

An Experimental Investigation of Unprofitable Games¹

John Morgan

*Woodrow Wilson School and Department of Economics, Princeton University,
Princeton, New Jersey 08544*

and

Martin Sefton

*School of Economics, University of Nottingham, Nottingham
NG7 2RD, United Kingdom*

Received September 21, 1999; published online April 8, 2002

We investigate behavior in two *unprofitable games*—where Maxmin strategies do not form a Nash equilibrium yet guarantee the same payoff as Nash equilibrium strategies—that vary in the riskiness of the Nash strategy. We find that arguments for the implausibility of Nash equilibrium are confirmed by our experiments; however, claims that this will lead to Maxmin play are not. Neither solution concept accounts for more than 53% of choices in either game. The results indicate that the tension between the Nash and Maxmin strategies does not resolve itself over the course of the experiment. Moreover, the relative performance of the solution concepts is sensitive to the riskiness of the Nash strategy. *Journal of Economic Literature* Classification Numbers: C72, C92. © 2002 Elsevier Science (USA)

Key Words: unprofitable games; Nash equilibrium, maxmin.

1. INTRODUCTION

Determining how rational individuals will play a particular game is perhaps the fundamental question that game theory seeks to answer. The most frequently employed solution concept for answering this particular question

¹ We are grateful to the anonymous referees as well as participants at the ESA conferences in Mannheim, 1998, and Tucson, 1998, and seminar participants at Indiana University, Purdue University, IUPUI, and University of Birmingham for helpful comments. We are especially grateful to Abdullah Yavas for his assistance in conducting the US experiments. Morgan's work was supported by the National Science Foundation and the Woodrow Wilson School.

	LEFT	RIGHT
UP	2, 6	4, 2
DOWN	6, 0	0, 4

FIG. 1. Aumann's example.

is Nash equilibrium. However, in some games it has been argued that Nash equilibrium is *not* a plausible prediction about how rational individuals would play.² Aumann (1985) offers the example shown in Fig. 1.

The strategies comprising the unique Nash equilibrium in this game consist of Row player choosing UP and DOWN with equal probability and Column player doing likewise with LEFT and RIGHT. If both players play the equilibrium strategies, then each earns an expected payoff of 3. Of course, if Column were instead only to choose RIGHT while Row continued to play her equilibrium strategy, then Row's expected payoff would be reduced to 2. Moreover, as long as Row plays the equilibrium strategy, Column is not disadvantaged in any way by switching to only playing RIGHT.

Aumann argues that the Nash equilibrium of this game is implausible since Row has a means of assuring herself an expected payoff of 3 regardless of the strategy played by Column. Were Row to play UP with probability $\frac{3}{4}$ then, for any choice by Column, Row obtains an expected payoff equal to 3—exactly what she obtained when both players were playing the equilibrium strategies—but without the risk of lower payoffs from Column doing unexpected things. Aumann suggests that the availability of this secure strategy undermines the Nash equilibrium prediction: "This risk is quite unnecessary, since player 1 has the Maxmin strategy $(\frac{3}{4}, \frac{1}{4})$ available, which assures him of 3 regardless; similarly player 2 has strategy $(\frac{1}{4}, \frac{3}{4})$. Under these circumstances it is hard to see why the players would use their equilibrium strategies" (p. 668).

In the terminology of Harsanyi (1966), the game in Fig. 1 is *unprofitable* to each player. That is, for each player no equilibrium yields more than the Maxmin payoff. In many unprofitable games Nash and Maxmin strategies do not coincide, and for these games Harsanyi (1964, 1966, 1977), Aumann and Maschler (1972), and Aumann (1985) have all argued that Nash equilibrium is a poor prescription/description of how rational players would play. In this paper, we examine this debate from an experimental perspective. Specifically, we examine whether or not Nash equilibrium profiles are reasonable descriptions of how experimental subjects play these games and, if not, what strategies are employed. To make identification of strategies more transparent, we consider symmetric unprofitable games that have

²Other solution concepts, for example, Rationalizability and Cautious Rationalizability (Pearce, 1984; Bernheim, 1984) might be applied, but in all the games we discuss these have no predictive power, as they do not rule out any strategy profiles.

a unique pure strategy Nash equilibrium distinct from the unique Maxmin pure strategy.

While the arguments against Nash equilibrium in this context appear reasonable, they still leave unanswered the question of how rational individuals play unprofitable games. Harsanyi (1977), in specifying postulates of rational play in games, suggests, “If a player cannot hope to obtain *more* than the Maxmin payoff, then that player should use a strategy fully assuring at least *that much*” (p. 116). This motivates the following:

A1. *Maxmin postulate*. In any game G unprofitable to you, always use a Maxmin strategy. (p. 116)

This postulate offers a sharp prediction at odds with Nash equilibrium for how rational players will play some unprofitable games.³ For precisely for this reason, it is somewhat controversial: van Damme (1980) has argued that a “good set of axioms” cannot include this postulate since any axiomatic theory of rational decision-making must yield equilibrium point solutions for all non-cooperative games. By having the Maxmin strategy occur as a pure strategy, our experimental design should give a clear indication of whether or not this postulate is a good description of actual behavior.

The arguments against Nash equilibrium play and in favor of Maxmin depend implicitly on the beliefs about “unexpected” play by an opponent. For instance, in Aumann’s example, Row is clearly worse off playing the Nash equilibrium strategy if Column plays only RIGHT, but is better off if Column plays only LEFT. Thus, the argument that Nash equilibrium play is risky essentially means that Row judges “excess” RIGHT play to be more likely than “excess” LEFT play.

To recommend Maxmin play requires more than this. In the above example, if Row really believes that excess RIGHT play is prevalent, her best alternative is not to play the Maxmin strategy but rather to play UP. Thus, to obtain the Maxmin prediction, it must be that weights Row places on the proportion of LEFT and RIGHT play change as her own strategy changes. That is, if Row counters excess RIGHT play by choosing UP, then she must believe that Column’s behavior will also change (in a detrimental fashion) to mostly LEFT play. Indeed, part of Harsanyi’s motivating argument (“cannot hope to obtain *more* than the Maxmin payoff—”) seems to suggest that, for any strategy, these shifting pessimistic beliefs are the natural ones to ascribe to players. This reasoning may be made more formal by using Lo’s (1996) game theoretic adaptation of the multiple prior

³In the remainder of the paper, we consider only unprofitable games where Nash and Maxmin strategies are distinct. In games where the two concepts coincide, for example, zero-sum games or games with dominant strategies, the Maxmin postulate does not challenge the equilibrium prediction.

framework of Gilboa and Schmeidler (1989). We discuss this in more detail in the next section.

We anticipated that the degree of “riskiness” associated with playing the Nash equilibrium was likely to play an important role in its success relative to the Maxmin postulate. To assess this hypothesis, we examined behavior in two unprofitable games that differ from one another in terms of a simple measure of riskiness. Both games are two-player 3×3 symmetric games with a unique pure Nash equilibrium strategy (labeled A), a unique Maxmin (pure) strategy (C), and a third strategy (B) which is a best response to Maxmin play. In both games, a Nash player benefits from B play by her opponent but is made worse off (relative to Maxmin) by C play. In the “Baseline game,” the Nash strategy delivers less than the Maxmin strategy if C play is more than two-thirds as likely as B play. In the second game, which we call the “Upside Game,” C play must be at least 2.67 times more likely than B play for the Nash strategy to yield a lower payoff than Maxmin. Thus, Nash play in the Baseline game may be viewed as riskier than in the Upside game since the proportion of unexpected C play can be quite low yet still lead to payoffs lower than that obtained by playing the Maxmin strategy. In the next section, we show that other measures of riskiness, based on multiple prior models, lead to the same conclusion.

By varying the riskiness of Nash play in our treatments, we examine a number of questions in the extant theoretical literature. First, does the riskiness of Nash equilibrium strategies undermine the descriptive power of Nash equilibrium? Our results suggest the answer to be yes. In our Baseline Game, only 14% of choices corresponded to the Nash strategy. In the Upside Game, the Nash strategy was modal, but still only accounted for 47% of choices.

Second, is Maxmin play a good description of how subjects play unprofitable games? We obtain a negative answer to this question. In the Baseline Game, the Maxmin strategy accounts for the majority (53%) of choices, but this still leaves a substantial proportion of non-Maxmin choices. In the Upside Game, the Maxmin strategy accounts for only 30% of choices.

Third, are pessimistic beliefs a good description of how subjects evaluate unprofitable games? We reject this contention. In the Upside Game, choices corresponding to Nash occur almost four times as often as in the Baseline game. This suggests that variation in the *upside* riskiness of Nash strategies does affect subjects’ propensities to choose them.

A better sense of the performances of Nash and Maxmin strategies is obtained by comparing them with alternative models. In this paper we consider three alternatives: a random play model, a mixed-strategy equilibrium model, and McKelvey and Palfrey’s (1995) quantal response equilibrium model (QRE). We find that, in the class of models considered, the QRE outperforms all other models in predictive power, according to

standard scoring rules applied to out-of-sample predictions, while the two models most widely discussed, pure Nash and Maxmin, do worst.

Our experiments have two especially noteworthy features. First, in our games, the Maxmin strategy is quite transparent since it is a *pure* strategy with a constant payoff, and there is a Nash equilibrium in pure strategies as well. This seems appropriate in light of the debate about whether subjects play mixed strategies (see Brown and Rosenthal, 1990; O'Neill, 1987; and Shachat, 1996). Second, we use two parameterizations that explicitly manipulate the riskiness of equilibrium strategies. This allows us to directly test the implicit Maxmin assumption that subjects entertain pessimistic beliefs when making choices.⁴

The remainder of the paper is organized as follows. In Section 2 we describe unprofitable games and discuss alternative solution concepts. Our experimental design and procedures are outlined in Section 3. We present our results, including an evaluation of various models, in Section 4. Section 5 concludes.

2. UNPROFITABLE GAMES

In this section, we describe the two unprofitable games that we used to examine behavior. As mentioned previously, the games we consider have a unique pure strategy Nash equilibrium and a different unique pure Maxmin strategy. In addition, we sought several other features to reduce the complexity of the experimental environment. First, we restricted our attention to symmetric games. Second, we sought a minimal number of pure strategies. Finally, we sought a minimal number of payoff levels in the payoff matrices of each of the games. Below, we establish that symmetric 3×3 games with three payoff levels best meet these criteria.

We begin by showing that, regardless of symmetry, having only two pure strategies is insufficient to meet our objective of having distinct Nash and Maxmin outcomes occur in pure strategies.

⁴To our knowledge, the only previous experimental investigations of unprofitable games are papers by Holler and Host (1990) and Ochs (1995). Both of these papers report results on 2×2 asymmetric unprofitable games with both Nash and Maxmin profiles occurring in mixed strategies. Holler and Host find “significant evidence in favor of maximin”; however, their subjects were given no financial incentives for their decisions, nor were they playing real opponents in making their strategy decisions. Ochs rejects both Maxmin and Nash profiles in describing aggregate subject behavior. Ochs also reports estimates made with the QRE model as well as simulations using a choice-reinforcement learning model and concludes that they both fit the data better than the static predictions.

	A	B	C
A	40, 40	60, 10	10, 40
B	10, 60	10, 10	60, 40
C	40, 10	40, 60	40, 40

FIG. 2. Baseline Game.

PROPOSITION 1. *There does not exist any 2×2 game with distinct Nash and Maxmin profiles in pure strategies.*⁵

Thus, we are forced to have three or more pure strategies to meet our criteria. The following proposition shows that with symmetric 3×3 games, we must have more than two payoff levels as well.

PROPOSITION 2. *There does not exist a symmetric 3×3 unprofitable game with two payoff levels and distinct Nash and Maxmin strategies.*

Our Baseline Game, shown in Fig. 2, establishes that there is a 3×3 unprofitable game with three payoff levels that does meet our criteria.

In this game the Maxmin strategy is C, which guarantees a payoff of 40. This is not a Nash equilibrium, as B is a best response to C. In turn, A is a best response to B. The unique pure Nash equilibrium is (A, A), which delivers a payoff of 40.

All of the implications of interest in the Baseline Game occur when subjects simply choose from among the pure strategies; thus, we do not need (or particularly want) subjects to be playing the mixed extension of the above normal-form game.⁶ If, however, we characterize equilibria arising in the mixed extension of the Baseline Game, we find that, in addition to the equilibrium (A, A), there is a symmetric mixed equilibrium where B is played with probability .4 and C with probability .6 as well as two asymmetric mixed-strategy equilibria.⁷

In light of this fact, it might be useful to have uniqueness of Nash equilibrium in the mixed extension of our candidate games. The following proposition shows that if a symmetric 3×3 unprofitable game has a pure strategy equilibrium and a distinct pure Maxmin strategy, then its mixed extension also has a symmetric mixed-strategy equilibrium.

PROPOSITION 3. *Every symmetric 3×3 unprofitable game with distinct Nash and Maxmin profiles in pure strategies has a symmetric Nash equilibrium in mixed strategies.*

⁵See Appendix A for a proof of this and other propositions in this section.

⁶The mixed extension of a normal form game is a game in which each player's strategy set consists of all probability distributions over pure strategies. See Osborne and Rubinstein (1994, Definition 32.1) for details.

⁷In an asymmetric equilibrium, one player chooses A with probability .4 and C with probability .6 while the other player chooses B with probability .6 and C with probability .4.

	A	B	C
A	40, 40	120, 10	10, 40
B	10, 120	10, 10	120, 40
C	40, 10	40, 120	40, 40

FIG. 3. Upside Game.

Proposition 3 shows that, in the class of games we study, the mixed extension *always* admits multiple equilibria; however, all of these mixed-strategy equilibria lead to the same expected payoff as the Maxmin strategy. Our view is that mixed-strategy predictions for the games studied are highly implausible. Notice that, for the mixed-strategy prediction to literally hold at the individual level, players must *actively* randomize over a set of pure strategies that includes the Maxmin strategy.⁸ It is hard to see why such a randomization would be willingly undertaken when a player can obtain the same (expected) payoff—with certainty—by simply playing the Maxmin strategy. That is, if there were any effort cost whatsoever to randomizing, a player would strictly prefer to choose Maxmin. This argument, along with several others, underlies a well-documented skepticism on theoretical grounds regarding mixed strategies.⁹ Furthermore, even in laboratory experiments where a unique mixed-strategy equilibrium corresponds with Maxmin play, these predictions have received, at best, limited support.¹⁰ For these reasons, we focus mainly on the pure strategy equilibrium prediction.

We vary the riskiness of Nash play, holding fixed the number of pure strategies and payoff levels, by increasing all of the 60 payoffs to 120. We refer to the resulting game as the Upside Game (see Fig. 3). Notice that this game is best response equivalent to the Baseline Game, and, hence, (A, A) is still the unique pure strategy equilibrium. Likewise, C is still the unique Maxmin strategy.¹¹

As mentioned previously, implicit assumptions about beliefs and riskiness of Nash play are at the heart of Harsanyi's Maxmin postulate and Aumann's critique of the Nash prediction. To formalize these ideas, we study the properties of our Baseline and Upside games, using a game-theoretic adaptation of Gilboa and Schmeidler's multiple prior model (Lo, 1996). To facilitate comparison, we adopt Lo's notation exactly. Let $S = S_1 \times S_2$ denote the set of pure strategy profiles for a given normal-form game. Let $g_i: S \rightarrow X$

⁸Of course, it could be that individuals all play pure strategies leading to a population with mixed-strategy proportions of each strategy.

⁹Rubinstein (1991) presents a wide-ranging discussion of these issues.

¹⁰See, for instance, Brown and Rosenthal (1990).

¹¹In the symmetric mixed-strategy equilibrium of the Upside Game A is played with probability 55/88, B is played with probability 9/88, and C is played with probability 24/88. This game has no asymmetric equilibria.

be the outcome function of the game for player i . Let $u_i: M(X) \rightarrow R$ represent i 's preference ordering over the set of lotteries of X , denoted by $M(X)$. Let $M(S_i)$ be the set of mixed strategies available to player i with typical element σ_i . Finally, i 's beliefs about the strategies that player j might select are given by the closed and convex set $B_i \subseteq M(S_j)$ with typical element p_i .

In this framework, a strategy σ_i is a best response to beliefs B_i if

$$\sigma_i \in \arg \max_{\sigma_i \in M(S_i)} \min_{p_i \in B_i} u_i(\sigma_i, p_i). \quad (1)$$

Equation (1) says that a strategy σ_i is a best response if it maximizes i 's payoff, given (pessimistic) multiple prior beliefs, B_i . With this definition in mind, the following is immediate:

Remark. Suppose that for all i , $B_i = M(S_j)$, then a best response for i is a Maxmin strategy.

The remark states that if a player's priors are completely diffuse, then she can do no better than to play a Maxmin strategy.¹²

To assess the relative riskiness of Nash strategies in the Baseline versus the Upside game with the use of this framework, we first identify the sets of beliefs against which A is a best response for each game. We let B_i^* and U_i^* denote these sets for the Baseline and Upside Games, respectively.

The set of beliefs against which A is preferred to B, C, and all convex combinations of A, B, and C in the Baseline Game is

$$B_i^* = \{\mathbf{p} | p_A \geq 1 - (5/3)p_B\},$$

where, in a slight abuse of notation, $\mathbf{p} = (p_A, p_B, 1 - p_A - p_B)$ denotes an element of the unit simplex, p_A denotes the probability assigned to strategy A, and p_B is the probability assigned to B.

For the Upside Game,

$$U_i^* = \{\mathbf{p} | p_A \geq 1 - (11/3)p_B, p_A \geq (11/14) - (11/7)p_B\},$$

where the first inequality ensures that A is preferred to B, and the second ensures that A is preferred to C.

One can readily verify that $B_i^* \subset U_i^*$. In other words, the set of beliefs where Nash play is a best response in the Baseline game is a strict subset of the set of beliefs where it is a best response in the Upside game, even when beliefs are "pessimistic" in the sense of the multiple priors model. Thus, by this measure, Nash play is less risky in the Upside Game than

¹²This remark is formalized as Proposition 1 in Lo (1996).

in the Baseline Game. However, for completely diffuse priors, Maxmin is predicted for both games.

Finally, notice that standard (single prior) beliefs are a special case of this framework where B_i is restricted to being a singleton. In that case, any belief that leads to a best response of A in the Baseline Game also leads to a best response of A in the Upside Game, but not viceversa. Thus, once again, we have that the Baseline game is riskier than the Upside game.

3. EXPERIMENTAL DESIGN AND PROCEDURES

The experiment consisted of 12 sessions, six conducted at the University of Newcastle (U.K.) in spring 1998 and six at Penn State University (U.S.) in summer 1998. Ten subjects participated in each session, and no subject appeared in more than one session. At each location, three sessions involved the Baseline game and three involved the Upside game. The U.K. and U.S. sessions allowed us to investigate the robustness of results across subject pools, since as far as was possible the same procedures were used in the two countries. For the U.K. sessions, subjects were recruited by an e-mail invitation to participate in a decision-making experiment. Participants were promised between £2 and £12 for a session lasting at most 75 minutes. For the U.S. sessions subjects were respondents to posters offering between \$3 and \$18 for a session lasting less than 1 hour. During the session, subjects accumulated points according to their decisions, and at the end of a session subjects were paid either 10p per 25 points (U.K.) or 25 ¢ per 40 points (U.S.).¹³

The following procedures were common to all sessions. At the beginning of a session, the 10 subjects were seated at computer terminals and given a set of instructions. These instructions were read aloud, and, at the end, subjects were given an opportunity to ask questions. Subjects then completed a quiz to ascertain that they understood how their choices translated into earnings.¹⁴

The session then consisted of 50 rounds. No communication between subjects was permitted, and all choices and information were transmitted via computer terminals. In each round, a subject chose from the options A, B, and C. When all subjects had made their choice, subjects were randomly paired and informed of their opponent's choice. Subjects then received point earnings according to their choice and the choice of the person with whom they were paired, according to one of the payoff matrices described in the previous section. At the end of each round subjects

¹³At the time of the experiments, £1 = \$1.65, so that the stakes are quite comparable.

¹⁴Appendix B contains copies of the instructions and quiz.

received additional feedback about the results of the round. Specifically, subjects were informed of the number of A's, B's, and C's chosen, and the average earnings of subjects who chose A, B, and C, respectively.

At the end of the session, subjects were paid in cash according to their accumulated point earnings, using the appropriate exchange rate. All sessions took fewer than 50 minutes, and average earnings were £7.44 (U.K. Baseline sessions), £9.26 (U.K. Upside), \$12.33 (U.S. Baseline), and \$14.53 (U.S. Upside). This is considerably more than outside earnings opportunities of the subjects. In fact, subjects completed a short post-experimental questionnaire and the question "Would you be willing to take part in other experiments of this sort?" received 119 out of 119 affirmative responses.¹⁵

Several features of the design are worth noting. Our goal was to choose the simplest possible unprofitable game with pure Nash and Maxmin strategies. As described above, this led us to restrict our attention to 3×3 symmetric games with three levels of payoff. We regarded this as sufficiently simple to ensure that subjects had a complete understanding of the game. In addition, we administered a quiz prior to the experiment to test subject understanding of the payoff matrix, and all subjects completed the quiz correctly before the decision-making part of the session began. We also tried to facilitate understanding by having subjects play the game repeatedly and providing them with feedback about the population proportions and average payoffs from playing each strategy in the previous period.

As we mentioned previously, our simple games still have some undesirable, but inevitable, features: the games have more than two levels of payoff and a mixed equilibrium. Despite this, to keep the design as simple as possible, we chose not to employ a binary lottery procedure to induce (theoretically) risk-neutral preferences.¹⁶

4. RESULTS

In this section, we examine aggregate behavior for evidence of venue or game treatment effects. Our main finding is that the game treatment (Baseline versus Upside) has a strong effect on behavior. The presence of a game effect is inconsistent with the notion that subject choices are based on extremely pessimistic beliefs. We then turn to a more detailed analysis assessing how well the various theoretical predictions of the Baseline and Upside game do in describing subject behavior. We find that

¹⁵There was one non-respondent.

¹⁶An additional consideration for eschewing the use of the binary lottery is that its success in inducing risk-neutral behavior in practice is decidedly mixed and may in fact be worse than simply using money directly (Selten *et al.*, 1999).

neither the pure strategy Nash nor the Maxmin predictions do especially well in describing subject behavior. We compare these benchmarks with three other hypotheses: random play, in which subjects choose each strategy with equal probability, the mixed Nash equilibrium, and QRE play. These fare much better than the benchmark predictions according to standard scoring rules. Indeed, in all sessions, random play outperforms both pure Nash and Maxmin. In the Baseline game, the mixed-strategy Nash equilibrium does reasonably well according to a Mean Square Deviation criterion, but fares less well in the Upside game. However, this finding is sensitive to the scoring rule (in particular, the extent to which the scoring rule penalizes models for attaching a zero probability to choices that are actually observed). Of the five models considered, overall the best performing is the QRE model. In the remainder of this section, we examine these results in more detail.

4.1. *Comparing Pure Nash and Maxmin Predictions*

Histograms of subject choices for each game and country are presented in Fig. 4. This figure highlights the fact that the country in which a given game is run makes little difference to the distribution of choices, but the game treatment has a marked effect. Consider first the Baseline game. The most frequently observed choice is C and the least frequently observed is A. (Although Fig. 4 pools the data from three sessions for each country, this same pattern is observed in all six sessions when considered separately.) Thus, at least for this game, the Maxmin outperforms the pure Nash strategy as a predictor of behavior. It should be noted, however, that although most choices are C's, and hence are consistent with Maxmin play, there is a substantial amount of B play (33% of choices across all six sessions), which is consistent with a best response to C. Thus, the Baseline game gives only limited support to the Maxmin concept as a description of observed behavior.

Next, consider the Upside game. In contrast to the Baseline game, the most frequently observed choice in the Upside game is A, the pure Nash strategy, and the least frequently observed choice is B. (Again, these same patterns are observed in all six sessions when considered separately.) However, while there is some support for the pure Nash strategy in the Upside sessions, the majority of choices are *not* the pure Nash strategy.

The Maxmin strategy (C) and the best response to the Maxmin strategy (B) are both chosen less frequently in the Upside game than in the Baseline game. Thus, the performance of the Maxmin prediction varies across games in a manner that seems to reflect differences in the riskiness of the Nash strategy as well as the increased incentive to choose strategy B against a Maxmin player.

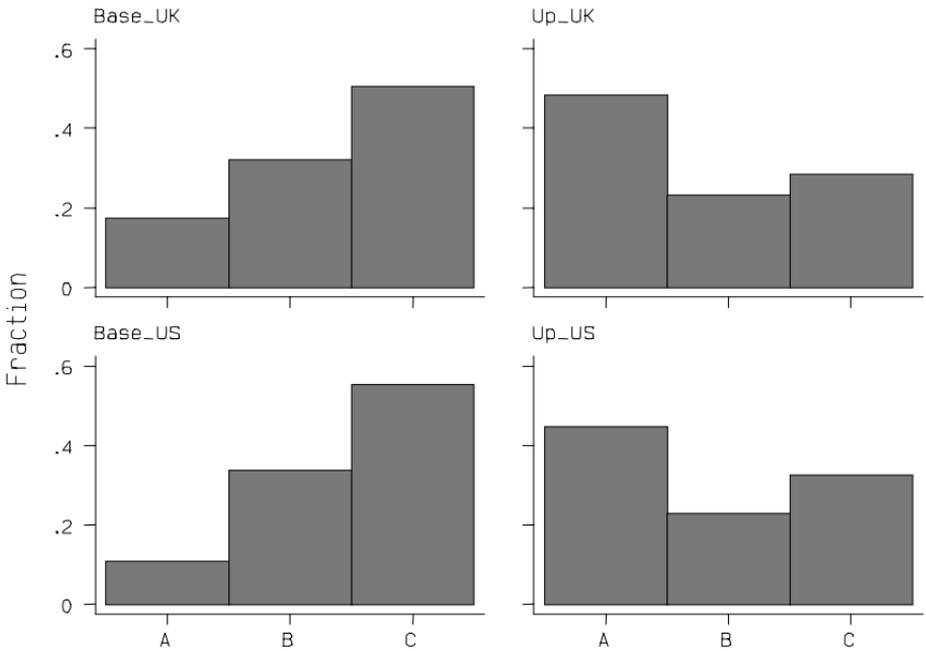


FIG. 4. Histograms of Choices by game and country.

We use two methods to determine whether choices are systematically related to venue or game. First, we treat each session as a single observation and use a permutation test to determine the presence of a venue effect.¹⁷ For both games, regardless of whether we use the proportion of A, B, or C play to be the summary statistic being compared, we fail to reject the null hypothesis that the data generation process is the same across venues. Second, we test for venue effects by using Fisher's exact test of the difference in empirical frequencies of choices A, B, and C with first-round data for a given game treatment.¹⁸ We find no significant difference in the proportions of first-round choices across countries for the Baseline game (p value = .351), or for the Upside game (p value = .110). Next, we examine the presence of game effects by using these same tests. Permutation tests indicate a significant game effect regardless of whether we use the proportions of A, B, or C choices as the summary statistic. Similarly, the

¹⁷Specifically, let Y_s be a summary statistic based on session s data. We conclude that the venue has a significant effect if either $\text{Max}\{Y_{US,1}, Y_{US,2}, Y_{US,3}\} < \text{Min}\{Y_{UK,1}, Y_{UK,2}, Y_{UK,3}\}$ or $\text{Max}\{Y_{UK,1}, Y_{UK,2}, Y_{UK,3}\} < \text{Min}\{Y_{US,1}, Y_{US,2}, Y_{US,3}\}$. If the data generation process is the same across venues, $\text{Pr}\{\text{Max}\{Y_{US,1}, Y_{US,2}, Y_{US,3}\} < \text{Min}\{Y_{UK,1}, Y_{UK,2}, Y_{UK,3}\}\} = 3!3!/6! = 0.05$, this is a 10% significance test.

¹⁸Since first-round choices may be viewed as independent, the Fisher exact test is appropriate.

Fisher exact test of the hypothesis that there is no game treatment effect is clearly rejected (P value = 0.002). Choice proportions in first-round data are given in Table I.

One might wonder whether the patterns of play observed above are stable across rounds. For instance, it is possible that learning effects lead to more Nash play later in the session. To study this, we recomputed the permutation tests described above, using the data from the last 25 rounds of each session. Again, we fail to detect venue effects but find significant game effects.¹⁹ In particular, exactly the same patterns of play are evident in the last 25 rounds as in all rounds taken together.

Although equilibrium and Maxmin payoffs are 40 points in both games, Table II reveals an interesting difference in earnings attained by subjects across the two games. In the Upside game subjects averaged 46.4 points per round, whereas in the Baseline game earnings averaged 38.3 points per round. In fact 53 of the 60 subjects in the Upside game earned more than 40 points per round, and only 7 earned less than 40 points per round. In contrast, only 14 subjects in the Baseline game earned more than 40 points per round, while 40 subjects earned less than 40 points per round. The remaining six subjects employed the Maxmin strategy, choosing C in every round. These subjects earned exactly 40 points, thus exceeding the average payoff among Baseline subjects. This again illustrates the difference in the

¹⁹Details are available from the authors upon request.

TABLE I
Proportions of Choices by Session

	All rounds			First round			Last 25 rounds		
	A	B	C	A	B	C	A	B	C
Baseline									
UK_1	16.0	31.6	52.4	10.0	50.0	40.0	10.4	34.8	54.8
UK_2	25.4	29.2	45.4	10.0	0.0	90.0	22.4	30.0	47.6
UK_3	10.6	35.2	54.2	0.0	30.0	70.0	14.0	34.4	51.6
US_1	11.8	33.8	54.4	10.0	10.0	80.0	9.6	35.2	55.2
US_2	5.0	35.0	60.0	10.0	30.0	60.0	7.2	37.6	55.2
US_3	16.0	32.6	51.4	40.0	40.0	20.0	13.6	33.2	53.2
Upside									
UK_1	42.6	24.0	33.4	20.0	10.0	70.0	52.0	17.6	30.4
UK_2	55.6	22.0	22.4	60.0	0.0	40.0	62.8	17.6	19.6
UK_3	46.8	23.6	29.6	50.0	10.0	40.0	46.8	24.0	29.2
US_1	47.8	22.8	29.4	40.0	30.0	30.0	50.8	21.2	28.0
US_2	36.4	26.8	33.8	50.0	30.0	20.0	47.6	26.4	26.0
US_3	46.8	19.0	34.2	30.0	20.0	50.0	44.4	18.4	37.2

TABLE II
Average Payoff Per Round by Session

Session	Baseline	Upside
UK_1	36.36	46.38
UK_2	37.26	44.00
UK_3	37.94	48.48
US_1	39.66	46.82
US_2	40.34	47.70
US_3	38.40	44.92
ALL	38.33	46.38

riskiness of the two game treatments. In both games subjects understood that choosing C guaranteed them a payoff of 40, but nearly all subjects attempted to earn more than this. In the Baseline games such attempts generally failed, and subjects who played the Maxmin strategy did relatively well. In the Upside game deviations from the Maxmin strategy yielded more than 40 points, so that attempts to improve upon the Maxmin payoff were generally successful.

A possible explanation for payoff differences is that the games differ in their potential for subjects to coordinate on strategies that maximize joint payoffs. Formally, suppose that subjects select symmetric strategies to maximize the sum of payoffs. In this case, the optimal strategy profile in the Baseline game consists of playing B with 20% probability and C with 80% probability. These strategies yield each subject an expected payoff of 42. In contrast, maximizing strategies in the Upside game give 36.4% probability to B and 63.6% to C, which yield each subject an expected payoff of 54.5. In all cases subjects achieved substantially less than the coordination payoffs. The decrease in the empirical frequency of C play in the Upside game relative to Baseline is qualitatively consistent with the predictions of joint profit maximization; however, the decrease in the empirical frequency of B play is not. In neither game is anything close to the empirical frequency of A play predicted. Thus, we find only weak support at best for the notion that subjects are employing collusive strategies in these games.

Overall, the results offer strong support for Aumann's argument that the Nash equilibrium is an implausible prediction in unprofitable games. Indeed, Nash play accounts for less than 50% of all choices in both games. Numerous responses in post-experiment questionnaires indicated that the riskiness of the Nash strategy relative to the safe Maxmin strategy seemed to play a significant role in subjects' thinking when they made their choices. Moreover, our results offer strong support for the hypothesis that subjects are more likely to play the Nash strategies in less risky, unprofitable games.

Moving from the Baseline to the Upside game, subjects are more than three times as likely to play the Nash strategy.

Our results do not support Harsanyi's prediction that subjects will play Maxmin in unprofitable games. Maxmin play accounts for only 53% of choices in the Baseline game and only 30% of choices in the Upside treatment. Furthermore, while subjects worried about the risk of A and B play, they found the upside potential of these strategies too strong to resist in many instances. That is, neither subjects behavior nor their questionnaire responses indicate the degree of pessimism about the opponent's strategies that is required to justify always playing Maxmin. In fact, subjects did hope to earn more than the 40 points associated with Maxmin play and would often take chances on A and B play in the Baseline game and even more so in the Upside game.

The change in the propensity to play Nash strategies is intriguing. Perhaps it is possible that when the Upside is made large enough, play converges to Nash. In particular, if we increase the 60 payoffs in the Baseline game to 6000 (say), we will not affect the pure Nash or Maxmin strategies, but we will massively increase the degree of weight subjects must place on unexpected C play to justify their switching away from Nash. We suspect that subjects would seldom find it desirable to forgo the "risk" associated with playing Nash for the "safety" of playing Maxmin.

Since we find the pure Nash equilibrium prediction wanting, but at the same time reject the notion that subjects play their Maxmin strategy instead, in the next subsection we focus on measuring the applicability of these solution concepts relative to some alternative models.

4.2. *Assessing Models of Behavior*

In assessing the performance of the pure Nash and Maxmin hypotheses, we compare them with three alternative models of subject behavior. The first alternative is the symmetric mixed Nash equilibrium, $(p^A, p^B, p^C) = (0, .4, .6)$ and $(p^A, p^B, p^C) = (55/88, 9/88, 24/88)$ for the Baseline and Upside games, respectively. The second alternative considered is denoted the "Random Play" model consisting of strategies $(p^A, p^B, p^C) = (1/3, 1/3, 1/3)$. The final model considered is a QRE model. In this model all subjects choose the j th choice with probability

$$p^j = \frac{\exp\{\lambda\pi^j\}}{\sum_{k \in \{A, B, C\}} \exp\{\lambda\pi^k\}},$$

where π^j is the expected payoff from the j th choice and is therefore a function of probability assessments about opponents' behavior. The quantal response equilibrium condition sets choice probabilities equal to these probability assessments, thus providing a set of probabilities for a given

TABLE III
Maximum Likelihood Estimates of QRE Models for Baseline Game

Session	λ	$se(\lambda)$	p^A	p^B	$-\ell$
UK.1	1.017	0.122	0.200	0.300	358.848
UK.2	0.603	0.129	0.278	0.278	376.886
UK.3	1.932	0.191	0.089	0.357	318.319
US.1	1.369	0.143	0.147	0.325	342.112
US.2	2.611	0.271	0.048	0.382	282.391
US.3	1.126	0.129	0.182	0.308	358.610
ALL	1.297	0.056	0.156	0.320	2083.512

value of λ .²⁰ The parameter λ can be interpreted as a sensitivity parameter: when $\lambda = 0$ all choices are equally probable, as λ increases more probability weight is assigned to those choices that give a higher expected payoff, and as λ approaches infinity the probability with which the expected payoff maximizing choice is made approaches one.

Since the QRE prediction is based on the subject data whereas the other models are not, we follow Camerer and Ho (1999) by estimating the model using the first 70% of the data, reserving the last 30% of the data for out-of-sample prediction.²¹ The maximum likelihood estimates for the QRE model, estimated from the first 35 rounds of data, are presented in Tables III and IV.

²⁰In fact in our games there are sometimes two solutions to the set of equations, each describing a path that converges (as $\lambda \rightarrow \infty$) to a symmetric Nash equilibrium. We restrict our attention to the path that converges to the mixed Nash equilibrium.

²¹An alternative employed by Chen and Tang (1998) is to use all observations to calibrate and validate the models. Since all of the models with the exception of Maxmin are special cases of QRE, it will always fit the data better. Using separate sets of observations to calibrate and validate the models avoids this problem.

TABLE IV
Maximum Likelihood Estimates of QRE Models for Upside Game

Session	λ	$se(\lambda)$	p^A	p^B	$-\ell$
UK.1	0.134	0.085	0.378	0.318	383.422
UK.2	0.658	0.143	0.515	0.228	357.680
UK.3	0.526	0.128	0.491	0.247	368.034
US.1	0.491	0.124	0.483	0.253	370.440
US.2	0.000	0.002	0.333	0.333	384.515
US.3	0.966	0.270	0.553	0.197	355.473
ALL	0.416	0.045	0.465	0.265	2242.253

To compare the predictive performance of these models we use the mean square deviation (MSD) scoring rule applied to the last 15 rounds of choices:

$$\text{MSD} = \frac{1}{30S} \sum_{t=36}^{50} \sum_{i=1}^S \sum_{j \in \{A,B,C\}} \left(d_i^j(t) - p_i^j(t) \right)^2,$$

where S is the number of subjects, $d_i^j(t)$ is a dummy variable indicating whether subject i chose j in round t , and $p_i^j(t)$ is the predicted probability with which subject i chose j in round t . Perfect predictions yield a score of $\text{MSD} = 0$, and the worst possible score is $\text{MSD} = 1$.

While widely accepted as a reasonable measure of the predictive power of models, the MSD measure has several limitations. For example, the MSD criterion favors models that “smear” probability. Suppose that we observe A in a given period. Then, a model that predicts $(.59, .205, .205)$ will outperform a model predicting $(.6, .4, 0)$ by this criterion. Given that the former model actually placed less weight on the observed outcome, it seems peculiar that, by allocating probability more evenly over unobserved choices, it receives a better score than the latter model.²² Thus, we also report minus the out-of-sample log-likelihood as another measure of predictive success (where, again, lower values indicate better predictive performance):

$$-\ell_{35+} = \sum_{t=36}^{50} \sum_{i=1}^S \sum_{j \in \{A,B,C\}} d_i^j(t) \ln(p_i^j(t)).$$

According to the MSD criterion, Maxmin predicts better than pure Nash in the Baseline sessions (Table V) and vice versa for the Upside sessions (Table VI). In this respect, the MSD scoring rule quantifies the information contained in Fig. 4 rather well. By comparison, in every session all of the other models presented in Tables V and VI achieve better MSD scores than either of the pure strategy predictions. There is no clear MSD ranking among the alternative models at the session level. In four of the six Baseline sessions QRE does best, while the mixed Nash does best in the other two sessions. Of the Upside sessions QRE does best in three, mixed Nash in two, and Random Play in the last.

To get a better sense of the overall performance of the alternative models we also computed the MSD scores after pooling the data from all sessions for a given game. Then the ranking of the models for the Baseline game is

$$\text{QRE} \succ_{\text{MSD}} \text{Mixed Nash} \succ_{\text{MSD}} \text{Random} \succ_{\text{MSD}} \text{Maxmin} \succ_{\text{MSD}} \text{Pure Nash},$$

²²We thank an anonymous referee for bringing this point to our attention.

TABLE V
 Predictive Success of Models for Baseline Game

Session	Maxmin		Pure Nash		Mixed Nash		Random Play		QRE	
	MSD	$-\ell_{35+}$	MSD	$-\ell_{35+}$	MSD	$-\ell_{35+}$	MSD	$-\ell_{35+}$	MSD	$-\ell_{35+}$
UK_1	0.440	—	0.933	—	0.273	—	0.333	164.792	0.283	141.714
UK_2	0.500	—	0.793	—	0.343	—	0.333	164.792	0.313	156.905
UK_3	0.487	—	0.853	—	0.317	—	0.333	164.792	0.303	151.227
US_1	0.447	—	0.947	—	0.270	—	0.333	164.792	0.277	134.659
US_2	0.460	—	0.940	—	0.277	—	0.333	164.792	0.273	130.601
US_3	0.440	—	0.893	—	0.290	—	0.333	164.792	0.287	142.704
ALL	0.463	—	0.893	—	0.297	—	0.333	988.751	0.290	855.769

where the relation \succ_{MSD} has the obvious meaning. For the Upside game, we have

$$\text{QRE} \succ_{\text{MSD}} \text{Mixed Nash} \succ_{\text{MSD}} \text{Random} \succ_{\text{MSD}} \text{Pure Nash} \succ_{\text{MSD}} \text{Maxmin}.$$

Thus, the QRE outperforms the other models in both games. In terms of the log-likelihood-based scoring rule, for the Baseline sessions, QRE outperforms the Random Play model, which in turn outperforms the other models. In the Upside sessions there is little to choose between Random Play and Mixed Nash, but overall the QRE model does best, giving the highest score in four of the six sessions. The out-of-sample log-likelihood is infinite if a choice is observed that is predicted to occur with zero probability. As indicated in Tables V and VI, this happens in the cases of the Nash and Maxmin predictions for all sessions, and in the case of the Mixed Nash prediction for all sessions of the Baseline game. Thus, the log-likelihood rule does not discriminate between these models for these sessions.

Overall, the tests confirm the implausibility of the Nash prediction in unprofitable games. However, the Maxmin solution concept fares poorly;

TABLE VI
 Predictive Success of Models for Upside Game

Session	Maxmin		Pure Nash		Mixed Nash		Random play		QRE	
	MSD	$-\ell_{35+}$	MSD	$-\ell_{35+}$	MSD	$-\ell_{35+}$	MSD	$-\ell_{35+}$	MSD	$-\ell_{35+}$
UK_1	0.687	—	0.467	—	0.303	151.109	0.333	164.792	0.323	160.144
UK_2	0.800	—	0.373	—	0.273	142.442	0.333	164.792	0.277	141.576
UK_3	0.687	—	0.560	—	0.353	176.451	0.333	164.792	0.327	161.638
US_1	0.727	—	0.500	—	0.327	166.045	0.333	164.792	0.313	155.913
US_2	0.760	—	0.487	—	0.327	167.329	0.333	164.792	0.333	164.834
US_3	0.567	—	0.653	—	0.380	184.137	0.333	164.792	0.360	174.524
ALL	0.703	—	0.507	—	0.327	987.512	0.333	988.751	0.313	940.586

thus there is little support for the Maxmin postulate. Of the models considered, QRE does best overall, while the two pure strategy predictions, pure Nash and Maxmin, do worst.

5. CONCLUSION

Our results strongly support Aumann's argument that the Nash equilibrium prediction is implausible in unprofitable games. Indeed, by the MSD scoring rule, the pure Nash prediction in the Baseline game scores below every alternative model, including a model where subjects choose strategies at random. In the Upside game, the pure Nash prediction only outperforms Maxmin. However, our results do not support Harsanyi's Maxmin Postulate. By our scoring rule, every alternative model outperforms Maxmin in the Upside game. In the Baseline game, Maxmin outperforms only the pure Nash prediction.

Our results also strongly support the hypothesis that the Nash prediction improves when Nash play is less risky. By our scoring rule, the pure Nash prediction performs significantly better in the less risky Upside game than it does in the Baseline game. Of course, with only two treatments, one can only guess whether, with the upside increased sufficiently, play would ultimately converge to the Nash prediction. In future research, we intend to explore this further by examining a variety of unprofitable game treatments that vary the potential upside associated with Nash play.

APPENDIX A: PROOFS

PROPOSITION 1. *There does not exist any 2×2 game with distinct Nash and Maxmin profiles in pure strategies.*

Proof. Consider the game described below.

	LEFT	RIGHT
UP	1, 1	c, b
DOWN	a, d	e, f

Without loss of generality, we suppose that the Nash equilibrium is UP, LEFT (with payoffs normalized to 1), so that $a \leq 1$ and $b \leq 1$. If $c \geq e$ then UP is a weakly dominant (and therefore also a Maxmin) strategy. So suppose $c < e$. Then for DOWN, RIGHT not to be an additional equilibrium, we must have $d > f$. But then LEFT is a weakly dominant (and Maxmin) strategy. Hence either UP or LEFT is a weakly dominant strategy and a Maxmin strategy.

PROPOSITION 2. *There does not exist a symmetric 3×3 unprofitable game with two levels of payoff and distinct Nash and Maxmin strategies.*

Proof. Consider the payoffs to the row player represented as a 3×3 binary matrix. Such a matrix must have the following properties to be a unprofitable game:

1. No rows consisting entirely of 0's.
2. No pure strategy equilibria consisting of payoffs equal to 1 for each player. If so, then the Maxmin strategies must consist of a row of 1's, but then this would be an equilibrium strategy, violating the requirement that Nash and Maxmin strategies be distinct. This implies
 - a. No 1's on the principal diagonal.
 - b. If the payoff to the row player from the strategy pair (x, y) is 1, then the payoff to the strategy pair (y, x) must be 0.

With these conditions, any 3×3 two outcome symmetric game may be represented by the matrix

A 3×3 Two-Outcome Game		
0	a	b
$1 - a$	0	c
$1 - b$	$1 - c$	0

There are two cases to consider: If $a = 1$, then the above facts imply $c = 1$ and $b = 0$. This reduces the game to rock-scissors-paper, where Nash and Maxmin strategies coincide. Likewise, if $a = 0$, the above facts imply $b = 1$ and $c = 0$, and again the game is rock-scissors-paper. ■

PROPOSITION 3. *Every symmetric 3×3 unprofitable game with distinct Nash and Maxmin profiles in pure strategies has a symmetric Nash equilibrium in mixed strategies.*

Proof. Consider the set of 3×3 symmetric unprofitable games with a unique pure strategy Nash equilibrium and a different unique pure Maxmin strategy. Let the pure strategy equilibrium be (A, A) , let the Maxmin strategy be C , and normalize the equilibrium payoff strategy to zero. Then we may represent any such game by the payoff matrix where $b \leq 0$ (otherwise (A, A) is not an equilibrium) and $\min\{a, c\} < 0$ (otherwise A is the Maxmin strategy):

	A	B	C
A	0,0	c, b	$a, 0$
B	b, c	d, d	$e, 0$
C	$0, a$	$0, e$	0,0

For (A, C) and (C, C) not to be additional equilibria we must have $e > \max\{a, 0\}$. Then, for (C, B) and (B, B) not to be additional equilibria we need $c > \max\{d, 0\}$. Since this implies $c > 0$, from the earlier inequality $\min\{a, c\} < 0$ we must have $a < 0$. Finally, for (B, A) not to be an equilibrium we need $b < 0$. In summary, the following constraints must hold: $c > 0$, $c > d$, $e > 0 > a$, and $b < 0$.

Case 1. $ec \geq ad$. Then consider the mixture $\sigma = (p, q, 1 - p - q)$, where

$$p = \frac{ec - ad}{(ec - ad) + ab - bc}, \quad q = \frac{ab}{(ec - ad) + ab - bc}.$$

Since $ec - ad \geq 0$, $ab > 0$, and $-bc \geq 0$, this mixture lies in the unit simplex and is a feasible mixed strategy. Against this strategy it is easily verified that A gives an expected payoff of 0, as does B and C. Thus any mixed strategy is a best response to σ and thus σ is a best response to itself. Hence σ is a full-support Nash equilibrium.

Case 2. $ec \leq ad$. (Note that since $ec > 0$ and $a < 0$ this inequality implies $d < 0$.) Then consider the mixture $\sigma' = (0, e/(e - d), d - /(e - d))$. Since $e > 0$ and $-d > 0$ this mixture lies in the unit simplex and is a feasible mixed strategy. Against this strategy the expected payoff to A is $(ec - ad)/(e - d) < 0$. The expected payoff to B or C against this strategy is 0. Thus any mixture over B and C is a best response to σ' , and thus σ' is a best response to itself. Hence σ' is a mixed-strategy Nash equilibrium.

APPENDIX B: INSTRUCTIONS

General Rules

This is an experiment in the economics of decision making. If you follow the instructions carefully and make good decisions you can earn a considerable amount of money. You will be paid in private and in cash at the end of the experiment.

There are ten people participating in this experiment. These instructions apply equally to all ten participants. It is important that you do not talk, or in any way try to communicate, with other people during the experiment. If you have a question, raise your hand and a monitor will come over to where you are sitting and answer your question in private.

The experiment will consist of 50 rounds. In each round you will be randomly matched with another person in the room. The matchings will change from round to round and you will not know with whom you are matched in any round.

In each round you will have an opportunity to earn points. At the end of the experiment, you will be paid an amount in cash that will be determined by the total number of points you earn from all rounds.

Description of Each Round

At the beginning of the first round you will see a screen like the one below:

```

ROUND 1      MY POINTS TO BEGINNING OF THIS ROUND 0

              MY CHOICE IS ___
              (SELECT A, B, OR C)

PRESS ENTER WHEN YOU ARE SATISFIED WITH YOUR CHOICE

```

The top line tells you the round number and the total number of points you have accumulated up to the beginning of the round.

You will make a decision by typing in a choice of A, B, or C. Until you press the Enter button you are free to change your selection as often as you like simply by typing in a different choice. When you are satisfied with your choice you will press the Enter key. Once you press the Enter key you will have made your decision for the round and it cannot be changed.

When all ten people have made their decisions you will see a screen displaying your choice, the choice of the person with whom you were matched, and your point earnings for the round. A sample screen is shown below (all entries in the sample screen are for illustrative purposes only):

```

              MY POINTS TO BEGINNING OF ROUND 1      0

You chose                A
Person with whom you were matched chose      B
You earned                120

              MY POINTS AT END OF ROUND 1          120

              Press space bar to continue

```

Your point earnings for the round will depend on your choice and the choice made by the person with whom you are matched. (Remember, the person with whom you are matched will change from round to round.) Specifically, your point earnings will be calculated according to the table below, which gives your point earnings for each possible choice combination. For example, if you choose B and the person with whom you are matched chooses A, you will earn 10 points. Notice that in this example, the person with whom you were matched earned 120 points, since this person chose A and was matched with you (who chose B).

		Person with Whom You Are Matched Chooses		
		A	B	C
You Choose	A	40	120	10
	B	10	10	120
	C	40	40	40

When you have read the screen displaying your point earnings you will press the space bar. This will conclude round one and you will go on to round two.

At the beginning of all subsequent rounds, you will see a screen like the one below (again, the numbers in the sample screen are only for illustrative purposes):

ROUND 2 MY POINTS TO BEGINNING OF THIS ROUND 120

Information from Previous Round			
Choice	A	B	C
Number	2	5	4
Average earnings	65	54	40

MY CHOICE IS ____
(SELECT A, B, OR C)

PRESS ENTER WHEN YOU ARE SATISFIED WITH YOUR CHOICE

This screen contains information about the results of the previous round. The line beginning with "Number" lists the number of participants choosing A, B, or C. The line beginning with "Average earnings" lists the average number of points earned by participants choosing A, B, or C.

The line beginning with "MY CHOICE IS ____" is for you to type in a choice. You should select A, B, or C. Until you press the Enter key, you are free to change your selection as often as you like by typing in a different choice. When you are satisfied with your choice you will press the Enter key. Once you have pressed the Enter key you will have made your decision for the round and it cannot be changed.

When all ten participants have made their decisions you will be informed of your point earnings for this round. Your point earnings will be calculated and displayed in the same way as for round one.

Ending the Experiment

At the end of round 50 you will be paid, in private and in cash, an amount determined by the total number of points you accumulated over all 50 rounds. You will be paid 25 cents for every 40 points earned.

If you have any questions raise your hand.

Quiz

If you chose C and the person with whom you are matched in that round chose A:

1. You would earn ____ points.
2. The person with whom you are matched would earn ____ points.

If you chose B and the person with whom you are matched in that round chose B as well:

3. You would earn ____ points.
4. The person with whom you are matched would earn ____ points.
5. If at the end of 50 rounds you earned 2,500 points, you would receive a payment of \$ ____.

REFERENCES

- Aumann, R. J. (1985). "On the Non-Transferable Utility Value: A Comment on the Roth-Shaffer Examples," *Econometrica* **53**, 667-677.
- Aumann, R. J., and Maschler, M. (1972). "Some Thoughts on the Minimax Principle," *Management Sci.* **18**, 54-63.
- Bernheim, D. (1984). "Rationalizable Strategic Behavior," *Econometrica* **52**, 1007-1028.
- Brown, J., and Rosenthal, R. (1990). "Testing the Minimax Hypothesis: A Reexamination of O'Neill's Game Experiment," *Econometrica* **58**, 1065-1081.
- Camerer, C., and Ho, T.-H. (1999). "Experience-Weighted Attraction Learning in Normal Form Games," *Econometrica* **67**, 827-874.
- Chen, Y., and Tang, F.-F. (1998). "Learning and Incentive-Compatible Mechanisms for Public Goods Provision: An Experimental Study," *J. Political Econ.* **106**, 633-662.
- Gilboa, I., and Schmeidler, D. (1989). "Maxmin Expected Utility with Non-unique Prior," *J. Math. Econom.* **18**, 141-153.
- Harsanyi, J. C. (1964). "A General Solution for Finite Non-Cooperative Games, Based on Risk-Dominance," in *Advances in Game Theory* (M. Dresher, L. S. Shapley, and A.W. Tucker, Eds.), Princeton, NJ: Princeton University Press.
- Harsanyi, J. C. (1966). "A General Theory of Rational Behavior in Game Situations," *Econometrica* **34**, 613-634.
- Harsanyi, J. C. (1977). *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*, Cambridge, U.K.: Cambridge University Press.
- Holler, M., and Host, V. (1990). "Maximin vs. Nash Equilibrium: Theoretical Results and Empirical Evidence," in *Optimal Decisions in Markets and Planned Economies* (R. Quandt and D. Triska, Eds.), pp. 245-255, Boulder, CO: Westview Press.
- Lo, K. C. (1996). "Equilibrium in Beliefs under Uncertainty," *J. Econom. Theory* **71**, 443-484.
- McKelvey, R., and Palfrey, T. (1995). "Quantal Response Equilibria for Normal Form Games," *Games Econom. Behav.* **10**, 6-38.
- Ochs, J. (1995). "Games with Unique Mixed Strategy Equilibria: An Experimental Study," *Games Econom. Behav.* **10**, 174-189.
- O'Neill, B. (1987). "Nonmetric Test of the Minimax Theory of Two-Person Zero-Sum Games," *Proc. Natl. Acad. Sci. USA* **84**, 2106-2109.
- Osborne, M., and Rubinstein, A. (1994). *A Course in Game Theory*, Cambridge, MA: MIT Press.
- Pearce, D. (1984). "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica* **52**, 1029-1050.
- Rubinstein, A. (1991). "Comments on the Interpretation of Game Theory," *Econometrica* **59**, 909-924.
- Selten, R., Sadrieh, A., and Abbink, K. (1999). "Money Does Not Induce Risk Neutral Behavior, but Binary Lotteries Do Even Worse," *Theory Decision* **46**, 213-252.
- Shachat, J. (1996). "Mixed Strategy Play and the Minimax Hypothesis," mimeo. van Damme, E. E. C. (1980). "Some Comments on Harsanyi's Postulates for Rational Behavior in Game Situations," *Methods Oper. Res.* **38**, 189-205.