Equivalence and Stooge Strategies in Zero-Sum Games

JOHN FOX
MELVIN GUYER
Mental Health Research Institute
University of Michigan

Classes of two-person zero-sum games termed "equivalent games" are defined. These are games with identical *value* and identical optimal mixed-strategies but with different matrix entries and thus different opportunities for exploiting a nonrational opponent. An experiment was conducted to investigate the strategy-choice behavior of subjects playing pairs of these "equivalent games." Also investigated was the extent to which subjects would exploit a programmed stooge as a function of the degree to which the stooge departed from his optimal strategy mix. The results indicated that subjects learned to exploit the nonrational play of the stooge opponent. The game factor, on the other hand, seemed to have no significant effect upon the strategy-choice behavior of the players. The implications of these results are discussed in light of questions raised by previous research on decision-making in 2 x 2 zero-sum games.

In this paper we intend to explore the manner in which decision-making in two person zero-sum games is affected by game characteristics other than those formal properties of games which are usually assumed to influence choice behavior. The question dealt with here involves "equivalent" games and the extent to which the behavior of naive subjects playing against programmed stooges is similar or different in pairs of such games. In particular, we wish to determine the ability of human decision makers to learn to exploit their opponents' departures from the rationally prescribed strategy-choice in zero-sum games without saddle points. This ability will be examined in relationship to two factors: one, the extent to which such exploitation produces payoffs greater than the value of the game; and two, in terms of the magnitude of an opponent's departure

AUTHORS' NOTE: The experimental work on which this paper is based was partially supported by National Institute of Health Grant 035940, and by the Mental Health Research Institute, University of Michigan. The authors would like to thank Professor Thomas Storer of the Department of Mathematics, University of Michigan, for his advice in the preparation of this paper.

Journal of Conflict Resolution, Vol. 17 No. 3, September 1973,

from his optimal strategy choice mix. (Row's optimal strategy choice in a zero-sum game without a saddle point is defined as the probability mixture of pure strategies which gives him his largest guaranteed expected return, given that he has no information concerning his opponent's strategy choices. This guaranteed expected return to Row is called the value of the game. The optimal mixed strategy for Column is similarly defined.)

Equivalent games, which are used in the present paper, are games having identical value and identical optimal mixed strategies but which have different entries in their payoff matrices. Such games allow one to vary the degree to which a player is rewarded for exploiting a nonrational opponent while keeping value and optimal strategy mix constant.

By using equivalent games and programmed stooge strategies, the effects of magnitude of reward for exploitation of the other and magnitude of the opponent's departure from his optimal strategy can be systematically studied.

Studies dealing with these effects are not reported in the literature on experimental gaming. Instead, previous studies on decision-making in two-person zero-sum games have addressed themselves to the questions of whether subjects can learn to play the strategy containing the saddle point in games which have them, or learn to play their optimal mixed strategy in zero-sum games without saddle points.

Thus for example, in a study by Morin (1960), 28 male subjects played a variety of two-person zero-sum games with saddle points against a stooge opponent who always selected his optimal strategy. It was found that, while subjects tended to make fewer nonoptimal strategy choices over time, there was still a significant number of nonoptimal strategy choices made. Subjects also seemed unable to analyze the strategic characteristics of the games by considering the viewpoint of their opponent and so were not able to effectively find their own best strategy choice.

In a paper addressed to the same question, Lieberman (1960) studied choice behavior in a zero-sum game with saddle point. It was found that the percentage of optimal strategy choices increased rapidly over trials, with half the subjects selecting their optimal strategy 100% of the time on the last 75 trials of the play, while the remaining subjects showed a continuing increase in their choice of the optimal strategy.

Brayer (1964) studied strategy choice behavior in a 3 x 3 game with saddle point. The variables of interest included (a) the opponent's strategy, (b) the player's expectation of his opponent's strategy, and (c) the value of the game. The manipulation of the subject's expectation of the other's strategy was accomplished by varying the instructions given to the subjects. In addition to being given different instructions concerning the other's strategy, the players in fact played against either of two different programmed stooges, either "rational" or "random." It was found that, when playing against a rational opponent, subjects almost without exception played their own minimax strategy. In playing against the random opponent, subjects did not show any consistent pattern of strategy-choice selection; however, in games whose value was large, subjects tended to make more minimax selections than in games whose value was small.

This brief review of experimental studies of zero-sum games with saddle points suggests the general conclusion that subjects are quite successful at finding and persisting in the optimal strategy choice in such games. Departures from the optimal strategy choice are most easily attributed to a boredom effect associated with the large number of trials typically run in such experiments.

As in the case of zero-sum games with saddle points, the relatively few studies which have been conducted on zero-sum games *without* saddle points also deal with the degree to which players succeed in finding the optimal strategy choice or, in this case, the optimal strategy-choice mixture.

Kaufman and Becker (1961) used five different zero-sum payoff matrices, all having the same value, to study the strategy choice behavior of naive subjects in games without saddle points. The study explored the questions of whether subjects learn, by playing, to make optimal strategy choices, and whether games with the same value, but calling for different optimal strategy mixes, produce differences in the players' performances. Twenty subjects played, each against the experimenter, who, given the subject's strategy, selected a counter-strategy that kept the subject's earnings at or below the value of the game. In the event that the subject chose the optimal strategy mix, the experimenter responded with his own minimax counter-strategy. If the subject chose a nonoptimal strategy mix, then the experimenter departed from his own minimax in such a way as to exploit S's nonoptimal choice. The further S departed from his optimal strategy, the greater the loss he could expect to suffer. The results of this study, in terms of the average deviation of actual strategies from the optimal strategies, indicated that a majority of the subjects learned to find the optimal solution to the games when playing repeated trials, with feedback, against an opponent who exploited the subject's departures from the optimal strategy mix.

In a 1962 study, Lieberman examined strategy-choice behavior in zero-sum games without a saddle point. Each of 20 undergraduates played 300 trials of a 2 x 2 game against the experimenter. The minimax solution to the game required that the subject play his 2 strategies with a .75–.25 mix. The experimenter's optimal strategy was a .25–.75 mix. Against 10 of the subjects, E played his optimal strategy mix for the entire 300 trials of play. Against the other 10 subjects, E played optimally for the first 100 trials and then switched to a random .50–.50 mix for the remaining 200 trials.

The results of the study indicate that the subjects who played against the always optimal experimenter did not attain their own optimal mix but instead did not significantly depart from a .50–.50 strategy mix. The 10 subjects who played against the experimenter who was initially optimal, then subsequently nonoptimal, played as the first group did for the initial 100 trials—the period during which the experimenter was optimal. After the one-hundredth trial, however, this second group of subjects began to quickly approach its own optimal strategy mix of .75–.25. Though this group could have exploited the experimenter's nonoptimal strategy by playing the mix 1.00–.00, it was not observed to do so. This result is somewhat difficult to interpret: because of the manner in which the study was designed, to approach the optimal strategy mix is also to increasingly exploit the opponent's departures from his optimal play. Thus, against an optimal opponent, the first group of subjects could and did receive the value of the game by playing any strategy mix, optimal or not. Against a nonoptimal opponent, subjects' strategy choice mix approaches the minimax, thus assuring them of receiving the value of the game, but no more than that.

In a more recent study, Messick (1967) evaluated the performance of subjects playing a 3 x 3 zero-sum game without saddle point. The experiment was designed to assess the subjects' performance against that of a programmed opponent who used one of the three different strategies: (a) minimax strategy; (b) a maximization-of-expected-value model with a complete memory of the subjects' past choices; (c) a maximization-of-expected-value model with a limited memory of the subjects' past choices. Fourteen subjects were assigned to each of 3 different opponent conditions and each subject played 150 repeated trials.

The results for this study are reported in terms of money earned by the player. Against the minimax opponent, subjects' earnings were not significantly different from the value of the game. On the other hand, subjects earned much more than the value of the game when playing against either of the two expected value-maximizing strategies. And they did significantly better against the model with limited memory, compared to the one with complete memory.

A study by Fox (1972) examined the behavior of naive subjects playing against either of two programmed stooges in a 2 x 2 zero-sum game without saddle point. In the one case, the programmed "other" (a computer) played the optimal mixed strategy for the game. In the second case, the "other" played a nonoptimal strategy choice mix. For each stooge condition, subjects were assigned to either a "high" or "low" information feedback group. In the high feedback condition, subjects were told of their earnings and their actual strategy choice mix for the previous twenty trials of play. In the low feedback condition, subjects were informed only of their earnings for each block of twenty trials.

The results show that subjects playing against the rational opponent tended to play closer to their rational strategy as the game progressed, and the two feedback conditions produced no differential effect. Also, as the game progressed, the subjects paired with the nonoptimal "other" tended to select a strategy-choice mix that increasingly exploited the "other" as the game progressed. This latter finding offers support for the hypothesis that subjects are responsive to the strategy-choice behavior of their opponent when they can profit from such responsiveness (as when playing against a nonrational opponent).

This brief review of research findings suggests that the emphasis in previous research has been on the extent to which a subject achieves normative strategy-choice behavior given a particular payoff matrix, while little study has been made of the extent to which the behavior of the "other" and characteristics of the social situation (payoff matrix) affect a player's strategy choices in the context of one or another particular zero-sum game environment.

In the present study, attention is directed to two somewhat different aspects of choice behavior in a zero-sum game. Here, we are concerned with whether subjects can learn to exploit a nonrational opponent and whether such exploitation is affected by the extent to which exploitation is rewarded.

Before discussing our experimental findings, we describe the type of zero-sum games used in this study; these games will be referred to as "equivalent games."

Equivalent Zero-sum Games

A pair of two-person zero-sum games will have the same solution (i.e., the same set of minimax strategies) if the payoffs of one game are a linear transformation of those of the other. This result follows from the usual assumption that utility is measured on an interval scale.[1] If two games have the same set of minimax strategies, they may be termed "strategically equivalent." In the present paper, we shall concern ourselves with games which both are strategically equivalent and have the same value. Let us call such games "strategy-value equivalent," or more simply, "equivalent."

We begin our exploration of equivalent 2 x 2 zero-sum games by fixing the values of each of three parameters: (i) ROW chooser's and (ii) COLUMN chooser's minimax strategies—which we term (y, 1–y) and (x, 1–x), respectively—and (iii) the value of the game (z), and solving for the entries (a,b,c,d) in the payoff matrix. The generalized 2 x 2 zero-sum game to which we refer is shown below as Game 1.

		A: a_1	A: a_2
B	b_1	a	b
B	b_2	c	d

Game I

The minimax strategies of both players may be expressed in terms of the several entries of the game payoff matrix (Rapoport, 1966):

$$x = p(a_1) = \frac{d - b}{(a + d) - (b + c)} \quad [1]$$

$$y = p(b_1) = \frac{d - c}{(a + d) - (b + c)} \quad [2]$$

Equations 1 and 2 may be simply solved for a and c in terms of b, d, x, and y, yielding:

$$a = \frac{x + y - 1}{x} b \quad + \quad \frac{1 - y}{x} d \quad [3]$$

$$c = \frac{y}{x} b \quad + \quad \frac{x - y}{x} d \quad [4]$$

1. Utility is, of course, invariant only under positive linear transformations. It is simple to show, however, that the solution of a 2 x 2 zero-sum game is unchanged by any linear transformation defined on the payoffs of the game.

The value of the game, z, may also be expressed as a linear combination of payoffs b and d:

$$z = by + d(1 - y) \qquad [5]$$

Solving 5 for b, we obtain:

$$b = \frac{z + yd - d}{y} \qquad [6]$$

Note that there remains one degree of freedom in our solution: having specified the two players' optimal strategies and the value of the game, we are free to choose one entry of the payoff matrix before a unique 2 x 2 zero-sum game is defined. Therefore, we are assured that classes of 2 x 2 zero-sum games will have more than one member.

Alternatively, we may wish to specify a class of 2 x 2 zero-sum games equivalent to a given game.[2] Since linear transformations of the payoffs of a game leave the game's minimax strategies unchanged, we seek (a class of) linear transformations which will not alter the game's value. Let us proceed by defining the linear transformation p' = rp + s on the payoffs of Game 1; if z' denotes the value of the transformed game, it can easily be

2. This alternative approach was suggested in a conversation with Henry Hamburger, University of California (Irvine). Dr. Hamburger also suggested the following reparametrization of our 4-parameter specification for 2 x 2 zero-sum games:

The counter-strategy function provides a fourth parameter for 2 x 2 zero-sum games without saddle points. Of course the payoffs themselves, a, b, c, and d constitute a 4-parameter specification of such games, but each of them individually is not particularly meaningful. More meaningful are x, y, and z, the optimal strategy mixtures for the two players and the value of the game. In addition to x, y, and z, a fourth parameter is needed to fully specify a game.

Referring to the counter-strategy function, we see that it is linear with slope: k(a-b-c+d) + b − d.

This slope is a measure of how much can be gained through "exploitation" of the nonoptimal play of the opponent by row per unit of change in his strategy mixture. Let us call this quantity the "opportunity rate" for row, the subject, and imagine k to be controlled by the experimenter through a stooge. Then the experimenter's "control of the opportunity rate"—that is, the change in opportunity rate per unit of change in stooge strategy mixture—is a − b − c + d.

Formally, simply define w by:

$$w = \frac{\partial}{\partial(k)} \; \frac{\partial}{\partial p(b_1)} \; E(\$) = a-b-c+d \qquad [a]$$

Solving equations 1, 2, 5, and [a] for a, b, c, and d in terms of w, x, y, and z yields:

shown that z' = rz + s. Setting z' = z and solving for the coefficient r or s, we get:

$$r = \frac{z - s}{z} \qquad [7]$$

or

$$s = z(1-r) \qquad [8]$$

To recapitulate, equivalent 2 x 2 zero-sum games may be constructed by (i) specifying minimax strategies for each player, the value of the game, and one entry of the payoff matrix, or (ii) performing a linear transformation (of the type defined above) on the payoffs of a given game.

Now we construct a pair of equivalent games; set x = 6/14, y = 3/14, and z = –4/14. Let d = 1. Substituting these values into equations 3, 4, and 6 produces Game 2. This is the game used by Fox (1972) in previous experiments.

B \ A	a_1	a_2
b_1	6	-5
b_2	-2	1

Game 2

B \ A	a_1	a_2
b_1	98/9	-26/3
b_2	-10/3	2

Game 3

If we set d = 2 and retain the previous values of x, y, and z, we obtain Game 3, which is, of course, equivalent to Game 2 by construction.

$$a = z + w\bar{x}\bar{y}$$
$$b = z - wx\bar{y}$$
$$c = z - w\bar{x}y$$
$$d = z + wxy \qquad \text{where } x = 1-\bar{x},\ y = 1-\bar{y}. \qquad [b]$$

This formulation shows that w, x, y, and z may be chosen independently. Although a choice of w, x, y, and z uniquely determines the values of a, b, c, and d, it is nevertheless true that interchange of rows or columns or players alters both sets of parameters without significantly altering the game. For example, interchange of rows corresponds to replacement of y by $\bar{y}$ and w by $-w$.

Alternatively, we could define the linear transformation p' = (16/9)p + 2/9 on the payoff matrix of Game 2. Note that the coefficients of this transformation are related as in equation 7 or 8.

"Nonrational" Counterstrategies

Although a player whose opponent plays a minimax strategy cannot influence the expected payoff of a 2 x 2 zero-sum game without saddle point, it is possible for the player to take advantage of his adversary's departures from minimax by appropriately adjusting his own strategy-choice behavior (compare with all this, Luce and Raiffa's 1957 geometric representation of zero-sum games). If his opponent adopts a nonoptimal strategy, it is to a player's advantage to choose one of his alternative pure strategies at all times, since expected payoff becomes a linear function of the proportion in which he mixes his strategy alternatives.[3]

Referring again to Game 1, let

$$k = p(a_1), [0 \leqslant k \leqslant 1];$$

then,

$$E(\$) = p(b_1)ka + p(b_1)(1-k)b + (1-p(b_1))kc + (1-p(b_1))(1-k)d \quad [9]$$

$$= p(b_1)(k(a-b-c+d) - d + b) + k(c-d) + d \quad [10]$$

For Games 2 and 3, equation 10 becomes,

Game 2:

$$E(\$) = p(b_1)(14k-6) - 3k + 1 \quad [11]$$

Game 3:

$$E(\$) = p(b_1)(24.9k - 10.67) - 5.33k + 2 \quad [12]$$

These equations suggest an interesting graphic representation for 2 x 2 zero-sum games. Let us define a two-dimensional coordinate space in which the horizontal axis represents the proportion of plays on which one of ROW's alternatives (let us say b_1) is chosen and the vertical axis

3. A line with nonzero slope defined on the unit interval will have its maximum at either 0 or 1.

represents the the expected payoff of the game to ROW. The choice of a specific random strategy by COLUMN chooser defines a line in the space. There are, of course, infinitely many such lines corresponding to the infinite set of mixed strategies ($p(a_1)$, $1-p(a_1)$) available to COLUMN. These lines intersect at the coordinates defined by the value of the game and ROW's minimax strategy. If $p(a_1) = x$, the slope of the function becomes identically O, and E(\$) = z. We may term this line the "counter-strategy function" for ROW in a particular game when his opponent plays a specific random strategy mixture.

For fixed $k \neq x$, the counter-strategy function for Game 4 has a steeper slope (positive or negative, depending on the direction of departure of k and x) than that for Game 2. Hence, ROW chooser in Game 3 can expect a greater return from departures of a given size from his minimax strategy (assuming that his opponent plays some nonrational mixed strategy) than can his counterpart playing Game 2. The situation for Games 2 and 3, for k = .6 and k = .7 is depicted in Figure 1.

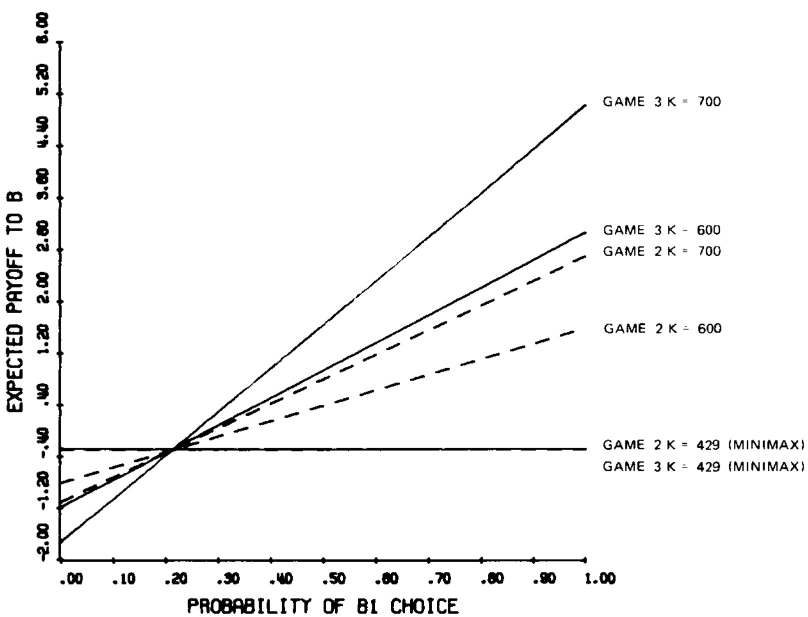

Figure 1.

Study Design

PURPOSE

The experiment reported here was designed to determine whether subjects respond (i.e., play) differently in equivalent games when their opponent departs from his prescribed minimax strategy.

SUBJECTS

Subjects were 32 students from the University of Michigan Paid Subject Pool, who were paid, as described below, on the basis of their performance in the experiment.

PROCEDURE

Subjects played a 2 x 2 zero-sum game without saddle point against a programmed stooge (a computer program). The instructions informed subjects that success in the game is related to intelligence and that the computer would "try" to make them lose as much money as possible. Payoffs were represented as "imaginary" dollars and, at the termination of the experimental session, subjects were paid at the rate of 1/2 cent (real money) to each imaginary dollar they had accumulated. Subjects' earnings depended upon experimental conditions as well as upon the manner in which they played. These earnings ranged from under $1.00 (real money) to more than $4.00; average earnings were about $2.00. An experimental session consisted of 200 plays of the game and typically lasted somewhat less than one hour.

Subjects were informed of the outcome of each play immediately after choosing one or the other of their strategy-choice alternatives; every twenty trials they received feedback about (a) their winnings over the last twenty plays; (b) the number of times each of their alternatives was chosen in the last twenty plays; (c) the current status of their bankroll. Subjects interacted with the computer through a teletype and, consequently, had a hard copy of the time course of the game.

GAMES

Experimental subjects played either Game 2 or Game 4. In all cases, subjects were ROW choosers while the computer program assumed the role of COLUMN chooser.

The properties of Game 2 have been discussed above. Game 4 is essentially identical to Game 3, with the exception that in Game 4 the entry in cell *a* has been rounded to 11. For practical purposes, however, Games 2 and 4 may be considered to be "equivalent" in the sense defined above.

		A: a_1	A: a_2
B	b_1	11	-8.67
B	b_2	-3.33	2

Game 4

EXPERIMENTAL CONDITIONS

Subjects were randomly assigned to one of four experimental conditions, eight subjects in each condition:

(1) Game 2, opponent departing relatively moderately from his minimax strategy (G2M).

(2) Game 2, relatively large opponent's departure from minimax (G2L).

(3) Game 4, moderate departure from minimax (G4M).

(4) Game 4, large departure from minimax (G4L).

In the "moderate departure" conditions, the computer assigned a probability of .6 to its a_1 choice; in the two "large departure" conditions, $p(a_1) = .7$. Subjects' opportunities for exploiting their opponents' departures from rational play in the several conditions are summarized in Figure 1 (Game 4 is substituted for Game 3).

Our experiment may, therefore, be described as a 2 x 2 x repeated measures design (in which trials are blocked in units of twenty) where game and opponent's strategy are between-subject factors.

Results

The presentation of the results of the experiment will be divided into three sections. First, we shall show how the intended experimental manipulations were in fact realized. Second, we shall examine data on the performance of subjects in the several experimental conditions as they relate to the two between-subjects factors and time. Last, we shall consider the relevance of our data to two questions that have been raised in the literature on zero-sum games:

(1) Do subjects show evidence of learning nonrational counter-strategies in zero-sum games?

(2) Do random fluctuations in a stooge's strategy-choice behavior (around some fixed mixed strategy) influence the choice behavior of experimental subjects?

VERIFYING THE EXPERIMENTAL MANIPULATIONS

For the manipulation associated with each of the experimental conditions to be successful, it is necessary that the programmed stooge play his assigned mixed strategy. The random numbers used by the program and the choices generated from these numbers were checked for apparent randomness before the experiment was run, following the methods described in Naylor et al. (1966).

To verify that the relationship between subject strategy (i.e., proportion of b_1 choices) and payoffs were consistent with our intentions, we regressed the actual earnings of subjects in twenty-trial blocks on their proportion of b_1 choices (again, in twenty-trial blocks) for each of the experimental conditions. Expected regression coefficients are computed from equation 10 and (substituting Game 4 for Game 3) describe the lines that appear in Figure 1. Results of the regression analyses are presented in Table 1.

TABLE 1
REGRESSIONS OF PAYOFFS ON STRATEGY ($p(b_1)$) FOR EACH EXPERIMENTAL GROUP

Group	*Coeff.*	*1. Observed Value*	*2. Expected Value*	*t(1-2)*	*p*	*$t(H_O)$*	*p*
G2M	Slope	48.46	48.00	.06	.96	5.87	.001
	Intercept	−20.21	−16.00	−.79	.42		
	r	.554					
G2L	Slope	60.63	76.00	−1.07	.29	4.24	.001
	Intercept	−15.18	−22.00	.58	.55		
	r	.433					
G4M	Slope	5.22	86.60	−3.59	.001	.23	.82
	Intercept	20.46	−23.96	3.22	.01		
	r	.026					
G4L	Slope	131.50	136.60	−.25	.80	6.42	.001
	Intercept	−35.32	−34.62	−.05	.96		
	r	.588					

TABLE 2
ALTERNATIVE EXPERIMENTAL DESIGNS

	Contrast Coefficients for Group			
Contrast	*G2M*	*G2L*	*G4M*	*G4L*
(A)				
Game	1	1	−1	−1
Stooge	1	−1	1	−1
G x S	1	−1	−1	1
(B)				
Game	0	1	0	−1
(Stooge)	1	−½	0	−½

Note that the regression coefficients for three of the regression lines do not significantly differ from the parameters of the theoretically expected lines. However, quite surprisingly, the relationship between payoffs and performance for G4M group subjects is strikingly different from our expectation. Indeed, for these subjects, there is not even a significant positive relationship between proportion of b_1 choices and winnings.

Since the experimental manipulation apparently failed in one of the conditions, yet was successful in the other three,[4] we are faced with a choice between alternative data-analysis strategies: (a) we might proceed with the analysis as originally planned; or (b) we could discard data for G4M subjects and analyze only data for the remaining three groups. Viewing these alternative approaches in terms of experimental designs, the "traditional" main effects and interaction effects analysis we originally intended implies the set of contrasts whose coefficients appear in part A of Table 2. If we eliminate the G4M group, we are left with two degrees of freedom among the three remaining groups. Since the primary question we wish to ask of our data concerns subjects' ability to differentiate "equivalent" games, one degree of freedom could be used to compare the G2L with the G4L group. We could then (partially) address the question of the impact of stooge strategy by comparing the G2M group to the G2L and G4L groups (see part B of Table 2).

We shall, in fact, combine these alternative analysis strategies in the following manner: first, we shall make descriptive comparisons implied by the set of contrasts in part B of Table 2, and then examine the behavior of

4. This result is quite curious, since the same series of random numbers was used to generate the stooge's choices in all experimental conditions. Therefore, the series of stooge choices were identical for G4M and G2M subjects.

G4M group subjects to determine whether their performance diverges from that of subjects in the other groups; finally, we shall let the results of our descriptive analysis govern our decision concerning the disposition of data from the G4M group.

SUBJECTS' PERFORMANCE

Figures 2 and 3 display subject performance for each experimental group, along with the performance of the programmed stooge. We note that subjects in the G4M group do not appear to exhibit behavior that is obviously divergent from that of subjects in the other experimental groups. The performance of G4M subjects is slightly below that of their counterparts in the G2M condition–an unexpected occurrence, which might be ascribed to failure of the experimental manipulation; however, a similar unexpected difference in performance holds true for G4L and G2L subjects. Similarly, any descriptive conclusions we might draw about the effect of stooge strategy would apparently be unchanged by the inclusion of the G4M group in the analysis. We shall, therefore, proceed as originally planned.

The analysis of variance for data in Figures 2 and 3 appears in Table 3. The assumptions underlying the repeated-measures model analysis of variance were checked against the data using a procedure described in Winer (1962: 369) and a satisfactory fit of model to data was observed. While the computer's strategy significantly affects the subject's behavior in the expected manner (i.e., the more the stooge departs from minimax, the more its departure is exploited), the game factor appears to have no significant effect on performance. Indeed, looking at Figures 2 and 3, we may note that such differences as are observed are not in the expected direction: subjects exploit their opponent's nonrational play *slightly less* in the game (Game 4) which provides greater opportunities for exploitation.

That our subjects' performance tends to improve over the course of the game is reflected in a significant linear trend over time. Looking at Figures 2 and 3, however, it becomes apparent that this improvement is not dramatic.

There is an aspect of the relationship between strategy-choice behavior and winnings, other than expected reward, which is possibly relevant to subjects' ability to take advantage of an opponent's departures from minimax; namely, the strength of association between strategy and payoff. We might suppose that subjects would be better able to take advantage of an opponent's nonoptimal play if their behavior were highly correlated with their winnings than if this relationship were weak—even if, in the first case, the slope of the "counter-strategy function" was less steep than in the second case. However, such an explanation of subjects' behavior in the present situation is not tenable (see correlations in Table 1). For example, G2L subjects do not perform at a lower level than G4L subjects, although the latter group is at an advantage both in terms of the slope of their counter-strategy function (131.5 versus 60.63) and with regard to the strength of association between strategy and payoff (r = .588 as compared to r = .433).

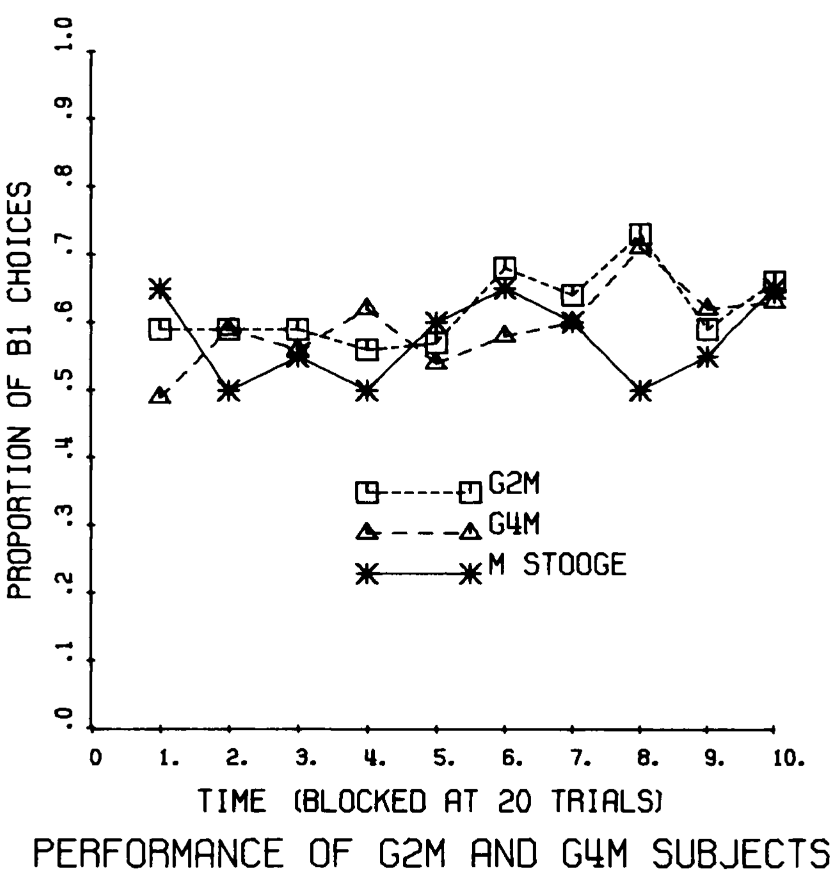

Figure 2.

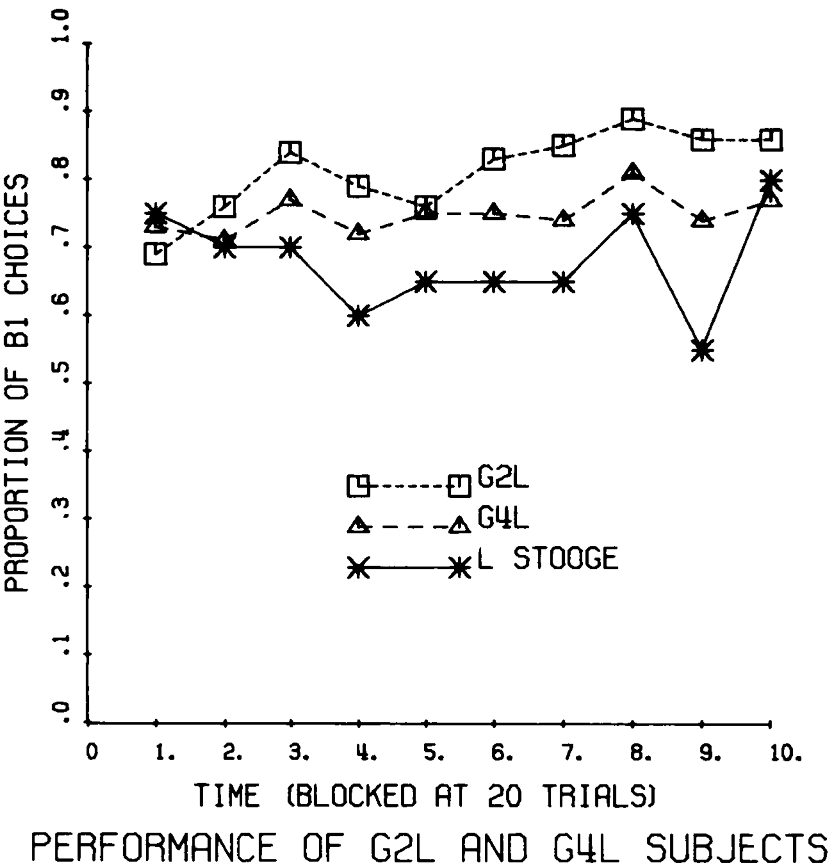

Figure 3.

TWO QUESTIONS POSED IN PREVIOUS RESEARCH

In this section, we shall consider how the results from our experiment relate to two questions posed in earlier experimental literature on zero-sum games:

(1) Can subjects learn to increasingly take advantage of an opponent's non-rational strategy-choice behavior (see Fox, 1972)?

(2) Are subjects responsive to random fluctuations in their opponent's sequence of choices (see Fox, 1972; Lieberman, 1962)?

TABLE 3
ANALYSIS OF VARIANCE OF SUBJECTS' PERFORMANCE

Source	*D.F.*	*Mean Square*	*F*	*Significance*
Between S's	*31*			
A (game)	1	0.15976	0.989	ns
B (Stooge)	1	2.45876	15.223	.01
AB	1	0.03101	0.192	ns
S's w/groups	28	0.16151		
Within S's	*288*			
T (time)	9	0.06220	6.894	.01
T (linear)	1	0.33245	23.777	.01
AT	9	0.00531	0.589	ns
AT (linear)	1	0.01220	0.878	ns
BT	9	0.00661	0.732	ns
BT (linear)	1	0.00043	0.031	ns
ABT	9	0.01313	1.455	ns
ABT (linear)	1	0.0347	2.484	ns
T x S's w/groups	252	0.00902		
T x S's w/groups (linear)	28	0.01398		

TABLE 4
CORRELATIONS AMONG MEASURES OF (A) ASSOCIATION OF PERFORMANCE AND PAYOFF AND (B) PERFORMANCE

Variables[a]	*1.*	*2.*	*3.*	*4.*	*5.*
1. r(s,$)	1.0				
2. b($,s)	0.8168	1.0			
3. b(s,t)	–.0933	–.2792	1.0		
4. Change	–.0195	–.2921	0.9150	1.0	
5. Adjusted change	0.0184	–.2693	0.6106	0.6023	1.0

a. Defintion of Variables
1. Correlation between $p(b_1)$ and payoff.
2. Slope of regression line of payoff on $p(b_1)$.
3. Slope of regression line of $p(b_1)$ on time.
4. Mean performance performance over last 5 20-trial blocks minus mean performance over first five 20-trial blocks.
5. Change divided by possible improvement from first half mean.

Our observation, that there is a significant linear trend in our subjects' performance in the direction of increasingly taking advantage of their opponent's nonoptimal play, lends credibility to the hypothesis (Fox, 1972) that subjects can learn counter-strategies in zero-sum games.

However, more is involved in the concept of learning than merely improvement in performance; such a change in behavior is usually predicated upon the reward contingencies of a learning situation. Hence, we would expect that subjects whose performance is more closely related to the magnitude of their winnings (or whose performance tends to have greater influence upon their winnings) should show greater gains in performance than their counterparts whose rewards are less contingent upon their behavior (or whose rewards are less sensitive to changes in behavior).

Table 4 presents intercorrelations among two measures of the relationship between a subject's strategy and his payoff (computed for each subject) and three measures of subjects' improvement over time. (How each measure was computed is explained in Table 4.) Note that, although the correlations among variables within each set of measures range from high to moderate, those between measures from different sets are quite low (and, with one exception, negative). In contrast, Fox (1972) found a significant positive association of r(s,$) and b(s,t) for subjects in his experiment. Our data, then, cast some doubt upon whether our subjects' improved performance can reasonably be described as a learning process.

Lieberman (1962) has advanced the hypothesis that subjects respond to random fluctuations in a stooge's play when the stooge follows some mixed strategy. Fox (1972) found some, although weak, support for Lieberman's hypothesis.

To assess, in the present study, the dependence of subject behavior on variations in the stooge's play, we look separately at subjects playing against the M and L stooge strategies. Since there was no significant effect of the game factor, G2L and G4L subjects are classed together, as are G2M and G4M subjects. An analysis of variance for each of these classes of subjects allows us to partition the between-trial block variance into three components: a component due to a linear trend, a component due to responsiveness to fluctuations in the stooge's play, and a residual component.[5] Results of these analyses appear in Table 5. Note that, while about half of the between-trial block sum of squares for G2M and G4M subjects is contained in the one-degree-of-freedom linear comparison, and more than half of between-trial-block variation for G2L and G4L subjects can be ascribed to a linear trend, very little variation is accounted for by the contrasts generated from the stooge's play.

5. These linear and "responsiveness" comparisons are not strictly orthogonal. However, prior to the experiment, the series of stooge choices was determined to be virtually uncorrelated with time. Hence, the linear and "responsiveness" comparisons

TABLE 5
ANALYSIS OF VARIANCE TABLES WITH LINEAR AND "RESPONSIVENESS" COMPARISONS FOR G2M-G4M AND G2L-G4L SUBJECTS

Source	*SS*	*df*	*MS*	*F*	*Significance*
G2M & G4M Subjects					
Between S's					
S's w/grps	2.68125	15	.17875		
Within S's					
T (time)	.36969	9	.04107	3.537	.01
T (linear)	.17850	1	.17850	9.712	.01
T $(p(a_1))$	.01639	1	.01639	1.319	ns
T x S's w/grps	1.56781	135	.01613		
TS (linear)	.27570	15	.01838		
TS $(p(a_1))$	.18645	15	.01243		
G2L & G4L Subjects					
Between S's					
S's w/grps	2.03185	15	.13546		
Within S's					
T (time)	.24952	9	.02772	4.293	.01
T (linear)	.15438	1	.15438	14.231	.01
T $(p(a_1))$	.00319	1	.00319	.370	ns
T x S's w/grps	.87173	135	.00646		
TS (linear)	1.16272	15	.01085		
Ts $(p(a_1))$	.12931	15	.00862		

Conclusion

The present study has attempted to clarify certain aspects of choice behavior in two-person zero-sum games without saddle points. An experiment was conducted to determine whether subjects could learn to exploit a nonoptimal opponent and whether games which are "equivalent," but which offer different incentives to exploit the other, are played differently.

The results of the experiment indicate that naive subjects do learn to exploit a programmed stooge opponent who plays a nonoptimal strategy mix and that the subject's performance in the game shows a significant improvement over trials. This improvement in performance requires that the player systematically depart from his own optimal strategy-choice mix so as to exploit the other's nonoptimal choices. With regard to the second

are nearly orthogonal. The "responsiveness" contrast coefficients are obtained from the stooge's twenty-trial block proportions by normalizing these proportions to a mean of 0.

variable, the results showed no significant difference in subjects' strategy-choice behavior between games having identical values and optimal strategies but differential payoffs for exploitation of a nonoptimal other.

The finding that equivalent games do not appear to be played differently suggests opportunities for further research. We may pose the question: To what parameters of a zero-sum game do subjects respond when they are playing against an opponent who plays nonoptimally? It is conceivable, for example, that subjects in such situations respond to the order of payoffs in the game matrix rather than to the numerical payoffs themselves.

REFERENCES

BRAYER, A. R. (1964) "An experimental analysis of some variables of minimax theory." Behavioral Sci. 9: 33-44.

FOX, J. D. (1972) "Learning of strategies in a simple two-person zero-sum game without saddle point." Behavioral Sci. 17: 300-308.

KAUFMAN, H. and G. M. BECKER (1961) "The empirical determination of games-theoretical strategies." J. of Experimental Psychology 61: 462-468.

LIEBERMAN, B. (1962) "Experimental studies of conflict in some two and three person games," in J. H. Griswell et al. (eds.) Mathematical Methods in Small Group Processes. Stanford: Stanford Univ. Press.

--- (1960) "Human behavior in a strictly determined 3 x 3 matrix game." Behavioral Sci. 5: 317-322.

LUCE, R. D. and H. RAIFFA (1957) Games and Decisions: Introduction and Critical Survey. New York: John Wiley.

MESSICK, D. M. (1967) "Interdependent decision strategies in zero-sum games: a computer-controlled study." Behavioral Sci. 12: 33-48.

MORIN, R. E. (1960) "Strategies in games with saddle points." Psych. Reports 7: 479-485.

NAYLOR, T. H., J. L. BALINFY, D. S. BORDICK, and K. CHU (1966) Computer Simulation Techniques. New York: John Wiley.

RAPOPORT, A. (1966) Two-Person Game Theory: The Essential Ideas. Ann Arbor: Univ. of Michigan Press.

--- and A. M. CHAMMAH (1965) Prisoner's Dilemma: A Study of Conflict and Cooperation. Ann Arbor: Univ. of Michigan Press.

VON NEUMANN, J. and O. MORGERNSTERN (1947) The Theory of Games and Economic Behavior. Princeton: Princeton Univ. Press.

WINER, B. J. (1962) Statistical Principles in Experimental Design. New York: McGraw-Hill.