overconfidence.

3. Experiment Design

This paper provides experimental evidence that individuals are overconfident in the domain of prospective memory. In the experiment presented here, subjects made choices between larger payments that were contingent on subjects remembering to claim them at a six month delay and smaller payments that were automatically sent to them at the same delay. A lower bound on each subject's subjective probability of remembering to claim the payment can be derived from his or her incentivized choices.

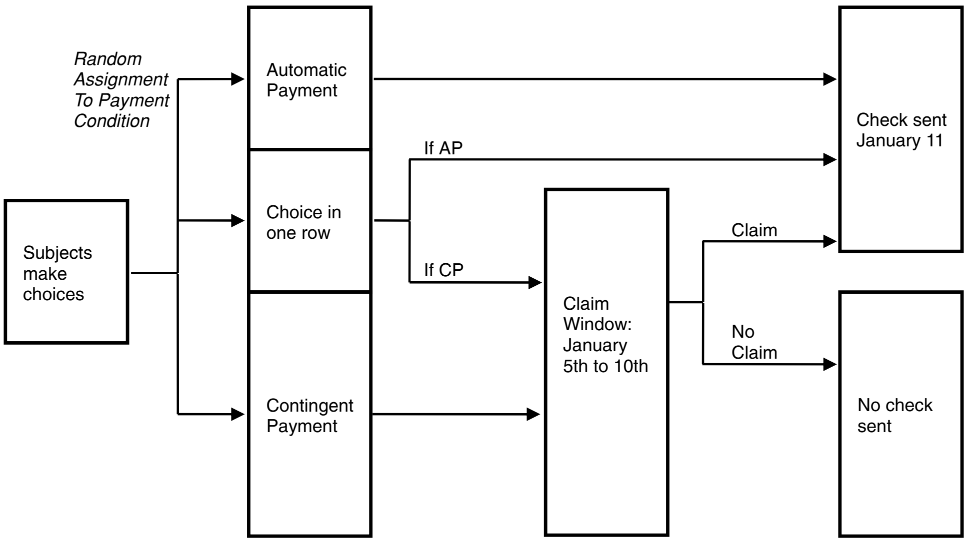

FIGURE 1. Experiment design.

This average lower bound on the subjective probability of remembering can then be compared to the actual claim rate.

3.1 Procedure

The design of the experiment is summarized in Figure 1. Subjects first read instructions describing the structure of the experiment.[4] Then, on each of 21 rows, subjects indicated their choice between the following two types of payments:

(1) *contingent payment: $20 contingent on remembering to claim the payment;*
(2) *automatic payment: $$x$ automatically sent to the subject;*

where $$x$ varied monotonically from $5 to $20 in $0.75 increments. See Table 1 for an excerpt of the instructions subjects saw.

After making their choices, subjects were randomly assigned to one of three payment conditions:

(1) *the choice they made in a randomly selected row;*
(2) *contingent payment regardless of choice;*
(3) *automatic payment regardless of choice.*

Random assignment was conducted by having subjects roll a ten-sided die twice, generating a two-digit number from 0 to 99 with uniform probability. For dice rolls from

4 See the online version of this article for the complete experimental materials.

TABLE 1. Excerpt from experiment instructions.

Please indicate your choices between the two rewards listed on each line by checking your preferred box. Rewards in both columns will be mailed at the same time, on 11 January 2006.

Row	Rewards in this column will be automatically mailed to you. No further action is required on your part.	Rewards in this column will be mailed to you if you email us with your contact information between Jan. 5 and Jan. 10, 2006. You must email us within this time period to receive your reward.
#0	____ $20	____ $20
#1	____ $19.25	____ $20
#2	____ $18.50	____ $20
...	...	...
#19	____ $5.75	____ $20
#20	____ $5	____ $20

0 to 20, subjects received the election they made in the row that was thus numbered. They therefore had a 21% chance of receiving the election they had made in one of the 21 rows. With 60% probability (dice rolls from 21 to 80), subjects were assigned to the $20 contingent payment condition, while with 19% probability (dice rolls from 81 to 99) they were assigned to receive an automatic payment of $20. These probabilities were chosen to facilitate a simple description of the tasks to subjects. All subjects then filled out a sheet giving their contact information. Subjects receiving the automatic payment also filled out a short questionnaire, discussed in the next section.

This random assignment regime balanced a number of objectives. Random assignment to the contingent payment condition was necessary to measure the average claim rate in the population and also increased statistical power. Random assignment to the automatic payment condition was included to reduce the potential that experimenter demand effects would result from very uneven treatment of the two payment types. Thus, subjects could be randomly assigned to either condition, reducing the likelihood that subjects inferred that one payment type was endorsed by the experimenter. Finally, randomly assigning some subjects their election in one row gave their choices real stakes.

Claiming a contingent payment constituted a memory task, as subjects were told that they could only claim the contingent payment by emailing the experimenter between January 5 and January 10, 2006. It was clearly indicated that claims outside this period would not be honored, and that the contingent payment would only be mailed if the subject emailed the experimenter. Subjects were not restricted from using memory aids (e.g. calendars, PDAs) and were given a reminder sheet with instructions on how and when to email to claim the payment. Both types of payments would be sent on the same date, so time preference is irrelevant to the choice between the two types of payments. In the instructions, subjects were reminded multiple times that both types of payments would be mailed simultaneously on January 11 and that they must contact the experimenters to receive the contingent payment.

Subjects were told the full structure of the experiment at the beginning. Subjects knew that they would make their choices and that a dice roll would then determine how they would be paid. They also knew the probabilities of each outcome. At the start, they were told the precise dates of the window for claiming the payment and the date on which the payments would be mailed.

3.2 Participants

Subjects were recruited from two different universities during summer 2005. The College Sample is comprised of 80 subjects who were summer session students at a large elite university, though some of these students were high-school students over age 18 taking summer classes. The MBA Sample is comprised of 266 subjects recruited at another large elite university in a different city. Although subjects were predominately first-year MBA students, nine college students and two economics Ph.D. students were also part of this sample.

The experiment was run on a drop-in basis: subjects were permitted to participate in the experiment by showing up at the room during specified times when the experiment was taking place. Slots were available for five days in July 2005 for the College Sample, and three days in August 2005 for the MBA Sample. In order to participate, subjects were required to produce identification showing affiliation with the university. Before participating in the prospective memory experiment, subjects first completed an unrelated experiment on investment portfolio allocation, as described in Choi, Laibson, and Madrian (2006). For participating in the experiment, subjects in the MBA Sample were paid \$20, and subjects in the College Sample were paid \$5. Subjects were also entered in a lottery to win a payment based on their allocations in the investment experiment.

3.3 A Subject's Decision Problem

A lower bound on the estimate of a subject's subjective probability of completing the memory task can be derived from that subject's choices. The solution for a risk-neutral agent's decision problem is a limiting case, as a lower bound for a risk-neutral agent is also a lower bound for a risk-averse agent. To choose a given uncertain contingent payment over a certain automatic payment, a risk-averse agent would require a higher subjective probability of remembering than would a risk-neutral agent. A risk-neutral agent chooses the contingent payment of \$20 over an automatic payment of \$$x$ if and only if its expected value is greater than that of the automatic payment. This implies that

$$\hat{p}\,(20 - costClaim) - costTry \geq x, \tag{1}$$

where $\hat{p}$ is the subject's subjective probability of completing the memory task, and *costTry* and *costClaim* are the (unobserved, non-negative) costs associated with attempting to remember to claim the payment, respectively, and actually claiming the payment. These costs may be stochastic and can therefore capture subjects' forecasts of the probability that email access is inconvenient or that time costs are higher than average. Unfortunately, these costs cannot be observed, but because the expected value of the contingent payment is decreasing in the unobserved costs, setting them to zero allows the derivation of a lower bound on $\hat{p}$. Rearranging then gives

$$\hat{p} \geq x/20$$

when the contingent payment is chosen. Since x varies across the 21 different choices subjects make, I define for each subject the lower bound on subjective probability of remembering as $\underline{\hat{p}} = x_{\max}/20$, where $x_{\max}$ is the highest value of the automatic payment for which the subject chooses the contingent payment. For example, a subject who chose the contingent payment only when the automatic payment was less than or equal to \$12 would have $x_{\max} = 12$ and $\underline{\hat{p}} = 0.60$. If a subject never chooses the contingent payment, $\underline{\hat{p}} = 0$.

Subjects are overconfident if their average estimate of the probability of recall is positively biased. I test for overconfidence by comparing the actual fraction of subjects who claim the payment to the expected fraction implied by their choices. If the mean $\underline{\hat{p}}$ of the sample is significantly greater than the fraction of the sample that claims the payment, then the sample is overconfident on average regarding its prospective memory ability.

4. Results

The experiment gives evidence of overconfidence on the memory task at the population level. Section 4.1 describes how the relevant samples were constructed. Section 4.2 examines population-wide behavior and shows that subjects were, on average, overconfident: subjects' choices implied that they would claim the payment about three-quarters of the time, but in fact, only half claimed the payment. Section 4.3 provides evidence that the derived $\underline{\hat{p}}$ reflects subjects' subjective probabilities: it is correlated with their unincentivized statements of confidence, and the average $\underline{\hat{p}}$ is near the average stated probability of claim. Section 4.4 examines how behavior varies with beliefs. It shows that overconfidence exists through most of the distribution of $\underline{\hat{p}}$ and shows that $\underline{\hat{p}}$ predicts probability of claim.

4.1 Sample Construction

I construct the Analysis Sample from all 193 subjects who were randomly assigned to the contingent payment, combining the College Sample and the MBA Sample. I remove the seven subjects in this sample whose answers are inconsistent, leaving 186 subjects in the Analysis Sample. Subjects' choices must be consistent in order to derive a subject's subjective probability of completing the memory task. Consistency requires that if a subject chooses the contingent payment over a given value of the automatic payment, then the subject should also choose the contingent payment over all smaller values of the automatic payment (and vice versa). Most subjects made consistent choices.[5] Results are unchanged even if they are included and assigned a $\underline{\hat{p}}$ of 0.

The Analysis Sample allows us to examine the behavior of a representative sample of the subject population. The Analysis Sample includes only those *randomly* assigned to the contingent payment condition, not those who were placed in contingent payment condition as a result of receiving their election in a randomly chosen row. This allows for the appropriate hypothesis test for population overconfidence: comparing the claim frequency of a random sample of the population (the group randomly assigned to the contingent payment condition) with the average estimate of the subjective probability of claim derived from those subjects' choices. Subjects assigned to receive the contingent payment as a result of their choices are excluded because any overconfidence among those who elected to receive the contingent payment could be matched by underconfidence in those who elected to receive the automatic payment, whose claim frequency cannot be observed.

I also construct the Automatic Payment Sample from the 61 subjects who were randomly assigned to the automatic payment condition. Again, to make the sample representative of the subject population, only subjects who were randomly assigned to automatic payment are included in this analysis; those who received an automatic payment as a result of getting one of their choices are excluded. These subjects were asked (hypothetically) how likely they would have been to remember to claim the contingent payment.

4.2 Population Overconfidence

The Analysis Sample had an average implied subjective probability of claim $\underline{\hat{p}}$ of 0.76, with a distribution as shown in Figure 2. The fraction of these subjects who claimed the payment was only 0.53, significantly and substantially lower. I reject the null hypothesis that the fraction claimed is equal to that implied by subjects' choices (unpaired t-test, $t = 5.61, p < 0.0001$).[6] Results are summarized in Table 2. Overconfidence is similarly found in both the MBA and college student subsamples, as well as for both men and women.

Because overconfidence is displayed by both the College Sample and the MBA Sample, this finding is not likely to be driven by a group-specific shock. Nor is this

5 In addition to the seven subjects dropped from the Analysis Sample, one other subject also made inconsistent choices. The eighth subject to make inconsistent choices was assigned to receive the automatic payment as a result of his/her choices and would not have been included in any analyses anyway. Of the eight inconsistent subjects, three only had one choice that was inconsistent. One subject simply alternated back and forth between the two types of payments, indicating disengagement from the task. The other four subjects followed no discernable pattern. Of the seven inconsistent subjects dropped from the Analysis Sample, three claimed the payment (43%).

6 The precise question, what is the probability of observing the claim fraction given subjects' subjective probabilities were correct, requires evaluating the probability of more than $180!/(90!\,90!) \approx 10^{52}$ possible subsamples, which is computationally infeasible. Bootstrapping indicates that given subjects' implied probability estimates, the probability of seeing a claim fraction less than or equal to 0.53 is below 1/100,000.

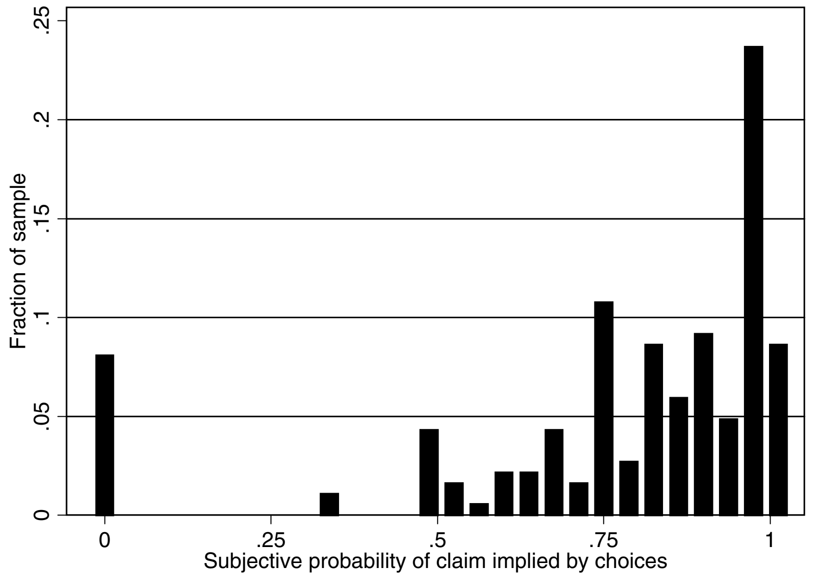

FIGURE 2. Distribution of subjective probability of claim ($\underline{\hat{p}}$). (Analysis Sample.)

TABLE 2. Test for population overconfidence.

	N	**Fraction Claimed**	**Average Implied Claim Probability**	***t*-statistic**
Analysis Sample	186	0.53	0.76 (0.04)	5.61***
. . .College Only	36	0.39	0.82 (0.04)	4.76***
. . .MBA Only	150	0.56	0.75 (0.02)	4.01***
. . .Men Only	114	0.54	0.72 (0.03)	3.29**
. . .Women Only	72	0.50	0.82 (0.02)	4.99***

** $p < 0.01$, *** $p < 0.001$.
Standard errors in parentheses.

result being driven by a misunderstanding of the instructions. Choosing a contingent payment over an automatic payment of the same amount is a weakly dominated action that should only be chosen if $\underline{\hat{p}} = 1$ and costs of claiming and trying to claim are zero. If the 16 subjects in the Analysis Sample (8.6%) who make this weakly dominated choice are excluded, the results are virtually unchanged: average $\underline{\hat{p}}$ is 0.74, the fraction claimed is 0.52, $t = 4.91$. Finally, this result is not driven by subjects attempting to claim the payment outside the specified claim window. Only five subjects attempted to claim payment after the claim window had passed. The analysis thus far counts them as having failed to complete the memory task. However, they may not have found credible the experimenters' claim that payment would be denied if they did not fulfill the requirement. Even if they are counted as having claimed the payment, the results are unchanged.

4.3 Relationship Between Choices and Statements

The $\underline{\hat{p}}$ derived from subjects' choices can be compared to an unincentivized elicitation of subjects' subjective probability of recall. I find that these measures are similar, indicating that the estimated $\underline{\hat{p}}$ is measuring subjects' subjective probability of claiming the payment. This similarity also indicates that the subjects' estimates of the costs of claiming and trying to claim the payment are low, since the estimated $\underline{\hat{p}}$ will diverge more from the true subjective probability of claim the larger such costs are.

Of the 61 subjects in the Automatic Payment Sample, 59 provided answers to how likely they thought they would have been to claim the \$20 payment. They answered the question: "On a scale of 0% to a 100%, how likely would it be that you would remember to contact us via email between Jan. 5 and Jan. 10, 2006 to claim a \$20 reward? (0% means you certainly would not remember and 100% means you certainly would remember)."

The average stated hypothetical probability of claim was 0.66, compared to an average $\underline{\hat{p}}$ for this sample of 0.69. This difference is not significant (paired t-test, $t = 0.56$). However, the correlation of these two measures is only 0.36. Since the average values of the two measures are similar, the low correlation is unlikely to be an artifact of one measure being a lower bound and the other being an estimate. Rather, it suggests that unincentivized statements of belief may be only distantly related to actual choices.

4.4 Relationship Between Beliefs and Behavior

While it is clear from the previous section that the population of subjects is on average overconfident, we can also examine how beliefs are related to behavior. Does variation in the subjective probability of claim predict claim behavior? Those who claimed the payment had a higher $\underline{\hat{p}}$ than those who did not: a 0.81 claim fraction versus a 0.71 claim fraction, a difference that is significant at the 5% level. Table 3 shows the results of a probit model for claim behavior. It demonstrates that $\underline{\hat{p}}$ statistically significantly predicts successful claim, indicating that subjects were able to predict whether they were more or less likely to claim the payment. The marginal effect at the mean of 0.348 indicates that that a 1 percentage point increase in $\underline{\hat{p}}$ from its mean of 0.76 to 0.77 raises the predicted probability of claim by (only) 0.3 percentage points, from a predicted 0.527 to 0.530 probability of claim.

This relationship is robust to excluding people with $\underline{\hat{p}} = 0$ or 1, so it is not being driven by extreme beliefs. As shown in Table 3, this relationship between $\underline{\hat{p}}$ and claim behavior holds with and without demographic controls, but none of the demographic controls predict claim when $\underline{\hat{p}}$ is controlled for.

Figure 3 plots average claim behavior by subjective probability of claim. Subjects are grouped by their subjective probability of claim into 17 groups, one for each value

TABLE 3. Relationship between beliefs and behavior: probit model and means.

	Dependent Variable: Remembered to Claim Payment			
	Coefficient	**Marginal Effect**	**Coefficient**	**Marginal Effect**
Subjective Probability of Claim $\underline{\hat{p}}$	0.873*	0.348*	1.009**	0.402**
	(0.353)	(0.140)	0.364	0.145)
In MBA Sample	–	–	0.584	0.228 ˆ
			(0.303)	(0.113)
Male	–	–	0.137	0.055 ˆ
			(0.202)	(0.080)
Age in Years			−0.0155	−0.006
			(0.027)	(0.011)
N	186		186	
Pseudo R^2	0.025		0.045	

* $p < 0.05$, ** $p < 0.01$
Analysis Sample. Standard errors in parentheses. Marginal effects are reported at the mean of the independent variables. The symbol ˆ indicates a dichotomous variable, for which the effect of a discrete change from 0 to 1 is reported.

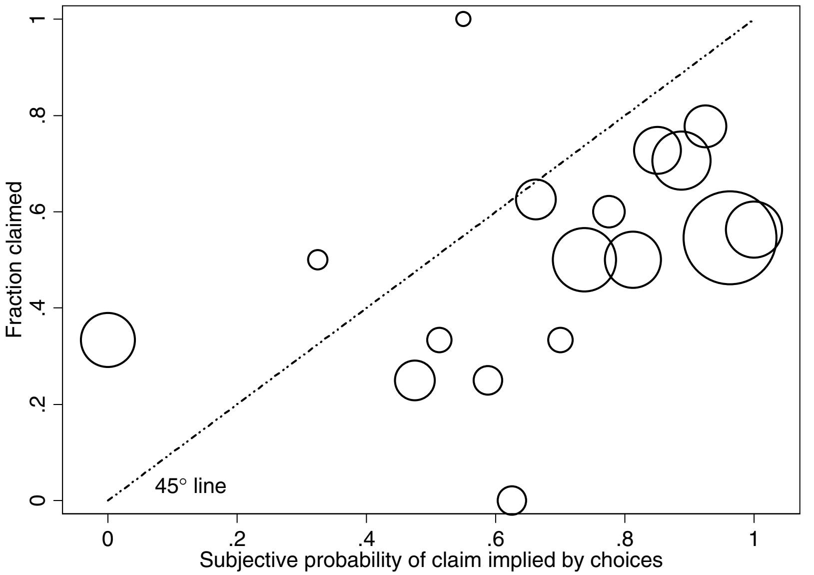

FIGURE 3. Average claim behavior by subjective probability of claim ($\underline{\hat{p}}$). (Analysis Sample. Marker size indicates the number of subjects in each $\underline{\hat{p}}$ value group.)

of $\underline{\hat{p}}$ (recall that $\underline{\hat{p}}$ is a discrete variable derived from a series of choices). This $\underline{\hat{p}}$ is plotted against the fraction of each group that claimed the payment. The size of the marker represents the size of the group, which can also be read from Figure 2. Note that all but three groups lie below the 45 degree line. This indicates overconfidence at each of those belief levels, as the fraction of the group that claimed is below the group's average subjective probability of claim. Two of the groups that lie above the

45 degree line contain small numbers of subjects: only one subject had a $\underline{\hat{p}}$ of 0.55, and (s)he claimed the payment, and only two subjects had a $\underline{\hat{p}}$ of 0.325, one of whom claimed the payment. Because these are small samples, the apparent underconfidence in these groups is likely attributable to noise, and we cannot reject the hypothesis that the claim rate is equal to the $\underline{\hat{p}}$ in each of these groups: the probability of the realizing the observed behavior given beliefs were accurate is 0.45 in the first case, and 0.44 in the second case.

A larger group of subjects (15, or 8% of the Analysis Sample) never chose the contingent payment and were therefore assigned a lower bound $\underline{\hat{p}} = 0$. This group of subjects claimed the payment 33% of the time, and a t-test indicates that this is significantly different from zero at the 5% level. While this may be taken to suggest underconfidence in this group, recall that $\underline{\hat{p}}$ is a lower bound on the subjective probability of claiming the payment: the true value is likely to be higher due to unobserved costs of effort in claiming the payment. This group may therefore simply have higher costs of claiming the payment, and we cannot rule out overconfidence, underconfidence, or accurate beliefs among this group.

5. Discussion

5.1 Sources of Overconfidence

The results show that subjects were overconfident about their performance on a prospective memory task. Prospective memory tasks in this experiment and in the marketplace involve both *remembering* and *acting*. Biased beliefs in either of these areas could bias estimates of prospective memory task performance. The following biased beliefs could explain subjects' overconfidence:

(1) Overestimating the probability of remembering they planned to make a claim;
(2) Overestimating the amount of effort they would be willing to provide to make a claim (but correctly estimating the amount of effort required);
(3) Underestimating the effort required to make a claim.

Simple calibrations indicate that subjects were likely to have overestimated the probability they would remember they planned to make a claim, rather than to have overestimated the probability they would claim conditional on remembering.

Note that claiming the payment requires experiencing a cost today (the effort to send the email) to receive a benefit in the future. Thus, while the choice between an automatic payment and a contingent payment is not an intertemporal choice, the decision to claim the contingent payment is. While procrastination can explain delay in sending the claim email from one day to the next within the five day claim window, on the last day of the claim window, subjects still face a choice between claiming and not claiming the payment.

Distinguish three time periods: at $t = 1$, subjects choose between contingent and automatic payments. At $t = 2$, subjects have the opportunity to claim the contingent payment, and at $t = 3$ subjects receive the payment if they remember to claim the payment. We can capture overestimation about the effort subjects thought they would be willing to provide by supposing subjects overestimate their future discount factor. Let the actual discount factor between periods 2 and 3 (the claim window and the receipt of payment) be β. However, at period 1 when choosing between the two payment forms, let subjects believe that the discount factor between periods 2 and 3 will be $\hat{\beta}$. We can capture underestimation about the effort subjects thought would be required by distinguishing between period-1 anticipated costs of claim, $\widehat{costClaim}$, and period-2 actual costs of claim, $costClaim$. Conditional on remembering, subjects anticipate they will claim if $20\hat{\beta} > \widehat{costClaim}$, but will only actually claim if $20\beta > costClaim$.

It takes about a week for subjects to receive the check in the mail, so that β is approximately a weekly discount factor. A standard exponential discounter must have β near one,[7] but an individual with present-bias (e.g. a hyperbolic discounter) may have β substantially below one. Suppose that $\beta = 0.8$ (a reasonable or even low value[8]).

Evidence indicates that anticipated costs of claiming, $\widehat{costClaim}$, were close to zero. This then implies that the anticipated probability of claim conditional on remembering was near one. Recall that the estimated $\underline{\hat{p}}$ is a lower bound on the subjective probability of claim: when subjects estimate high costs of claiming the payment, $\underline{\hat{p}}$ will be lower than the true subjective probability of claim (see Section 3.3). The similarly between $\underline{\hat{p}}$ and the stated subjective probability of remembering in Automatic Payment Sample indicates that the subjective estimates of the costs of claiming the payment were very low (see Section 4.3).

Because of the low costs involved in claiming the contingent payment, it is unlikely that subjects remembered they could claim the reward but chose not to. Claiming the contingent payment required low effort and could be done in less than five minutes: to receive \$20, subjects need only include their names in an email to the experimenters. If a subject does not claim conditional on remembering, it must be that $20\beta > costClaim$, or that the costs of sending the email were above \$16. The probability that subjects find it worthwhile to claim conditional on remembering (i.e. that $20\beta > costClaim$) is therefore also likely to be near one.

To make the most generous case against overestimation of the probability of remembering, continue to suppose that subjects think they are certain to make a claim conditional on remembering. Subjects' choices imply a subjective probability of remembering of 76%. Only 53% claimed. For the memory beliefs to be correct, the fraction of subjects that failed to claim conditional on remembering (the probability

7 If an exponential discounter had a weekly discount factor of 0.99, he would attach implausibly low value to the future. For example, his five year discount factor would be $0.99\hat{}(52 * 5) = 0.073$.

8 The measurement of time discounting is controversial. Frederick, Loewenstein, and O'Donoghue (2002) present a range of discount rate estimates from a number of experimental studies. The median annual discount factor was well above 0.8, implying that the weekly discount factor must also be substantially above 0.8. Also note that if subjects have access to credit markets and view money as fungible, the discount factor should be bounded by the rate of interest.

that $20\beta > costClaim$) must be about 30% ($= 1 - 0.53/0.76$). This would implausibly require that over a quarter of subjects faced a cost of more than \$16 to send an email to claim the payment (and moreover had predicted that the cost of doing so would be substantially lower!). For even partial awareness of discounting and costs of claiming ($\hat{\beta} < 1$ and $\widehat{costClaim} > 0$), the fraction of subjects that failed to claim conditional on remembering must be even higher for memory beliefs to have been accurate. This simple calibration indicates that overconfidence on the memory task is likely to be driven by overestimation of memory, rather than underestimation of costs or overestimation of effort provision.

5.2 Alternative Explanations

There are two alternative explanations for the results that do not rely on overconfidence: lack of experimenter credibility, and a hedging motive on the part of subjects. Both these explanations predict that the actual probability of remembering to claim the payment should diverge from the $\underline{\hat{p}}$ elicited from choices. Yet recall from Section 4.3, when subjects in the Automatic Payment Sample were asked how likely it was they would have remembered to claim the payment, their responses were similar to the $\underline{\hat{p}}$ elicited from choices.

5.2.1. Credibility. Lack of credibility could create biases in either direction. For instance, subjects may have been hesitant to take the contingent payment for fear that the experimenters would not honor the attempted claim (e.g. the experimenter might claim not to have received the email on time). In this case subjects would be even more overconfident than the prior analysis indicates. However, if subjects did not find the experimenters' promise to mail the automatic payments credible, the estimate of $\underline{\hat{p}}$ will be biased upwards. There are good reasons to believe that this bias, if any, is small. First, the aforementioned similarity between our estimate of $\underline{\hat{p}}$ and the subjects' statements of the probability they would remember to claim the payment indicates that the estimated $\underline{\hat{p}}$ is not too different from the actual $\underline{\hat{p}}$. Second, researcher credibility was likely to be high. At each university, the experiment only permitted students enrolled at that university to participate. Prominent researchers affiliated with that university were listed as contacts on the informed consent forms, further enhancing credibility. Subjects knew they would receive a reminder sheet, giving them a record of having participated in the experiment. Furthermore, the instructions indicated precisely the date on which the payments would be mailed.

Calibrations also indicate that implausibly low levels of credibility would be necessary to explain the observed overconfidence. Suppose a subject believed that if she took no action, the experimenters would only send the automatic payment with probability q, but that if they did not send it automatically but she remembered to contact them about it, they would certainly send the payment. Assume her costs of remembering and claiming the payment are small so it would be worth trying to remember to claim the payment even in the Automatic Payment condition, and set the costs to zero for simplicity.[9] Also let her probability of remembering to claim the payment be the same regardless of whether she is remembering to claim the contingent payment or the automatic payment. Then, modifying equation (1), the subject would choose the contingent payment of \$20 over an automatic payment of \$$x$ if and only if

$$\hat{p}\,(20) \geq qx + (1-q)\,\hat{p}x.$$

This would then imply that the true value of $\hat{p}$ is related to the estimated $\underline{\hat{p}} = x_{\max}/20$ as follows:

$$\frac{x_{\max}}{20} = \hat{p}_{\text{true}} + \frac{(1-q)\,(20-x_{\max})\,x_{\max}}{20\,(20-(1-q)\,x_{\max})},$$

where the second term is the bias term and $x_{\max}$ is the highest value of the automatic payment for which the subject chooses the contingent payment. For moderate levels of credibility, this bias term is extremely small: with $x_{\max} = 15$, and $q = 0.9$, the bias term is less than 0.02, clearly not large enough to explain the dramatic 0.23 gap between implied beliefs and realized probability of claim.[10] Even for rather dismal levels of credibility, this bias does not reverse the finding of overconfidence: if subjects thought there was only a 50% chance they would receive the payment automatically ($q = 0.5$), the bias term is only 0.15.

5.2.2. Risk and Insurance. Throughout this paper, $\underline{\hat{p}}$ has been derived assuming that subjects are maximizing expected value (or are risk averse). If subjects are risk seeking, then the estimated $\underline{\hat{p}}$ is biased upwards, but experiments typically find risk aversion even over small stakes. It is also possible, but very unlikely, that subjects had a hedging motive and chose the contingent payment as a form of insurance. In this scenario, subjects expect the stochastic cost of time to be low when marginal utility of consumption is high, and therefore choose the contingent payment even though it has lower expected value. This seems implausible. The cost of time to claim the payment is small (sending an email). Further, participants are full-time students whose opportunity cost of time is not likely to vary much with the marginal utility of consumption. Finally, the similarity in average levels of $\underline{\hat{p}}$ and the hypothetical memory estimate provides evidence that subjects were actually making the choice based on their probability of remembering, not purchasing insurance.

6. Conclusion

This study finds a substantial bias in subjects' forecasts: while their choices imply a forecast of a 76% claim rate, they claim only 53% of the time. Because of the

9 Recall that $\underline{\hat{p}}$ is a lower bound on the subjective probability of remembering because we do not observe the costs involved with claiming the payments. If these costs are large, then the actual subjective probability is much higher than the estimated $\underline{\hat{p}}$, indicating that subjects are even more overconfident and any bias due to lack of credibility is unlikely to overturn the results.

10 Similar or smaller figures are found for the range of values of $x_{\max}$.

low costs of claiming the payment, the overconfidence displayed here is likely to be overconfidence in memory rather than overconfidence in future effort. It remains to be tested whether this finding of overconfidence can be generalized to other populations or tasks. Although subjects in this experiment are likely to have repeated practice with natural prospective memory tasks such as paying bills on time or turning in papers before deadlines, the actual task used in the experiment may have been unfamiliar to the subjects. Subjects in the experiment faced incentives to reveal their preferences truthfully, but only received one of their choices about one-fifth of the time, and so it is possible (albeit unlikely) that overconfidence might be reduced if the stakes were higher or they had a higher probability of getting one of their choices.

The existence of overconfidence in prospective memory has welfare implications in many domains, as prospective memory tasks are widespread in the workplace as well as private life. The limitations of bounded prospective memory may be mitigated by the use of memory aids, but overconfident individuals will underutilize these aids relative to the optimum. This finding of overconfidence can enable economists to develop more accurate models of markets in which agents are not prompted to make a decision, but must remember to do so on their own. Further, recognizing that individuals may be overconfident can help economists avoid making mistaken inferences about preferences from observed behavior.

Supporting Information

Additional Supporting Information may be found in the online version of this article:

Appendix S1. Experimental materials (Word document).
Appendix S2. Experimental materials (PDF document).
Appendix S3. Paper results (.do document).
Appendix S4. Raw over-confidence data (.dta document).

Please note: Blackwell Publishing are not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

References

Bernheim, B. Douglas, Lorenzo Forni, Jagadeesh Gokhale, and Laurence J. Kotlikoff (2003). "The Mismatch between Life Insurance Holdings and Financial Vulnerabilities: Evidence from the Health and Retirement Study." *American Economic Review*, 93, 354–365.

Camerer, Colin, and Robin M. Hogarth (1999). "The Effects of Financial Incentives in Experiments: A Review and Capital–Labor–Production Framework." *Journal of Risk and Uncertainty*, 19, 7–42.

Camerer, Colin, and Dan Lovallo (1999). "Overconfidence and Excess Entry: An Experimental Approach." *American Economic Review*, 89, 306–318.

Choi, James J., David Laibson, and Brigitte C. Madrian (2006). "Why Does the Law of One Price

Fail? An Experiment on Index Mutual Funds." National Bureau of Economic Research Working Paper No. 12261.

DellaVigna, Stefano, and Ulrike Malmendier (2004). "Contract Design and Self-Control: Theory and Evidence." *Quarterly Journal of Economics*, 119, 353–402.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue (2002). "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature*, 40, 351–401.

Gabaix, Xavier, and David Laibson (2006). "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *Quarterly Journal of Economics*, 121, 505–540.

Gigerenzer, Gerd, Ulrich Hoffrage, and Heinz Kleinbölting (1991). "Probabilistic Mental Models: A Brunswikian Theory of Confidence." *Psychological Review*, 98, 506–528.

Guynn, Melissa J., Mark A. McDaniel, and Gilles O. Einstein (1998). "Prospective Memory: When Reminders Fail." *Memory and Cognition*, 26, 287–298.

Hoelzl, Erik, and Aldo Rustichini (2005). "Overconfident: Do You Put Your Money On It?" *Economic Journal*, 115, 305–318.

Holman, Jeffrey, and Farhan Zaidi (2004). "The Economics of Prospective Memory." Working Paper, University of California, Berkeley.

Madrian, Brigitte C., and Dennis F. Shea (2001). "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *Quarterly Journal of Economics* 116, 1149–1187.

Moore, Don, and Daylian Cain (2007). "Overconfidence and Underconfidence: When and Why People Underestimate (and Overestimate) the Competition." *Organizational Behavior and Human Decision Processes*, 103, 197–213.

O'Donoghue, Ted, and Matthew Rabin (2001). "Choice and Procrastination." *Quarterly Journal of Economics*, 116, 121–160.

Sandroni, Alvaro, and Francesco Squintani (2007). "Overconfidence, Insurance, and Paternalism." *American Economic Review*, 97, 1994–2004.

Scheinkman, Jose, and Wei Xiong (2003). "Overconfidence and Speculative Bubbles." *Journal of Political Economy*, 111, 1183–1219.

Soll, Jack, and Joshua Klayman (2004). "Overconfidence in Interval Estimates." *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 30, 299–314.

Svenson, Ola (1981). "Are We All Less Risky and More Skillful Than Our Fellow Drivers?" *Acta Psychologica*, 47, 143–148.

Weinstein, Neil (1980). "Unrealistic Optimism about Future Life Events." *Journal of Personality and Social Psychology*, 39, 806–820.