

Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox

ROBERT W. ROSENTHAL

Bell Telephone Laboratories, Murray Hill, New Jersey 07974

Received April 8, 1980; revised July 2, 1980

1. INTRODUCTION

The thesis of this paper is that finite, noncooperative games possessing both complete and perfect information ought to be treated like one-player decision problems. That is, players ought to assign at every move subjective probabilities to every subsequent choice in the game and ought to make decisions via backward induction. This view is in contrast with the game-theoretic approach of Nash equilibrium.

After expanding on this view for games in the abstract in Sections 2 and 3, attention is turned in Section 4 to an example due to Reinhard Selten, called the chain-store paradox, which possesses the flavor of a situation involving a predatory-pricing monopolist. It is argued that for the chain-store game the decision-analytic approach leads, under certain assumptions, to more realistic outcomes than the standard Nash-equilibrium approach.

2. AN ILLUSTRATIVE EXAMPLE

Consider the game tree in Fig. 1 to be played only once in which the outcomes are assumed to be expressed in U.S. dollars and x and y are dollar values known to both players. In words, if Player 1 chooses Left, he receives x dollars and 2 receives y dollars; if Player 1 chooses Right, then the outcome is either $(0, 0)$ or $(1 \text{ million}, 1)$ depending on Player 2's choice. Assuming that each player's von Neumann-Morgenstern utilities for the outcomes are ordered in the same way as the dollar values and assuming that the game is played under conditions of complete information (i.e., each player knows the rules, the dollar payoffs for both, von Neumann-Morgenstern utility images of both players' payoffs, the fact that the other player knows all of this, the fact that the other knows that he knows, etc.) what should Player 1 do? More directly, for what values of x and y would you as Player 1 be indifferent between your two choices? For




FIG. 1. Example 1.

most of those to whom I have posed this question, the answer depends on (among other things) what is known about Player 2. When I reply that Player 2 is known to be a very sensible person whose identity will always remain secret and to whom Player 1's identity will always remain secret, the answer invariably involves some x -value strictly less than 1 million.

I interpret this answer to mean that when the stakes are sufficiently high people will usually not act as though the hypothesis that other people are utility maximizers is absolutely dependable. Moves of other players simply can never be anticipated with certainty. If that is the view of real game players, what are its implications for the theory of games in which nearly all of the models and solution concepts seem to be based on the idea that all players can be assumed to be acting in their own best interests?

One might respond that the hypothesis of complete information is too strong to be ever realized in life, that utility payoffs of others (for example) can never be known with certainty and that 1's selection of Left over Right when $x < 1$ million simply indicates a natural equilibrium strategy in the imagined incomplete-information game. My feeling is that while probably no games are played under conditions of complete information in the real world, there are some situations in which the information is complete enough that the situations can reasonably be modeled as complete-information games; and that for any extensive-form model it should be possible to imagine that and reason as though the model corresponded closely enough to some hypothetical game situation. In my view, therefore, appealing to the lack of realism of the complete-information game would not be sufficient to explain preferences of Left over Right in Example 1.

It might also be pointed out that the Left strategy for Player 1 is in fact part of a Nash equilibrium whenever $x \geq 0$ and that game theory therefore admits considerations of the aforementioned sort. But 1's Left is only in equilibrium against 2's Left strategy; a situation I interpret to mean that 1 believes 2's (possibly nonexplicit) threat to play Left, a threat which can only be believed if 2 is possibly not a utility maximizer. Note that this strategy combination is in equilibrium whenever $x \geq 0$ and is otherwise insensitive to the magnitude of x . (Nash equilibria containing threats which do not get carried out are examples of what are called imperfect equilibria. In some recent game-theoretic work it has been suggested that attention should focus only on perfect equilibria or special subsets of the set of perfect equilibria. See [3, 7, 8].)

The point of view which I shall adopt is that Player 1 should consider his Right strategy to be a risky one and that the risk should be treated in the same way that it would be if Player 2 were replaced by chance in this game. That is, Player 1 should arrive at some subjective probability distribution over 2's strategy choice and should play his Right strategy if the expected utility to him from the lottery exceeds his certain utility for his Left strategy. Moreover, I will argue in Section 3 that such an approach is appropriate for the analysis of all finite games with perfect information. Note that the uncertainty should concern Player 2's choice, not his utilities for the various outcomes. Although this distinction is not significant in Example 1, where 1's positing appropriate uncertainty over Player 2's payoffs and assuming he does not carry out threats can be seen to be equivalent to 1's positing uncertainty over 2's choice, the distinction does become significant in multi-move games.

3. PERFECT-INFORMATION GAMES

A finite n -person extensive-form game with perfect information consists of:

1. A tree with a distinguished node (called the origin). Node b follows a if there is a path from the origin to b through a . A terminal node is a node with no followers.
2. A partition of the nonterminal nodes into sets labeled $0, 1, 2, \dots, n$. A node in set $i > 0$ corresponds to a move by Player i . A node in set 0 is a move by chance.
3. For every node in set 0 , a probability distribution over the immediate followers of that node. The interpretation should be clear.
4. For every terminal node, a vector in R^n . The i th component of the vector is interpreted as Player i 's von Neumann–Morgenstern utility for the outcome represented by the terminal node.

For games with perfect information, no extra informational complications arise. When a player is to choose an immediate follower from any of his nodes, he knows for which node he is making the choice. Thus, he knows perfectly all the choices which have been made at preceding nodes. There are no possibilities of secret or simultaneous moves. The assumption of complete information is that everything about the extensive form is known to all the players of the game.

A pure strategy for Player i (>0) is a function which assigns to each of Player i 's nodes one of its immediate followers. A Nash equilibrium in pure strategies is a pure strategy combination (one strategy for each player) with the property that no player can improve his expected utility payoff by

switching to some other pure strategy if the strategies of all other players are held fixed. It is well known (see, e.g., [1]) that every finite perfect-information game possesses at least one Nash equilibrium in pure strategies and that such an equilibrium can always be constructed by working backward through the tree, selecting a best (defined inductively) immediate follower at every personal node and taking expected utilities at every chance node. Nash equilibria obtained in this way will be termed principal equilibria. Nonprincipal pure-strategy equilibria may exist, as has already been observed in Example 1. Nash equilibria in randomized strategies (i.e., probability distributions over pure strategies), which have been widely studied, are of lesser interest for games with perfect information; since there seem to be no intuitive reasons to randomize consciously when the outcome of the randomization will be evident to all players who have subsequent moves.

Although a player may not consciously intend to randomize and may in fact intend to reason inductively in order to select his strategy in a game with perfect information, it may be that for some reason his intention is not realized. If players with preceding moves realize that this may be the case, they would be wise to consider it when making their choices.

Why might players believe that players moving subsequently might deviate from their principal-equilibrium strategies? (Incidentally, I do not rule out the possibility that a player may view his own subsequent choices as random.) One possibility is a mistake. Another is that subsequent players may decide not to analyze the tree completely because the time and effort required are not justified by the potential gains. (In this case it could be argued that the game in question is only part of the more complicated problem in which decisions about time and effort are also made. Such considerations quickly lead to unmanageably large models and even questions about infinite regressions. It seems far simpler to study the game of perfect information with appropriate allowance made for moves which are not part of principal equilibria.) A third possibility (though related to the previous one) is that the player in question recognizes that the game is actually played under conditions of incomplete information, but that a model incorporating the uncertainties would be unmanageable for him. To keep matters simple, he assumes a game with complete information but modifies his principal-equilibrium calculations in a manner he feels to be intuitively appropriate. Still a fourth possibility is the belief that other players may be acting according to some other view of what is rational. This view may not be known with certainty.

Consider Example 2 (see Fig. 2), where the origin is at the left. If Player 2 is actually called upon to move in this game, he has clear evidence that Player 1 is not playing according to any Nash-equilibrium strategy. Should he ignore this evidence, or should he consider the possibility that Player 1




FIG. 2. Example 2.

might not act in his own (and hence Player 2's) interest at his next move? (See [2, pp. 80–81] for a related discussion.)

Because the various reasons given for imputing estimates of the subsequent move probabilities involve a considerable degree of subjectivity, I shall not suggest a universal scheme whereby players may logically deduce these estimates. Rather, I appeal to the decision-analysis paradigm in which each player arrives at subjective probability estimates by considering what objective lotteries he views as indifferent to the move in question by the player in question (see, e.g., [5]), where the outcomes of the objective lotteries are the same, respectively, as the outcomes of the personal moves. This is admittedly a major incomplete aspect of the method of analysis which I am proposing. On the other hand, the arguments which I put forth in this paper may be viewed as an appeal to include, at least in situations with perfect information, whatever aspects of a real situation are left unmodeled by the extensive form. Without making very special assumptions about the situation in question, it is difficult to imagine that a useful theory is possible. Indeed, in Example 2 it would seem to be a hopeless task to say anything sensible about Player 2's decision without extra unmodeled information. (It can be argued that all aspects of a situation including, for example, perceived psychological attitudes of other players should be expressed in the extensive form. Although possible in principle, this would seem prohibitively complicated in practice.)

In some cases, it is possible to gain new qualitative insights from this approach without being too specific about subjective probabilities. Consider Example 3 (see Fig. 3). In Example 3 the leftmost node is the origin, and each player's choices at each node are called Right and Down, respectively. All the Nash equilibria in this example involve Player 1 picking Down on his first move. The principal equilibrium has both players selecting Down on their first move. Suppose now as an illustration that both players have identical views about each others' propensities to deviate from principal Nash behavior and that the views of both are common knowledge (i.e., each knows




FIG. 3. Example 3.

that the other knows, etc.). These views are that the probability assigned to the better of two choices, at any stage, is $\min(1, 0.5 + 0.4D)$, where D is the utility difference between the two choices. (Of course, a von Neumann–Morgenstern utility construction does not fix a scale for measurement. One must imagine that the coefficient (0.4) is the one that applies for the utility scale chosen in the representation of the game given above.) Thus, at the last node, although Player 2 would (no doubt) like to choose Down with probability 1; it is predicted that he will choose Down with probability 0.9 and right with probability 0.1. At the penultimate node, therefore, Player 1 can expect utility of 8 from Down and 7.3 from Right. He would therefore also elect Down but is predicted as choosing Down with probability 0.78 and Right with probability 0.22. At the preceding move 2 expects 9 from Down and $(0.78(8) + 0.22(11))$ from Right if he believes himself incapable of subsequent error or $(0.78(8) + (0.22)(10.9))$ from Right if he ascribes the common view of error to his own moves. Under either hypothesis, assuming common knowledge, the probabilities of Right continue to increase as the induction proceeds back through the tree, becoming one before the origin is reached.

Although it might be difficult to believe in the specific construction for the subjective probabilities given above and also in the common-knowledge assumption; it is clear that if all players consider the subjective probabilities to be sufficiently dependent on the utility differences and if the form of these considerations is common knowledge (even if the specifics of the other players' functions are not common knowledge) then there is a class of games for which this type of compounding effect occurs. As I shall argue in the next section, some important economic models fall into such a class. Of course, even if subjective probabilities are not related to the size of the utility differences but depend on other features, the moves prescribed by a decision-analysis approach may differ from those of every Nash equilibrium.

It is the perfect-information assumption which enables us to use the logic of backward induction. Still, it is possible that similar reasoning could be used (to some extent at least) in extensive-form games with imperfect information. If we believe that Nash equilibria are not reasonable or realistic in perfect-information games, the same could be true for the same reasons in games with imperfect information. If a game is "deep" (in the sense that some paths from the origin to terminal nodes contain many arcs) but only a small part of the tree is affected by information imperfections, then it might be possible to look at only those small parts of the tree as subproblems in traditional game-theoretic ways and to look at the grand problem more according to the decision-analytic mode.

4. PREDATORY PRICING AND THE CHAIN-STORE PARADOX

Consider a single-product industry containing a single firm setting a monopoly price and making corresponding profits. If another firm enters the industry, there is enough business for both to survive, but profits for the monopolist will be severely cut. A possible strategy for the monopolist is to cut price whenever a new firm enters to a level low enough that the entrant will sustain losses in every period. This strategy presumably reduces the monopolist's profits even further for those periods when the low price is in effect (perhaps even causing losses for the monopolist). If the monopolist drives the entrant out, and if no new firm enters, however, the monopolist can resume his high-price, high-profit position. Price cutting to drive other firms out of an industry or to deter entry is called predatory pricing and is a subject of considerable interest to economists. The proper role of anti-trust procedures in controlling predatory pricing is a matter of some debate among both economists and law makers (see, for example, [4, 6]).

The question of whether or not and to what extent the threat of predatory pricing is likely to succeed in deterring entry under various conditions is a little-understood but important aspect of the debate. It is also a question with a strong game-theoretic flavor.

In [7], Reinhard Selten describes an example of a game, called the chain-store paradox (version one), which seems to capture the essential elements of the deterrence question. In Selten's game there is one chain store, Player A, with branches in 20 towns. In each town there is one potential entrant i ($i = 1, \dots, 20$). Furthermore, for each i there is only one date at which he can enter. These dates are sequential with Player 1's date first, then 2's, etc. If Player i enters on his date, A with full knowledge has the choice of whether to respond aggressively or cooperatively. The payoffs to Players i and A (although it is only a partial payoff to A who is concerned with the sum of his 20 payoffs) are given below:

i 's decision	A's decision in period i	i 's payoff	A's payoff
In	Cooperate	2	2
In	Don't cooperate	0	0
Out	—	1	5

At Player i 's entry date, he is fully aware of all previous moves.

There are two features of this game which deserve comment with respect to the predatory-pricing interpretation. One is the assignment of unique entry dates to each firm. This assumption seems to have been made to keep the

specification of the game as simple as possible and does not appear to distort seriously the deterrence question. The second feature is the finite-time horizon assumption. Firms are often assumed to plan for the infinite horizon (with discounting where appropriate). Furthermore, it is well known that in many repeated games there is a qualitative difference between the structure of the set of Nash equilibria in the finitely-repeated case (even with a sufficiently large number of repetitions) and that in the infinitely-repeated case. Since all of the discussion thus far has been directed at finite-move games, it would be improper to argue that any insights gained by reasoning inductively on the finite, chain-store game are valid for a situation in which players perceive an infinite horizon. But it would also seem that there should be very little difference between what a firm would actually do early in a large-but-finite horizon predatory-pricing situation and what that same firm would do in a similar infinite-horizon situation.

The thrust of Selten's analysis of the chain-store game is that while the principal Nash equilibrium of the game requires all Players i to enter and A to cooperate, one's intuition suggests perhaps that the threat of the aggressive response might suffice to deter entry, at least until late in the game. Selten then proceeds to an interesting discussion of the game from many different points of view.

The point which I wish to make is that there is a considerable similarity between the chain-store game and Example 3. If the discussion of the previous sections has been convincing, then the approach advocated here for games with perfect information may be of use in resolving the "paradox" and in analyzing predatory pricing and other deterrence situations.

Let us look at the chain-store game under the assumption that the size of utility differences affects subjective probabilities in the same direction as in Example 3. At each of Player A 's last nodes he gains two units of utility by picking the cooperative response. Suppose, however, that Player 20 views that choice as less than certain. At 20's move, therefore, his expected utility gain is less than one from choosing In. Player A therefore has an expected gain of less than 2 if he responds aggressively to 19's choice of In. If 19 considers the aggressive response to him to have higher probability than 20 considered the aggressive response, then he has even less to gain than 20 by choosing In. As in Example 3, the effects may quickly compound if the probabilities are appropriate monotone functions of the expected utility differences. The reader will be spared any further details. It should be obvious that for a wide class of methods for forming subjective probabilities, the outcome of the game will be that entry is deterred for a large portion of the potential entrants.

REFERENCES

1. H. W. KUHN, Extensive games and the problem of information, in "Contributions to the Theory of Games, Vol. II" (Kuhn and Tucker, Eds.), Annals of Mathematics Studies, No. 28, pp. 193-216, Princeton Univ. Press, Princeton, N. J.
2. R. D. LUCE AND H. RAIFFA, "Games and Decisions," Wiley, New York, 1957.
3. R. B. MYERSON, Refinements of the Nash equilibrium concept, *Internat. J. Game Theory* 7 (1978), 73-80.
4. R. A. POSNER, "Anti-Trust Law: An Economic Perspective," Chap. 8, Univ. of Chicago Press, Chicago, 1976.
5. H. RAIFFA, "Decision Analysis," Addison-Wesley, Reading, Mass., 1968.
6. F. M. SCHERER, "Industrial Market Structure and Economic Performance," Chap. 8, Rand McNally, Chicago, 1970.
7. R. SELTEN, The chain-store paradox, *Theory and Decision* 9 (1978), 127-159.
8. R. SELTEN, Reexamination of the perfectness concept for equilibrium points in extensive games, *Intern. J. Game Theory* 4 (1975), 25-55.