

Discussion Paper No. 533

A NON-METRIC TEST OF THE MINIMAX THEORY

by

Barry O'Neill*

June 1982

Department of Industrial Engineering and Management Sciences
Northwestern University
Evanston, Illinois 60201

The minimax theory for two-person zerosum games was tested using a matrix that was especially easy for the subjects to understand, and whose solution did not depend on quantitative assumptions about their utility functions for money. Players' average relative frequencies for the moves and their proportion of wins were almost exactly as predicted by minimax, but subject-to-subject variability was too high. This suggests that people may deviate from minimax theory, since their opponents have limited information and are imperfect record-keepers, but they do not deviate so much that their own payoffs are diminished.

Von Neumann and Morgenstern's minimax solution (1944) is generally accepted as the correct way to play two-person zerosum games. "Correct" is meant here in the prescriptive sense of what players should do if they are rational.

Does the theory describe the behaviour of real players? A number of experiments have been reported but the results have been equivocal or negative (see Rapoport, Guyer and Gordon, 1977, Chs. 21 and 24, and Section 7 of the present paper). This question has wide importance since many models in the social sciences, particularly economics, are based on minimax theory or its generalization, the theory of Nash equilibria (Nash, 1950).

A problem in past research has been finding an experimental design that accurately tests the theory. Here we describe a procedure that avoids two previous difficulties.

The present experiment involves a special game matrix, chosen for two reasons. First, it tests the theory without making assumptions about the exact shape of the players' utility functions for money. Second, it is easy for the subjects to comprehend. It is unique in being the simplest game possible, except for some that are trivial to solve. (The definition of simplicity will be given in section 3.)

Our subjects' behaviour was close to minimax, and we will argue that this confirmatory evidence should weigh strongly against past failings of the theory, since the design used here is more appropriate.

In Sections 2 and 3 the research problems posed by utility measurement and the complexity of experimental games are discussed. Sections 4 and 5 detail the procedure and results. Section 6 compares minimax with some stochastic theories previously proposed for zerosum games, and Section 7 discusses possible explanations of the data.

The strategies recommended by the minimax solution depend on the subjects' utilities for the money payoffs. These may be different from the payoffs themselves, so any empirical test must either determine or make assumptions about the players' utility functions. To our knowledge all past researchers have assumed explicitly or implicitly, that utility in zerosum games depends only on the player's own payoff, and is a linear function of that payoff. This seems counterintuitive since it rules out such motives as a desire to equalize payoffs or to beat the opponent's earnings. Also, we have seen no empirical evidence that utility is linear in money, even if the range of payoffs is restricted to small amounts.

An alternative approach would be to assess each player's utility function individually, and then design a game matrix with payoffs calculated to be zerosum in utilities. We know of no research that has done this, and very likely the reason is that utilities would be different in a game than in a single-person decision problem. When people interact, various competitive or altruistic motives arise, so their utilities depend on their own and their opponent's payoffs.

The experiment described here uses a different approach. A matrix is constructed with the property that a player's minimax strategy is invariant over all reasonable utility functions. The game is shown in Figure 1.

FIGURE 1 HERE

Both players should use the mixed strategy vector $(.4, .2, .2, .2)$, in which case player 1 will win 40% of the time and have an expected payoff of -2μ . This holds for any pair of utility functions $u_1(x_1, x_2)$, $u_2(x_1, x_2)$, where x_1 and x_2 are the money payoffs in the matrix. The only requirements are that each player would rather win an amount than

+5¢	-5¢	-5¢	-5¢
-5¢	-5¢	+5¢	+5¢
-5¢	+5¢	-5¢	+5¢
-5¢	+5¢	+5¢	-5¢




Figure 1: Representations of the experimental game as a matrix, bicolored graphs, and directed graphs.

lose that amount, and that utility depends only on the payoffs and not on the cell where those payoffs appear, an assumption implicit in our notation.

This invariance over utility functions occurs because there are only two types of outcomes, a win or a loss. The minimax solution is constant for positive linear transformations of the players' utilities. If their utility functions satisfy $u_i(5¢, -5¢) > u_i(-5¢, 5¢)$, then a positive linear transformation can always be performed to bring the players' utilities into coincidence with the money values. When the experimenter uses the payoffs to solve the game, the result will be the same as when the subjects use the utilities, provided that both follow minimax theory.

To clarify the structure of the game we can also represent it as a bicolored graph, as shown in Figure 1. Player 1 chooses a triangular shaped node and player 2 a circle. If the two nodes are adjacent on the graph, then player 1 wins, and otherwise player 2. There are generally two representations of a game as a bicolored graph: in one, adjacency represents a win by player 1, in the other, by player 2. Since the matrix of Figure 1 gives both players the same number of strategies it can also be represented by a directed graph in which a player can choose any node, and player 1 wins by being within one step of player 2.

There are games other than Figure 1 that show invariance over utility functions -- any game in which each player has exactly two payoff levels will do -- but the game used here is most appropriate because it is the smallest non-trivial one, as will be described in the next section.

3. The Problem of Simplicity

In our view, the games used in past experiments have been too complicated for subjects to understand. In some ways a 2x2 game is

simpler than the situations they deal with in daily life, but a laboratory game confronts them with an unfamiliar situation. Many have difficulty processing numbers, especially when they must look at these numbers from both their own and their opponent's viewpoints.

One way we can achieve simplicity is to restrict the payoffs to two levels, as discussed above. The game becomes a purely structural entity, determined by the relationship of wins to losses, so the subjects are freed from calculations involving relative magnitudes.

After that, we can look for games with the smallest number of strategies. Dominated or duplicated strategies should not occur, since otherwise the game could be immediately reduced to a smaller one.

Of course the game should not be trivial to solve. An example of one that is too simple for our purposes is the children's pastime Scissors, Paper and Stone (Opie and Opie, 1970). A choice of "Scissors" beats an opponent's choice of "Paper", "Paper" beats "Stone", and "Stone" beats "Scissors". This is symmetrical in strategies, and therefore the mixed strategy solution is clearly $(1/3, 1/3, 1/3)$. We want a game that cannot be solved by symmetry, so that players will be forced to use the full logic of the minimax solution.

These ideas are incorporated in the following criteria:

(C1) The game is in normal form.

(C2) There are exactly two levels of payoff for each player.

(C3) No player has any two strategies identical.

(C4) No player has a dominated strategy.

(C5) The game is not completely symmetrical in strategies.

(C6) The game has the smallest number of strategies per player, consistent with (C1)-(C5).

It is surprising that there is a unique game satisfying (C1)-(C6),

assuming of course that games equivalent by positive linear utility transformations, or permutations of the players or strategies, are regarded as identical. It is the game of Figure 1. Some larger games satisfying appear in the Appendix, along with the method for generating these games.

4. Procedure

The game was played by fifty students working in twenty-five pairs. They were recruited from the Northwestern University student body by posted advertisements and personal contacts. Each served in only one session, and students who knew each other were not allowed to participate in the same pair. The sessions lasted about half an hour and subjects received their winnings as payment.

The players sat at a table opposite each other. Each held four cards, joker, ace, two and three. A large board across the table prevented them from seeing the backs of their opponent's cards. They were read the following instructions.

"We are interested in how people play a simple game. I will teach you the rules of the game, then have you play about 15 hands to make sure you are clear about the rules. Then you will play a series of hands for money at 5c per hand.

"The rules are as follows:

1. Each player has four cards, ace, two, three, and joker.
2. Each player will start with \$2.50 in nickels for the series of hands.
3. When I say 'ready' each of you will select a card from your hand and place it face down on the table. When I say 'turn', turn your cards face up and determine the winner. (I will be recording the cards as played.)
4. The winner should announce, 'I win', and collect 5c from the other player.
5. Then return the card to your hand.

"Are there any questions?

"Now to determine the winner:" [Subjects were shown a placard giving these rules :]

[S's name] wins if there is a match of jokers (two jokers played)

or mismatch of number-cards (2, 3, for example)

[S's name] wins if there is a match of number-cards (3, 3, for example)

or mismatch of a joker (1 joker, 1 number card)

Thus the game was presented in English and the rules were learned by practice, without a matrix, or graph.

The game was first practiced 15 times, then played 105 times for real money. Subjects went through at their own speed.

Based on some pretrials, the device of having the players themselves figure out who won seemed to be useful in that it increased their involvement in the game. Players tended to focus their attention on each other rather than on the experimenter. If they happened to make an error in determining the winner, they were corrected by the experimenter.

In a post-session questionnaire, all subjects answered that they had understood the rules of the game well.

5. Results

The number-cards (ace, two and three) are strategically equivalent to each other, so each should be used with equal probability. Therefore we will usually group these three moves in the analysis and look only at the relative proportions of jokers versus number-cards.

The number of jokers and number of wins for each subject are listed in Table 1. There were 5250 total moves made in the experiment (50 subjects x 105 plays each). The proportion of jokers was .394, compared to minimax theory's prediction of .400. Looking at the two types of players separately, those in the role of player 1 chose proportion .362 jokers, and those in the role of player 2 chose .426, compared to predictions of .400 for both by the minimax theory.

TABLE 1 HIPE

These three values are close to the predictions of the theory and the discrepancies are not statistically significant. Using a t-test for comparing a sample distribution of unknown variance to a mean of .4, the p values, two-tailed, are .412, .051 and .231 respectively.

Subjects showed no significant tendency to move toward or away from the minimax solution during the course of the sessions. Separating the first 52 moves from the last 53, 30 subjects were closer to the minimax proportion of jokers in the first part, and the other 20 were closer in the second part. This is not significantly different from the expected numbers of 25 and 25.

Although the mean number of jokers fitted minimax's prediction, some finer details differed from the theory. First of all, the variance in the number of jokers from subject to subject was larger than would be expected, as shown in Table 2. This was true for each role individually and for player-types combined. According to minimax, the number of jokers produced by a player will follow a binomial distribution with mean of $105 \times .4 = 42$. This distribution is approximately normal so we may use a χ^2 -test for a difference between observed and predicted variance. For subjects in the role of player 1, player 2, and combined, the values of χ^2 are 91.7 (df=24), 119.8 (df=24) and 234.5 (df=49), all of which are significant at least at the .001 level using a two-tailed test.

TABLE 2 HERE

Thus the subject-to-subject variation in the frequency of jokers is too large for a binomial distribution. It seems as though a subject's probability p of a joker is not fixed at .400, but changes over the individuals, centered approximately on .400. It is useful to have an estimate of this variability separate from that introduced by sampling only a finite number of moves of each player. We will regard each subject's probability of a joker as a random variable following a

distribution from the Beta family. This type has shapes that are plausible for our data, and are convenient to use in calculations. It is assumed that each subject samples a value of p and uses it throughout the game. Given that the mean of p is .400, the parameters of the Beta can then be written r and $3r/2$. Applying maximum likelihood methods, using the 50 observed frequencies of jokers, gives $r = 11.5$. The upper and lower quartiles of such a Beta distribution are $p=.340$ and $.460$, so 50% of the subjects are estimated to have p 's in this interval, rather than all having $p=.4$ as minimax states.

Another difference from the predictions of minimax is found in the numbers of runs, i.e., unbroken series of jokers or number-cards. Many people produced unusually many runs, meaning that they had a tendency to switch back and forth more quickly than if their choices were independent of past moves. The number of runs arising from a long random sequence is approximately normal, so for each player, the expectation and variance of that player's number of runs, given the observed proportion of jokers used, was calculated, and the observed number of runs was converted to a z-score. The mean was .843, compared with a null hypothesis of 0, indicating significantly more runs than chance ($t=3.02$, $df=49$, $p<.001$).

Looking at the frequencies of the three types of number-cards, we can judge the prediction that they were produced with equal likelihood. Subjects in the role of player 1 produced 578 aces, 565 twos and 532 threes, and those in the role of player 2 produced 593 aces, 470 twos and 446 threes, the latter being significantly different from equiprobability ($\chi^2=24.7$, $df=2$, $p<.001$). The best explanation we can offer is that the players were attracted by the positive connotations of an ace. In hindsight the use of an ace may have been a fault in the experimental design.

The proportions of wins by each player were strikingly close to the

predicted value. Proportions observed were (.401, .599) compared to minimax's (.400, .600) for players 1 and 2 respectively.

A finding of interest is that there was no evidence of differential skill in playing the game. Skill means that some subjects in the first role would do significantly better than 40% wins, and others would do significantly worse. This would show up as an increase in the variance of the number of wins of those in the role of player 1, compared to the variance of a binomial distribution with $n=105$ and $p=.4$. The sample standard deviation of the number of wins was 6.7, not significantly different from 5.02, that of a binomial ($\chi^2=31.8$, $df=24$, $p > .1$). This means that one of the players won or lost about 1.7 games further away from chance than would be expected.

If a player uses a certain move and wins, is that move more likely than average to be repeated on the next game? For our subjects it turns out that it is somewhat less likely. To determine this let a and b be the probabilities that players 1 and 2 respectively use a joker. Then the probabilities that a move that wins will be followed by a repetition of that move are $a+b-2ab$ and $1-a-b+2ab$ for the two players, given that successive moves are independent. The parameters a and b were estimated by the observed relative frequencies of jokers used, and for each player these numbers were compared with the observed number of repetitions of winning moves. Of the 50 players, 18 repeated more often than expected and 32 repeated less often, a difference that is statistically significant ($p < .05$). A large majority, 19 out of 25, of subjects in the role of player 2 tended to avoid repeating a move after a win.

Our players seemed to feel that a move that has just succeeded should be avoided in the immediate future. This is related to the single-person decision phenomenon of the Gambler's Fallacy, and negative recency in probability learning research (Estes, 1964; Jones, 1971). In a probability learning experiment, subjects must guess which of two

outcomes of a random process will occur. It is observed that they tend to switch their guesses following a success especially after a run of one type of outcome.

The negative recency behaviour of our subjects will be important in evaluating mathematical learning theory models for games discussed in the next section.

6. Learning Theory Models for Games

A natural theoretical system to apply to game-playing is mathematical learning theory. It was developed 1950's and 1960's to predict single-person learning situations, but can be extended naturally to two-person interactions. "Learning" is meant here in a very simple sense. A subject must choose from several responses, one of which will be rewarded and the others not. This decision is made again and again in a series of trials, and the theory's aim is to predict how the subject's probability of making each choice varies as a function of past reinforcements. In probability learning, the subarea of learning theory most applicable to games with mixed-strategy solutions, reinforcement is given according to some random process, and the statistical properties of the subjects' responses are observed.

To extend the theory to two-person situations, we regard each player as learning which strategy to use trial and error, by reacting to past successes and failures. The two players' experiences interact, of course, because the success of a move depends on what the other's move is.

Mathematical learning models are antithetical to game theory models, because they ignore the concepts of thinking and planning. The subjects are not assumed to be trying to outwit each other, or even to keep track of what the other has done, but only to modify their probabilities of moves depending on whether they themselves have

benefitted from using a certain move. This represents the influence of behaviorism in psychology, with its rejection of unobservable cognitive processes. A number of learning theory models have been applied to two-person game situations. They have had some success, and this poses a very basic challenge not only to minimax theory, but the whole approach of postulating rationality as a theory of human behaviour.

In this section we will summarize some learning theory models that have been applied to games. We have tried to include all published applications, on the condition that they are probabilistic. Some models in the literature, for example, some offered by Messick (1967a,b), predicted determinate moves on each trial as a function of past experience. These were not included, since they seem too vulnerable to exploitation by an opponent.

Only a summary of the assumptions of the models will be given here. For more information the review by Sternberg (1963) may be consulted.

A fundamental issue arises when we extend learning models to two-person interactions. Note that in game situations more than one response may be rewarded. For example, if player 1 chooses an ace, either a joker or an ace is a win for player 2, and it is not clear which move to regard as reinforced for player 2. Ideally we could postulate that both are reinforced, but no existing mathematical learning model seems to allow for the probability of more than one response to increase after a single play. One way of avoiding multiple reinforcements on a single trial would be to have an experimental game with only two moves per player. One or the other would be labelled a win, but not both. However this leads to a 2x2 game that is either completely symmetrical and therefore trivial, or to one with more than two levels of reinforcement for a single player. No stochastic theory of learning yet proposed seems to deal with different strengths of reinforcement. (Estes (1960) discussed this problem in the context of

two-person interactions, and proposed a non-stochastic model which was later tested by Messick (1965).)

We will take the same approach as past researchers, that a player is reinforced for a type of move, joker or number-card, if he or she uses that type and wins, or uses the other type and loses. If a player wins but would have won as well by using the other type of move, the latter is not regarded as reinforced. We suggest that only the move used to win is salient in the player's mind.

The Linear Operator Model

The linear operator model was introduced in the context of games by Atkinson and Suppes (1958).

Suppose two responses, R_1 and R_2 , are possible, and the subject has probability p_n of making response R_1 at trial n . If the experimenter reinforces response R_1 , then the linear operator model postulates that the probability p_{n+1} of response R_1 at the next trial is a linear function of p_n :

$$p_{n+1} = p_n + \theta(1-p_n) \text{ with } \theta \in (0,1].$$

If the experimenter had reinforced the other choice R_2 , a different operator would have been applied to p_n , an operator whose form can be deduced from the definition of the first operator, if we assume that the two responses can be learned equally quickly. If R_1 is reinforced then R_2 's probability must decrease from $q_n (= 1-p_n)$ to θq_n , since R_1 and R_2 are complementary events. Thus by, symmetry, if R_1 is not reinforced, its own revised probability must be

$$p_{n+1} = \theta p_n$$

The parameter θ measures the subject's responsiveness to past experience. For simplicity we will assume that the two players learn at the same rate, i.e., have identical values of θ although our conclusions about the adequacy of the model for our results can be shown

to be the same without this assumption.

Let the probabilities on trial n of players 1 and 2 choosing a joker be a_n and b_n respectively, and let w_n be the probability that they both choose a joker. Then the probability a_{n+1} for player 1 on the subsequent move can be calculated:

$$\begin{aligned} \text{Prob}(1 \text{ uses } J \text{ at } n+1) &= \text{Prob}(1 \text{ uses } J \text{ at } n+1 / JJ \text{ at } n) \times \text{Prob}(JJ \text{ at } n) \\ &+ \text{Prob}(1 \text{ uses } J \text{ at } n+1 / JN \text{ at } n) \times \text{Prob}(JN \text{ at } n), \text{ etc.} \end{aligned}$$

The conditional probabilities can be determined by noting who wins at play n , which move they made, and therefore whether it was reinforced. These probabilities are shown in Table 3.

TABLE 3 HERE

The probabilities of the pair of moves made at play n , JJ, JN, NJ and NN are w_n , $a_n - w_n$, $b_n - w_n$ and $1 - a_n - b_n + w_n$, respectively. It can be calculated that the asymptotes a and b , of a_n and b_n are $a = \frac{1}{2}$ and $b \in [.375, .400)$, where the value approached by b_n depends on the learning parameter θ . The set of possible values of a and b are compared with the observed values in Figure 2.

FIGURE 2 HERE

They are far from the observed values compared with minimax's predictions.

Atkinson and Suppes (1958) suggested the use of Luce's Beta model to deal with two-person games. Like the linear operator model, probabilities are modified depending on which move is reinforced, but the equations are different. In the case of reinforcement in a single-person decision situation, Luce (1959) suggested the following operators for the cases of reinforcement and nonreinforcement, respectively:

$$P_{n+1} = \beta P_n / (1 - P_n + \beta P_n)$$

$$P_{n+1} = P_n / (P_n + \beta - \beta P_n)$$

The parameter β measures the subject's responsiveness to past experience.

This results in an "independence of irrelevant alternatives" property: if a player decides to choose a response from some specified subset, then the relative likelihood of each response in the subset is independent of the unconditional probabilities of responses outside the subset.

Applying these operators to the two-person case by setting up a table analogous to Table 3, gives a set of recursive equations for the probability of responses at trial n . The equations for the asymptotic values are too difficult to solve analytically, and therefore were approximated for various values of β by iteration on a computer. The set of possible values is shown in Figure 2.

Again, all possible values are farther from the data than is the minimax prediction.

Suppes and Atkinson's Markov model

In the same way that linear operator theory is related to behaviour modification, Suppes and Atkinson's model (1960), also known as single-stimulus sampling theory, is an outgrowth of the study of

classical conditioning. Each player is regarded as being ready to make one or the other move at each play. If that move wins, the move is repeated, but if it loses, the player's disposition changes with probability θ to the alternative. Applying their equations (pp.26-27), the asymptotes for our game are (.500, .400).

Like the other two learning theory models, the Markov model performs poorly on our data, as shown in Figure 2. This is in contrast to the positive results of Suppes and Atkinson's series of experiments. Possible reasons are that in most of their experimental procedures subjects did not know they were facing each other, but thought they were interacting with a chance reinforcer. They were often not told the payoffs associated with the outcomes, or the opponent's moves, and further uncertainty was introduced into the linkage of the move with the payoff, by having the payoffs depend only probabilistically on the outcomes. Although these features are common in daily life, they are far from the assumptions of game theory, and it is understandable that a model based in learning theory gave better predictions. In their one experiment whose design involved fulling informing the subjects about the payoffs, and the existence of an opponent (Group E), learning theory predictions were not close to the observed proportions.

A probability matching model

A phenomenon reported in probability learning experiments is probability matching: subjects tend to guess an outcome with a probability that is identical to its likelihood of being reinforced (Estes, 1964). Although this is irrational from a utility-maximizing perspective, the fact that it has been widely observed, and the fact that another feature of our data, negative recency, also occurs in probability learning experiments, prompt us to test a model suggested by Estes (1958), based on probability matching. He conducted an

experimental game and found remarkably close agreement. Matching becomes more complicated in the case of two people, since each is trying to duplicate probabilities influenced by the other, not fixed externally. To determine the predictions of the theory for our game, the fourth and fifth columns of Table 3 were used to determine the probabilities of reinforcement of each move as a function of a and b. The basic idea of probability matching states that these are equal to the probabilities of using the moves:

$$a = ab + (1-a)b + (1-a)(1-b)/3$$

$$b = (1-a)b + 2(1-a)(1-b)/3$$

The solution is (.500, .400), identical to Suppes and Atkinson's model, and relatively far from the data. One possible reason that this model worked in Estes' experiment but not here, is that in the former only one pair of subjects was used, and they had previously participated in a series of single-person learning experiments. Thus they may have been unintentionally "primed" to produce probability matching.

Assessment of the stochastic models

Overall, we can say that the learning theory models do not fit the subjects' behaviour. First, there is the discrepancy in the predicted proportions of moves. In the plane of the players' probabilities, the linear operator model ranges from .140 to .147 in Euclidean distance from the data, Luce's model from .062 to .147 and the probability matching and Markov models lie .140 away. This compares with the minimax's distance of .046.

There is a further problem in the learning theories' predictions about the sequence of moves. They imply that when a response is reinforced its probability rises, but this was not observed in the data. Instead we found negative recency, which is tantamount to




Figure 2: Observed probabilities of a joker compared with minimax and learning theory predictions.

"anti-learning". Subjects avoid successful moves, at least in the short term. Of course there is a plausible explanation — they are trying to double-think each other, to avoid moves they think the opponent is expecting. But this is directly counter to the assumptions of learning theory, and in this way these models are incorrect in their first premises. (For a model that allows negative recency, see Restle (1961)).

7. Discussion

Game theory deals with social interaction, not with individual responses. Therefore the most accurate test of the theory involves situations where subjects are facing other subjects, rather than programmed strategies. Many researchers, however, have had subjects play against computers or stooges. While this may be useful in increasing experimental control, it is not clear that natural *game-playing* behaviour is being elicited. Unless the experimenter knows how an opponent behaves, it is impossible to program a computer to act like one.

A further disadvantage of this approach is that subjects are usually isolated in cubicles. They do not see an opponent, of course, because there is none, and feel less involved in a competition.

Subjects should know that they are interacting with an opponent and know the payoffs involved. In several experiments, subjects have been told they are acting in a random environment, when they were actually playing against an opponent. We do not believe it is valid to use the results as a test of *game* theory.

Past experiments were surveyed according to these guidelines, and the results are displayed in Table 4. We have tried to include all *game* experiments that used pairs of real subjects informed of their situation. Further criteria are that the subjects chose moves directly,

e.g., did not choose mixed strategies which were later implemented by the researcher, and also that the number of moves was "small". The last condition eliminates some experiments involving duels and Colonel Blotto games.

TABLE 4 HERE

Past results have been counter to the minimax theory, but we regard the present experiment as evidence that it may be true after all. Here the game was simple, the subjects met face-to-face in an involving situation, and the test was free of metric assumptions about utilities.

Positive evidence was the correct proportions of strategies used, and the correct proportions of wins. Negative evidence was the correct dependence among successive moves, and the high variance in proportion of jokers from subject to subject. It seems that minimax theory was confirmed in the large, but not in the small.

This is puzzling at first: how could the overall proportions have followed the theory, when the individual moves that generated them did not? We suggest that the following occurred: players were constrained to follow minimax in its gross statistics because these were more available to the opponent. However, at each move players felt free to invent patterns, follow hunches, or do a number of other things that introduced dependencies and variance into the sequence of plays. They could do this without significant danger because the opponent had a limited ability to calculate all the relevant probabilities, especially when only a small sample of moves was available. But a large deviation from the overall minimax proportion was easier to notice, so players avoided the risk of loss by sticking close to the minimax predictions.

This is more plausible when we consider the well-known feature of the minimax solution, that if one player follows it, the opponent is free to deviate without loss. Thus the closer one gets to the minimax

the less incentive there is for the other to follow it too. We can then expect a certain amount of variation, but it would be centered on the minimax probability.

If this explanation is true we would not expect each player's proportion of wins to be affected. Neither would decide to go so far from minimax that their proportion of wins decreased. This is precisely what was found: player 1 won almost exactly 40% of the time.

Appendix

Table 5 gives the number of game matrices with only two levels of payoffs. These structures satisfy (C1) and (C2), but possibly violate (C3)-(C6), in that they may have dominated or duplicated strategies or be symmetrical in moves.

TABLE 5 HERE

The numbers were determined by applying Harary's formula for enumerating undirected bicolored graphs (Harary, 1958). Two-level games and bicolored graphs are isomorphic, as can be seen by identifying strategies with nodes, and a win for player 1 with adjacency of the two strategies used, as in Figure 1. Player 1 needs to be specified to define the isomorphism, and will be chosen as the player with fewer moves, or in the case of equal numbers of moves, the player with the fewer winning strategies.

For certain values the isomorphism fails: if the two players have the same number of strategies and same number of wins, there can be two graphs yielding two games that are equivalent to each other but with the players switched. In this case Harary's formula gives an upper bound. The only such inexact value in the table is 55, an upper bound for the number of 4x4 games with eight wins for a player. Case-by-case listing shows that there are only 29 distinct games in that category.

To find the games satisfying the additional requirements of (C3)-(C5), we can proceed as follows. First, list all possible rows such a matrix could have. There are 2^n of these: RRR...R, CRR...R, ..., CCC...C. Construct a graph with nodes representing these rows and an edge connecting two nodes if one row dominates the other. Identify all the independent sets (sets of nodes with no edge connecting any pair) of up to n nodes. Each independent set will correspond to a game without dominated or duplicated rows. There may be dominated or duplicated

columns, so they must be inspected pairwise and the game eliminated if these appear.

It is helpful to use certain shortcuts, e.g., to divide the search into parts, looking first for games with exactly one row containing exactly one win for the row chooser. This row can be specified to be RCC...C, and the other rows, CRC...C, ..., CCC...R, eliminated from the graph. Next we look at rows with exactly two wins for player 1. This substantially reduces each graph's size. The method is tedious but can be performed by hand for games as large as 5x5. For larger games it can be computerized and independent sets of strategies determined by an algorithm such as that given by Bron and Kersboch (1973).

The results of applying this procedure for $m = 1...5$, and $n=1...5$ are shown in Figure 3. Along with each game is a graphical way of representing it, and also its solution by minimax theory. The list includes 1 4x4 , 1 4x5 and 7 5x5 games.

FIGURE 3 HERE


Two 5x5 games are of special interest in that they can be represented very simply as undirected graphs, as shown in Figure 3.

Similar lists of two-person nonzerosum games, and three-person 2x2x2 games are given by O'Neill (1976) and O'Neill (1982).

4x4 Game

	$\frac{2}{5}$	$\frac{1}{5}$	$\frac{1}{5}$	$\frac{1}{5}$
...	+	-	-	-
2/5	-	-	+	+
1/5	-	+	-	+
1/5	-	+	+	-


Value = $-1/5$



4x5 Game

	0	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$
1/4	-	-	-	+	+
1/4	-	+	+	-	-
1/4	+	-	+	-	+
1/4	+	+	-	+	-


Value = 0



5x5 Games


	$\frac{3}{7}$	$\frac{1}{7}$	$\frac{1}{7}$	$\frac{1}{7}$	$\frac{1}{7}$
3/7	+	-	-	-	-
1/7	-	-	+	+	+
1/7	-	+	-	+	+
1/7	-	+	+	-	+
1/7	-	+	+	+	-

Value = $-1/7$



	$\frac{1}{9}$	$\frac{1}{9}$	$\frac{2}{9}$	$\frac{2}{9}$	$\frac{3}{9}$
1/9	+	+	+	-	-
1/9	+	+	-	+	-
2/9	+	-	-	-	+
2/9	-	+	-	-	+
3/9	-	-	+	+	+

Value = $-1/9$



	$\frac{1}{7}$	$\frac{1}{7}$	$\frac{1}{7}$	$\frac{2}{7}$	$\frac{2}{7}$
1/7	+	+	+	-	-
1/7	+	-	-	+	-
1/7	+	-	-	-	+
2/7	-	+	-	-	+
2/7	-	-	+	+	-

Value = $-1/7$








Figure 3, cont., over

5x5 Games (cont.)




	$\frac{1}{2}$	0	$\frac{1}{2}$	0	0
1/2	-	-	+	+	+
0	-	+	+	-	-
1/2	+	+	-	-	-
0	+	-	-	-	+
0	+	-	-	+	-
	Value = 0				



	$\frac{3}{8}$	$\frac{1}{8}$	$\frac{1}{8}$	$\frac{1}{8}$	$\frac{2}{8}$
3/8	+	-	-	-	-
1/8	-	+	-	-	+
1/8	-	-	+	-	+
1/8	-	-	-	+	+
2/8	-	+	+	+	-
	Value = -1/4				



	$\frac{1}{7}$	$\frac{1}{7}$	$\frac{1}{7}$	$\frac{2}{7}$	$\frac{2}{7}$
1/7	-	+	+	-	-
1/7	+	-	+	-	-
1/7	+	+	-	-	-
2/7	-	-	-	+	-
2/7	-	-	-	-	+
	Value = -3/7				



	$\frac{1}{4}$	$\frac{1}{4}$	0	$\frac{1}{4}$	$\frac{1}{4}$
1/4	+	+	+	-	-
1/4	+	-	-	+	-
0	+	-	-	-	+
1/4	-	+	-	-	+
1/4	-	-	+	+	+
	Value = 0				





Figure 3: Games satisfying (C1)-(C5) for up to five strategies per player. "+" represents a win for player 1, "-" for player 2. Minimax strategies appear by the rows and columns. The value given is the value to player 1, for win = +1, loss = -1. Graphical representations are as described in Section 2.

References

- Atkinson, R. and Suppes, P. An analysis of two-person game situations in terms of statistical learning theory. Journal of Experimental Psychology 55. 1958, 369-378.
- Bron, C., and Kerbosch, J. Algorithm 457 — finding all cliques of an undirected graph. Communications of the A.C.M. 16, 1973, 575.
- Estes, W. Of models and men. American Psychologist, 12. 1957, 609-617.
- Estes, W. Theoretical treatments of differential reward in multiple choice situations. in J. Criswell, H. Solomon and P. Suppes, eds., Mathematical Methods in Small Group Processes. Stanford: Stanford University Press. 1962
- Estes, W. Probability learning, in A.W. Melton, ed., Categories of Human Learning. New York: Academic Press. 1964.
- Harary, F. On the number of bicolored graphs, Pacific Journal of Mathematics. 8. 1958. 743-755.
- Jones, M.R. From probability learning to sequential processing: a critical review. Psychological Bulletin. 76, 1971, 153-158.
- Malcolm, D. and Lieberman, B. The behavior of responsive individuals playing a two-person zerosum game requiring the use of mixed strategies. Psychonomic Science 2. 1965, 373-374.
- Messick, D. Interdependent decision strategies in zerosum games. Dissertation Abstracts. 27, 4147, 1967.
- Nash, J. Equilibrium points in N-person games. Proceedings of the National Academy of Sciences 36, 48-49, 1950.
- O'Neill, B. Variable-threat bargaining: a test of the Raiffa-Nash theory. Mathematical Psychology Program Preprint. University of Michigan, 1976.
- O'Neill, B. The simplest three-person games. Center for Mathematical Studies in Economics and Management Science discussion paper. Northwestern University, 1982.
- Opie, I. and Opie, P. Children's Games in Street and Playground. Oxford: Clarendon Press. 1969.
- Rapoport, A, Guyer, M. and Gordon, D. The 2x2 Game. Ann Arbor: University of Michigan, 1976.
- Restle, F. The Psychology of Judgement and Choice. New York: Wiley, 1961.
- Sakaguchi, M. Reports on experimental games. Japan Union of Science and Engineering, 7, 1960, 156-165.
- Sternberg, S. Stochastic learning theory, in R. Luce, R. Bush and E. Galanter, eds., Handbook of Mathematical Psychology, Vol.II, 1963, 1-120.

Suppes, P. and Atkinson, R. Applications of a Markov Model to Multiperson Interactions. Stanford: Stanford University Press, 1960

Von Neumann, J., and Morgenstern, O. Theory of Games and Economic Behavior. New York: Wiley and Sons, 3rd edition, 1964.

(number of jokers used by player 1,
number of jokers used by player 2,
number of wins by player 1)

Range: (0-105, 0-105, 0-105)
Minimax prediction: (42,42,42)

Observed values:

(19,37,41)	(46,58,31)	(57,58,41)	(35,76,44)	(49,47,36)
(41,47,44)	(32,37,50)	(34,31,49)	(31,36,38)	(44,43,40)
(32,39,39)	(51,45,57)	(28,56,43)	(32,34,49)	(48,39,34)
(46,40,35)	(38,43,35)	(41,43,51)	(34,53,41)	(45,52,46)
(39,34,50)	(48,36,42)	(17,44,47)	(27,39,52)	(35,62,35)

	Mean prop'n of jokers		Wins	
	Obs'd	Pred'd	Obs'd	Pred'd
Subjects, Pl'r 1	.362	.400	.401	.400
Subjects, Pl'r 2	.426	.400		
Combined	.394	.400		

Table 1. Numbers of jokers for each subject-pair and means of proportions.

No. of Jokers	No. of Subjects		Difference
	Observed	Predicted	
0-4	0	.00	.00
5-9	0	.00	.00
10-14	0	.00	.00
15-19	2	.00	+2.00
20-24	1	.01	+1.99
25-29	2	.27	+1.73
30-34	8	3.09	+4.91
35-39	11	12.24	-1.24
40-44	8	19.07	-11.07
45-49	9	11.96	-2.96
50-54	3	3.08	-.08
55-59	4	.32	+3.68
60-64	1	.01	+.99
65-69	0	.00	.00
70-74	0	.00	.00
75-79	1	.00	+1.00
80-84	0	.00	.00
85-89	0	.00	.00
90-94	0	.00	.00
95-99	0	.00	.00
100-105	0	.00	.00
Stand. Dev.	11.00	5.02	5.98

Table 2. Number of subjects per number of jokers used.

Move at n	Winner	Move reinf'd		Prob(Joker at n+1, given moves at last play)		
		Player:	Player:	Player:		
1	2	1	2	1	2	
J	J	1	J	N	$(1-\theta)a_n + \theta$	$(1-\theta)b_n$
J	N	2	N	N	$(1-\theta)a_n$	$(1-\theta)b_n$
N	J	2	J	J	$(1-\theta)a_n + \theta$	$(1-\theta)b_n + \theta$
N	N	*1	N	J	$(1-\theta)a_n$	$(1-\theta)b_n + \theta$
		2	J	N	$(1-\theta)a_n + \theta$	$(1-\theta)b_n$

Table 3. Operators for the linear operator model. "J" — Joker

"J" — Joker, "N" — Number-card.

* — first row occurs with probability 2/3, and second occurs with probability 1/3.

Source	No. of pairs	No. of plays	Minimax strat'y	Prop'n of moves: Average	Sample S.D.
Estes (1957, Figure 7)	1	50	.667 .667	.50 .59	-- --
Salzguchi (1960)	1	60	.433, .367, .150 .333, .533, .167	.500, .333, .167 .500, .333, .167	-- --
Suppes and Atkinson (1960, group E)	48	210	.375 .875	.691 .684	.113 .149
Malcolm and Lieberman (1965)	9	200	.750 .750	.694 .571	.074 .093
Frenkel, (Rapoport et al., 1976, game #75)	96	1	.909* .909	.86 .50*	.39 .49
present study	25	105	.400 .400	.362 .426	.093 .107

Table 4. Results of game experiments. (* -- these numbers were reversed in the original publication due to a typographical error.)

	Size of Game						
	2x2	2x3	2x4	3x3	3x4	4x4	
0	1	1	1	1	1	1	
1	1	1	1	1	1	1	
Wins for Player 1	2	3	3	3	3	3	
	3	1	3	3	6	6	
	4	1	3	6	7	11	16
	5		1	3		13	21
	6		1	3		17	39
	7			1		13	44
	8			1		11	55* (29)
	9					6	
	10					3	
	11					1	
	12					1	

Table 5. Number of games satisfying (C1) and (C2) (two-level games in normal form). * - upper bound.