

THEORY OF THE RELUCTANT DUELIST

By DANIEL ELLSBERG*

We wish to find the mathematically complete principles which define "rational behavior" for the participants in a social economy, and to derive from them the general characteristics of that behavior.

Thus von Neumann and Morgenstern defined their goal in the *Theory of Games and Economic Behavior*.¹

Their analysis of the problem, their model, their approach to it were brilliantly conceived; they invented an array of new concepts and new techniques of analysis to examine what has come to be known as the game situation. But they aimed beyond. With respect at least to the special situation of the "zero-sum two-person" game they offered not merely a new way of looking at an old problem but "a precise theory . . . which gives complete answers to all questions."²

Rarely is it kind to remind a theorist of such a statement ten years later. Yet the fact is that for more than a decade their solution of the special problem has stood without serious challenge. Critics of game theory have indeed questioned the assumptions, the concepts, the importance of the model; but they have published few complaints about the conclusions drawn from the model. Game theorists have gone on to new problems. They have abandoned some assumptions and developed more complex and versatile models. But they have rarely derived results so elegant, determinate or general as those claimed for the two-person zero-sum game. The solution associated with that model has come to represent game theory: its most solid achievement, the best measure of its promise. No one has cared to question its status as *the* theory of the two-person zero-sum game.

Is it a satisfactory theory, within the limits of its own model and assumptions? I do not believe it is. I do not think it would be useful in

* This article was originally part of an honors thesis entitled "Theories of Rational Choice Under Uncertainty" submitted at Harvard College, April 1952. The author wishes particularly to thank John Chipman (who supervised the original writing), Carl Kaysen, Wassily Leontief, Martin Shubik and Oskar Morgenstern for their valuable criticism. The author is presently serving with the United States Marine Corps, on leave of absence from Harvard University.

¹ John von Neumann and Oskar Morgenstern (Princeton, 1944), p. 31. All citations in this paper will refer to this work, unless otherwise noted.

² P. 101.

predicting behavior in the situations it considers, nor does it seem acceptable as a general norm of behavior.

Von Neumann and Morgenstern approach a particular context of rational choice under uncertainty: in which the outcome of an action is uncertain because it depends on the interaction of a small number of conflicting wills. Because rational choice has been undefined for situations involving uncertainty, orthodox theory, premised on rational behavior, has been correspondingly "indeterminate" for these situations.

By *certainty* is meant a situation in which each available action is associated in the actor's mind with a single, certain consequence. The rule of rational choice under certainty, to which we will refer frequently, requires him to choose that action whose consequence he most prefers. The rule is meaningless when an individual must act under uncertainty, *i.e.*, when he associates with a given action a *set* of possible outcomes, some of which may be favorable and some unfavorable.

It is not difficult to devise various rules of choice that can be applied in this situation, any one of which might be termed, more or less arbitrarily, "rational." To discover a rule with general usefulness in predicting or prescribing behavior under uncertainty is something else again. It is this at which von Neumann and Morgenstern aim. "The superiority of 'rational behavior' over any other kind is to be established . . . for all conceivable situations—including those where 'the others' behave irrationally, in the sense of the standards which the theory will set for them."³

We might consider a principle to be a "useful" definition of rational choice under uncertainty if most people who were rational under certainty would reject any decision inconsistent with the principle.⁴ With this criterion we can probe the von Neumann-Morgenstern conclusions. If it should appear that a large number of reasonable people will accept some decisions inconsistent with the particular rule that the authors propose, and if their reasons are not random or foolish, then the von Neumann-Morgenstern principle could not be satisfactory as the unique definition of rational behavior in the game situation.

Some familiarity with the model and concepts of the zero-sum two-person game will be taken for granted, but the abstract model can be described briefly. Player A selects a strategy i from the set of strategies open to him by the rules of the game. Simultaneously, in ignorance of A's choice, player B selects a strategy j from his set of admissible strategies. Then, after the choices are revealed, A receives an amount a_{ij} and B receives an amount $-a_{ij}$ (or, B pays A an amount a_{ij}). The

³ P. 32.

⁴ This proposition has been adapted from one in an unpublished paper by Jacob Marschak.

subscripts indicate that the payoff is a function of both strategies. The rules of the game prescribe a pair of outcomes corresponding to each possible pair of strategies. The sum of the outcomes is zero; what one player wins, the other loses.

In this model, each player makes but one "move." Thus, the analysis applies directly to such simple games as matching pennies, in which each player chooses between the alternatives, heads or tails. To generalize the results to more complex games such as poker or chess, the authors interpret the player's single move as the choice of a strategy, a concept which they define: "a plan which specifies what choices he will make in every possible situation, for every possible actual information which he may possess at that moment. . . ."⁵ When both players have chosen strategies in this sense, the outcome of the game is determined. Thus, complex games can be analyzed in static terms, as though the outcome were determined by a single choice on the part of each player.

To "divide the difficulties" of the analysis, the authors make some important simplifying assumptions. First, the outcomes are represented not in utilities, cardinal or otherwise, but unequivocally in money.⁶ Second, von Neumann and Morgenstern abstract from uncertainties concerning the rules of the game. Each player knows with certainty: (a) what strategies he is allowed; (b) what strategies his opponent is allowed; (c) the outcome corresponding to any pair of opposing strategies. Finally, there is a significant tacit assumption that each player knows his opponent has the same rules of the game in mind.

The strategies and payoff function, which comprise the rules of the game, can be represented by a matrix, each row corresponding to one of A's possible strategies, each column, one of B's strategies, the matrix elements being the payoffs corresponding to pairs of opposing strategies. There is no need to show B's outcomes explicitly, since they are merely the negatives of A's.

The model described above is the "normalized" version of the game. It expresses just those elements of uncertainty which von Neumann and Morgenstern wish to emphasize. Since each player must choose in ignorance of his opponent's choice, and since the outcome of any strategy depends on that unknown choice, there is a *set* of possible outcomes corresponding to each possible strategy, rather than a single, certain outcome. The problem is to prescribe a unique "rational" choice among these *sets* of uncertain outcomes.

In special cases the choice may be easy. If the payoff function hap-

⁵ P. 79.

⁶ P. 8.

pens to be such that the outcome of one particular strategy is better than the outcome of another for every one of the opponent's possible strategies, the first will be said to "dominate" the second. In terms of the matrix, if each element in one row is greater than the corresponding element in another row, the first strategy dominates the second. To choose a strategy which is dominated by another would be to accept an outcome which is certain to be less favorable than if the dominant strategy were to be chosen. The rule seems indicated that the rational player will never choose a dominated strategy.

The hard choices come when (a) sets of outcomes overlap, so that the opponent's choice determines whether the outcome of one strategy is better or worse than that of another, and when (b) the opponent's choice is uncertain, being made simultaneously with one's own. These are conditions of the normalized game, which is the target of the analysis.

Von Neumann and Morgenstern approach this target indirectly, via two models which depart from the conditions of the normalized game and which are in themselves of much less significance. They explain their approach:

The introduction of these two games . . . achieves this: It ought to be evident by common sense—and we shall also establish it by an exact discussion—that for [these games] the "best way of playing"—*i.e.* the concept of rational behavior—has a clear meaning.⁷

In one of these modified games, called the minorant game, A must make his choice first, after which B chooses in full knowledge of A's choice. Since B, in this game, acts under certainty, the basic principle of rationality under certainty prescribes his choice. Given strategy i by A, then to each strategy available to B there corresponds a single, certain outcome, and his unique rational choice is that strategy associated with the outcome: $\text{Min}_j a_{ij}$, where i is given.⁸ That is, he should pick the column corresponding to the minimum element in the row selected by A.

Now if A, moving first, knows literally nothing about B's "mentality," then A must choose under uncertainty. But in this model, if A has the relatively small scrap of information that B is rational under certainty, then A too acts under certainty. Operationally, the statement that B is known to be rational under certainty is equivalent to the statement that B, moving second, is certain to choose the minimum element in any row picked by A. If A thus considers it impossible

⁷ P. 100.

⁸ B will be described as minimizing a_{ij} , since his outcome is $-a_{ij}$. He could equally well be described as maximizing $-a_{ij}$.

that a given strategy should have any but its minimum outcome, the rule is inevitable that it is irrational for A to pay any attention to the $m(n-1)$ matrix elements which are not row minima. A should choose the "maximin" strategy corresponding to the largest of the row minima.

In the second special model, the majorant game, B must choose before A, who then chooses with certainty of the outcome. As above, if B does know A to be rational under certainty, this is equivalent to knowing that elements which are not column maxima are not possible outcomes. Hence under this special assumption B also acts under certainty, associating a single, certain outcome ($\text{Max}_i a_{ij}$, for given j) with each strategy. The only strategy which is rationally consistent with his belief about A is his "minimax" strategy, *i.e.*, the strategy corresponding to the outcome $\text{Min}_j \text{Max}_i a_{ij}$, which guarantees him the best of the "possible" outcomes.

Thus, the authors' assertion is justified. In these special models, in which one player possesses, and is known to possess, knowledge of the other's choice, certain rules of rational choice for both players do appear uniquely valid.

In the normalized game, to which we now return, neither player knows with certainty his opponent's choice beforehand, since both choose simultaneously. The principle of choice which von Neumann and Morgenstern propose is essentially this: each player should choose *as though* he were moving first in a minorant (or majorant) game, and *as if* he were certain that his opponent were rational and informed. Thus, player A should consider only the minimum element in each row: *i.e.*, the worst that could happen to him under that strategy. He should then choose the strategy with best minimum outcome. This particular outcome may be expressed as $\text{Max}_i \text{Min}_j a_{ij}$, his maximin strategy; the corresponding policy for player B is to choose the column with the lowest maximum outcome ($\text{Min}_j \text{Max}_i a_{ij}$), or minimax.⁹

The nature of this prescription differs markedly from that applying in the minorant or majorant games. In those, the first player knows that the second acts with certainty of the outcome, and knows (we assume) that the second is rational under certainty; thus he too acts under certainty. The conclusion that he should choose his minimax strategy then follows directly from the principle of rational choice under certainty (there is, in fact, no problem of uncertainty). It is

⁹ For convenience, this rule of choice for the normalized game will hereafter be referred to as the minimax principle. Either player may be said to choose his minimax strategy—*i.e.*, the one prescribed by this principle—though the context may indicate, for a particular player, that a maximin strategy is involved.

precisely this property that is missing from the normalized game. There, knowledge that one's opponent is rational under certainty has no immediate bearing on the outcome to be expected from different strategies, for it is known that the opponent himself must act under uncertainty. The fact that one's own choice is hidden from the opponent means that the opponent's choice, and hence the outcome of one's strategy, is uncertain.

Von Neumann and Morgenstern instruct the player in the normalized game to act as if he were certain of the consequences, although in fact he is not and cannot be certain. Note that these particular "as if" assumptions are not *mandatory* on the players merely because of the game situation and the conflict of interest between them. Each player knows that his opponent would like to inflict the maximum possible loss on him; but he also knows that his opponent, moving simultaneously and in ignorance of his own choice, cannot be certain of succeeding.

Unless the player believes that his opponent is gifted with extra-sensory perception, the knowledge that he is hostile, "reasonable" and informed cannot make his strategy certain. Uncertainty is a state of mind, a property of belief or expectation; if it is present it cannot simply be "assumed away." The "as if" or minimax policy proposed by von Neumann and Morgenstern is not a method for exorcising uncertainty but is one among several principles for acting in the presence of uncertainty. Its provisional nature is mentioned here not as a criticism but because that aspect is obscured in *The Theory of Games*. Von Neumann and Morgenstern leave the impression that the minimax principle for the normalized game follows logically (and inevitably) from the similar principle in the minorant and majorant games: hence, that it is derived eventually from the notion of rational choice under certainty.

They do not, indeed, begin in that vein. They first present the advantages (from a conservative point of view) of the minimax principle, concluding: "It is *reasonable* to define a good way for 1 to play the game" as that strategy which guarantees him at least the maximin outcome. Similarly, "it is reasonable to define a good way for 2 to play the game as one which guarantees him a gain" corresponding to the minimax outcome.¹⁰ Or as paraphrased above, it is reasonable to define a good way for a player to behave: play as though moving first in a minorant or majorant game. This method might indeed be considered sound conservative behavior. But the authors continue:

¹⁰ Both quotations are from p. 108; italics added.

“So we have:

(14:C:a) *The* good way (strategy) for 1 to play the game” (my italics) is maximin. And:

“(14:C:b) *The* good way (strategy) for 2 to play the game” (my italics) is minimax. The next paragraph begins: “Finally, our definition of *the* good way of playing, as stated at the beginning of this section, yields immediately. . . .”¹¹ The fact is that their statement at the beginning of the section did not define *the* good way of playing. It defined *a* good way. Yet the authors are ready to start the next section with: “(14:C:a) — (14:C:f) . . . settle everything as far as the strictly determined two-person games are concerned.”¹²

It is not to confront the authors with a petty lapse that the metamorphosis of *a* into *the* has been plotted in detail. That passage plays no minor role. Keystone of the whole “determinate” theory of the normalized game is a uniquely valid principle of rational behavior: enthroned, in the citation above, by a bit of sleight-of-hand. It cannot simply be taken for granted (in fact, it does not seem to be true) that what is uniquely reasonable in the minorant or majorant games will still be uniquely reasonable in the normalized game.

Rejecting any such easy conclusion, it still remains to judge the von Neumann-Morgenstern principle on its merits. The maximin rule does offer a type of security—the certainty of achieving an outcome which is at least better than the worst possible (*i.e.*, the lowest element in the matrix). But this security is purchased at a price. Along with it goes the certainty that the outcome will not *exceed* a certain sum (namely, the best element in the set containing the maximin outcome). This upper limit may be only slightly better than the maximin outcome, which in turn may be only a shade better than the worst possible outcome. At the same time, other strategies may offer the possibility of dazzlingly superior outcomes, combined with minimum outcomes barely below the maximin. With such a payoff function, it is not obvious that simple “reasonableness” prescribes uniquely the choice of the maximin strategy.

The key question of this paper may be phrased: Is it useful to call a player irrational because he decides to use a nonminimax strategy? Consider the payoff matrix shown in Figure 1. In this game, A’s maximin strategy is A-2; B’s minimax strategy is B-2. According to von Neumann and Morgenstern, these are “the rational” choices for A and B. Any other choice would expose the player to the chance of losing 10. On the other hand, the “rational” strategy also guarantees that the

¹¹ All quotations are from *loc. cit.*; italics added.

¹² P. 109.

player will not get more than 0.

Would most people who were rational under certainty reject any other choice of strategy? Suppose that A were to play a nonmaximin strategy, A-1 or A-3. If B played his "rational" strategy B-2, A would do exactly as well as if he had used his own "rational" strategy A-2. If B were not certain to use B-2, then A would stand to win or to lose 10. Player A might prefer this uncertainty to the certainty of winning 0.

A similar argument holds for B. In this game, both might use non-minimax strategies even though each knew his opponent to be rational under certainty and informed as to the payoff matrix. And there seem to be no convincing grounds for saying that these choices would be unreasonable.

In this game there is no way for one player to be sure of "punish-

	B-1	B-2	B-3
A-1	10	0	-10
A-2	0	0	0
A-3	-10	0	10

FIGURE 1

ing" the other for using a "bad" strategy; in fact, to have a chance of inflicting any loss on the other he must use a "nonrational" strategy himself.¹³ There is another implication that deserves some thought. In this game¹⁴ the behavior of a man "rational" in the von Neumann-Morgenstern sense would be unaffected if every element in the payoff matrix were multiplied by a constant. But are not most people interested in comparing the differential gains that they might make (by choosing a nonminimax strategy over a minimax strategy) with the differential losses they would risk? A player who was willing to accept the uncertainty of receiving either 10¢, 0¢ or -10¢ might be unwilling to risk the loss of \$100, even if combined with the possibility of winning \$100.¹⁵ Such a player, in contrast to the von Neumann-

¹³ This matrix is by no means a mere oddity. Von Neumann and Morgenstern spend considerable time analyzing games with precisely these characteristics; e.g., see the case on page 164: "If the opponent played the good strategy, then the player's mistake would not matter."

¹⁴ And in every "specially strictly determined" game, as defined below.

¹⁵ This point would not be met by replacing the money outcomes which the authors do use by "von Neumann-Morgenstern utilities," of the sort they discuss in their opening pages. The latter are relevant, if they can be found at all, only to situations involving

Morgenstern "rational" player, is taking into account outcomes other than minimum payoffs. He does not believe, and is not assuming, that his opponent is certain to succeed in enforcing the minimum. Is this irrational? Unlike the minorant and majorant games, there is no clear basis for certainty that the nonminimum outcomes are impossible.

It would appear that von Neumann and Morgenstern fail their own criterion; their rule of rational behavior fails to be superior in face of the possibility that the opponent may behave "irrationally." They do make an effort to avoid this test:

It is possible to argue that in a zero-sum two-person game the rationality of the opponent can be assumed, because the irrationality of his opponent can never harm a player. Indeed, since there are only two players and since the sum is zero, every loss which the opponent—irrationally—inflicts upon himself, necessarily causes an equal gain to the other player.¹⁶

The defense is inadequate. Their conception of "harm" seems to exclude any element of "opportunity cost," "regret," any notion of the pain incurred in passing up a real chance of great gains or in discovering, afterwards, that one could have done much better than he did (by risking slightly worse). The very mention of the possibility that an opponent will violate any given set of rules suggests that any element in the whole matrix is *possible*. If there is indeed a chance that the opponent will make a "mistake," why not help him to inflict a *large* loss on himself?

The particular game discussed above belongs to a class of games known as "specially strictly determined."¹⁷ With respect to these games at least, von Neumann and Morgenstern regard it as obvious and unquestionable that their minimax principle is solely rational. The criticisms presented so far might tend to unpin such faith in the unique claims of the rule, even in this most favorable context. It is time now to consider their efforts to extend the application of their principle and their concept of the "solution" or "value" of a game.

In the minorant and majorant games the assumptions (a) that both players are rational under certainty, and (b) that the player moving first knows that his opponent is rational under certainty, are sufficient to make the outcome of the play uniquely determined. Given these assumptions, the outcome to A in the minorant game will be $v_1 =$

risk, *i.e.*, known probabilities. (See "Classic and Current Notions of 'Measurable Utility,'" *Econ. Jour.*, Sept. 1954, LXIV, 528-56.) The situation described so far involves no probabilities. The authors insist that it is conceptually impossible to measure uncertainty as to the opponent's choice in terms of numerical probabilities.

¹⁶ P. 128.

¹⁷ Defined below.

$\text{Max}_i \text{Min}_j a_{ij}$ and the outcome to B is $-v_1$. In the majorant game (A moving second) the outcome to A is $v_2 = \text{Min}_j \text{Max}_i a_{ij}$, the outcome to B is $-v_2$. These payoffs are plausibly defined as "values" of the games for the players in two distinct senses: (a) if the assumptions apply, these are the outcomes that will actually result; (b) given the assumptions, they represent the maximum amounts which the players rationally should be willing to pay for the privilege of playing the game.

It can be proven that v_1 (maximin) is always less than or equal to v_2 (minimax). If $v_1 = v_2$, $\text{Max}_i \text{Min}_j a_{ij} = \text{Min}_j \text{Max}_i a_{ij}$, a "saddlepoint" is said to exist in the payoff function. This condition is of no interest at all in the minorant and majorant games, so far as the players' choices are concerned. However, it does play a role in the authors' attempt to derive a numerical value of a play in the normalized game.

They follow two parallel lines of argument. The first proceeds thus: (a) definite "values" can be assigned to the minorant and majorant games: (b) moving first is less advantageous than moving second, and moving simultaneously (as in the normalized game) must lie in between; (c) therefore if a "value" for the normalized game can be found at all, it must lie between the values of the minorant and majorant games, *i.e.*, v must be between v_1 and v_2 ; (d) thus, if the payoff matrix has a saddlepoint, $v_1 = v_2$, this outcome must constitute the unique value v of the normalized game.¹⁸

This is supported by an "heuristic" argument that the numbers v_1 and v_2 , defined as above but no longer associated with the minorant or majorant games, have a practical significance in connection with the normalized game. Although in this game both players choose simultaneously,

It is nevertheless conceivable that one of the players, say 2, "finds out" his adversary; *i.e.*, that he has somehow acquired the knowledge as to what his adversary's strategy is. The basis for this knowledge does not concern us. . . .¹⁹

They assert that in this case, conditions "become exactly the same as if" the game were a minorant game. Likewise, if player 1 "finds out" his adversary, conditions become "exactly the same as if" the game were the majorant game.²⁰ In either of these cases, they claim, the "value" of the normalized game becomes a "well-defined quantity":

¹⁸ A game with a saddlepoint corresponding to two strategies of the type we have considered so far ("pure" strategies) is said to be "specially strictly determined." The matrix in Figure 1 is an example.

¹⁹ P. 105.

²⁰ P. 106.

v_1 in the first case, v_2 in the second. Moreover, finding out is better than being found out, and the case in which neither occurs is in between; therefore the value of the normalized game must be bounded by v_1 and v_2 , and if they are equal it is uniquely determined.

In both arguments the key proposition is that v , if it can be defined at all, must lie between v_1 and v_2 , so that if they are equal the value is uniquely determined. Yet each overlooks a vital difference between the minorant-majorant games and the normalized game. In talking about the latter, we must assume that neither player is certain beforehand that he will be found out. After all, if B, for example, knew with certainty that he would be found out, it would not be "as if" A and B were playing the majorant game; they *would* be playing the majorant game.

Recognizing this, our whole previous discussion points to the possibility that B, *even though reasonable and informed*, might be "found" playing some nonminimax strategy. The potential reward to A of "finding out" B is thus not limited to v_2 (minimax). With foreknowledge in the normalized game, A might be able to achieve the very highest outcome in the matrix.

In other words, if A had a crystal ball that foretold B's strategies, a normalized game would not become "exactly the same as" a majorant game; it would be better. Even without the crystal ball, A might well prefer to play the normalized version of a game (which might offer the chance, if not the certainty, of higher payoff) to the majorant game: which is to say that v_2 is not a meaningful upper bound to the "value" of the normalized game for A.

Similarly, the possibility that B may find out A implies that the final outcome may range anywhere from maximin down to "minimin," the lowest element in the matrix. Under these circumstances, B might be willing to pay more than $-v_1$ to play the game. In terms of the payoff function in Figure 1, the "value" to either player of the minorant or majorant game corresponding to the matrix is 0. Yet either might be willing to pay, say, 1 for the privilege of playing the normalized game, which offers a chance of winning 10; and one of them might end up with an outcome of 10, rather than 0. Thus the value of the normalized game, either in the sense of actual outcome or of reasonable "worth" to the player, is not necessarily bounded by v_1 - v_2 .

This conclusion is crippling to the von Neumann-Morgenstern argument. If v_1 and v_2 separately have little relevance to the normalized game, they are no more relevant when they happen to be equal. This removes most of the interest from the question, regarded by von Neumann and Morgenstern as the central problem of the two-person zero-sum game, as to the general conditions for the existence of a

saddlepoint. Von Neumann's early solution to this problem does nothing to increase the *significance* of the saddlepoint.²¹

It is true that the existence of a saddlepoint is not entirely without interest, given certain hypothetical conditions:

1. If a player actually did expect with certainty that his opponent would "find him out," then he could do no better (in games with a saddlepoint) than to use his maximin strategy; in effect he would be playing a minorant game. But it would seem distinctly paranoid to feel certain that one would be found out on a single play (and von Neumann and Morgenstern insist again and again that their analysis is developed entirely with reference to a single play).²²

2. Although von Neumann and Morgenstern never consider aspects of a dynamic sequence of plays of two-person games, it might be argued that in such a sequence a saddlepoint would represent an equilibrium position, if both players had particular expectations. Specifically, if each player expected with certainty that his opponent would *go on* playing his minimax strategy, he himself would have no incentive to depart from his own minimax strategy. Aside from the fact that this seems definitely to involve a belief in the opponent's "rationality" such an argument, to earn a hearing, should be accompanied by assurances as to stability. If one player, for whatever reason, were in one play to depart from his minimax strategy, this in itself would tend to destroy the expectations which gave the saddlepoint the nature of an equilibrium. Henceforth both players would have incentives to use non-minimax strategies. Thus, let one small "shock" displace the outcome from the saddlepoint and it would show no tendency to return. Even as a dynamic equilibrium the saddlepoint would have no obvious interest, for it would be unstable.

Moreover, any appeal to dynamic considerations opens the door to new reasons for the employment of nonminimax strategies, strategies of a type the authors necessarily failed to consider: *e.g.*, strategies chosen to confuse the opponent as to one's own intentions, rationality or knowledge of the payoff matrix. Creating doubts by deliberately

²¹ He proves that if "mixed" strategies (probability combinations of the "pure" strategies we have considered) are regarded as included in the payoff function, every game will have a saddlepoint. But the saddlepoint is certainly no more significant when it corresponds to a pair of mixed strategies than when the game is "specially strictly determined," as is the example we have discussed.

²² "We have always insisted that our theory is a static one, and that we analyze the course of one play and not that of a sequence of successive plays" (p. 146); "We have repeatedly professed that our considerations must be applicable to one isolated play and also that they are strictly statical" (p. 189n). This paper has considered the theory on the authors' own terms, as applying to a single play. Besides providing a foundation for a dynamic theory, the case is far from trivial. There are many important real situations in which only one play of a game is possible.

erratic or "foolish" choices, one could tempt the opponent to pursue (for sound, profit-seeking motives) into the regions where big killings were possible. There would also be strategies designed to "find out" the opponent's future intentions or pattern of play. The fact is that there is not, in any real sense, a dynamic theory of games. To debate whether the saddlepoint may not be a stationary solution in this non-existent theory seems premature.

3. A saddlepoint represents an outcome v such that by acting appropriately A can be sure of receiving at least v no matter what B does, and B can keep A from receiving more than v no matter what A does. This fact is frequently cited as the basis for calling the saddlepoint the "value" of the normalized game. Yet to imply either that the saddlepoint will be achieved or that it represents the "worth" of the game to the players is to specify a very particular sort of player. Both must be "cautious pessimists," exclusively concerned with best *guaranteed* income.

Perhaps the bulk of recent work on the theory of the two-person zero-sum game has been concerned with the numerical computation of von Neumann's saddlepoint "solution." The question raised by our discussion is: Just what problem does this "solution" solve? We can make at least a partial answer as to what it does *not* solve. If no assumptions are made about the psychology of the players other than that they are reasonable and informed, then the saddlepoint represents neither an outcome that will necessarily be achieved nor the maximum amount that a player might reasonably offer to play the game. It fits neither of the two senses in which an outcome might usefully be defined as the "value" of a game.

In particular, we must reject the authors' statement: "Nor are our results for one player based upon any belief in the rational conduct of the other—a point the importance of which we have repeatedly stressed."²³ In nearly all games, if the possibility is considered that the opponent will not be "rational" in the von Neumann-Morgenstern sense, there will be nonminimax strategies which offer the chance of doing better than one could possibly do by choosing the minimax strategy. Surely there are many players, rational under certainty, who would regard these "bad" strategies as "superior," given this possibility.

At one point in their book, briefly, von Neumann and Morgenstern concede this aspect of their theory in a highly significant (and little noticed) sentence:

While our good strategies are perfect from the defensive point of view,

²³ P. 160.

they will (in general) not get the maximum out of the opponent's (possible) mistakes—*i.e.*, they are not calculated for the offensive.²⁴

This statement is decisive in establishing the character and the limitations of the theory. Swiftly leaving behind the admission so casually introduced, the authors hasten to point out: "It should be remembered, however, that . . . a theory of the offensive, in this sense, is not possible without essentially new ideas."²⁵ This may be no recommendation of the old ideas. Is it not likely that what they term a "theory of the offensive" is precisely what would appeal to many as a theory of rationality? More broadly, should there not be within a proper theory of rational behavior room for both offensive and defensive points of view? When did "rational" become synonymous with "defensive"?

Consider any matrix in which each row contains at least one negative element: *i.e.*, a game in which each one of A's strategies offers him a chance of some loss. The matrix in Figure 1 would be an example if its middle row were deleted (it would then look like the payoff function for matching dimes). Suppose that A were offered a new strategy which gave him the certainty of standing pat, neither winning nor losing. This would amount to adding a row of zeros to the matrix (as in Figure 1). To make the new strategy concrete, we might let the row of zeros correspond to the option of passing up particular plays of the game without penalty. If he were a disciple of von Neumann and Morgenstern, A's problem of choice would be solved. He would never play. No matter how slight the possible losses or how rich the potential gains with his other strategies, A would clutch at the row of zeros.

One might well ask: Why bother to play the game at all, if one prefers the certainty of zero to the chance of winning or losing? This question once was put to a prominent game theorist; his unconsidered reply, presumably intended as no more than a partial answer, was that in many situations one *must* play a game, even against one's wishes.

The vital orientation of game theory is implicit in that remark. If we should suppose—as no game theorist has in fact proposed—that the game models under consideration all represent uncertainty situations in which an individual is forced, reluctantly, to make decisions, the rationale for the minimax principle becomes immediately far more convincing. The behavior of their "rational" player may well be described as that of a man whose sole concern is to come out with as little loss as possible. The minimax strategy, to him, is the least

²⁴ P. 164. In other words, if B should use a nonminimax strategy, A generally could not enforce the maximum element in that column by choosing his own maximin strategy.

²⁵ P. 164.

ominous choice in a game he would rather not play. His is not the attitude, to be sure, of one playing a game for entertainment or profit. It is, in fact, the psychology of a timid man pressed into a duel.

This is not to deny that cautious pessimists do exist or that a defensive policy is often desirable. A theory of reluctant duelists is not a small achievement. But it could not be reliable in predicting behavior in situations corresponding to the zero-sum two-person game; nor is it plausible that players should be advised to conform to it against their inclinations. It is certainly not a theory of games. It is not a theory of rational behavior under game-uncertainty; that theory lies in the future. If it comes, I believe it will show an immense debt to the insights and theoretical framework provided by von Neumann and Morgenstern. But it will not come the faster for a misbelief that its place was filled a dozen years ago.