An Experimental Investigation of Unprofitable Games*

by

John Morgan

Woodrow Wilson School, Princeton University

and

Martin Sefton

Department of Economics, University of Newcastle

Abstract

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 relative performance of the solution concepts is sensitive to the riskiness of the Nash strategy. We compare the predictive performance of several alternative models including dynamic models that account for subjects' experiences from previous rounds do even better. In particular, an experience weighted attraction model captures the heterogeneity of strategies and outperforms all other models according to the standard scoring rules we employ.

Keywords: Unprofitable games, Nash Equilibrium, Maxmin, Learning models.

* We are grateful to participants at the ESA conferences in Mannheim, 1998, and Tucson, 1998, 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. Address all correspondence to Morgan at Robertson Hall, Princeton, NJ 08544. E-mail: rjmorgan@princeton.edu.

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 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. [1] Aumann (1985) offers the following example:

Figure 1. Aumann's Example		
	LEFT	RIGHT
UP	2, 6	4, 2
DOWN	6, 0	0, 4

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 ¾ 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

[1] Other solution concepts, for example Rationalizability and Cautious Rationalizability (Pearce, 1984 and 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.

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 (¾, ¼) available, which assures him of 3 regardless; similarly player 2 has strategy (¼, ¾). 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 Figure 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 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 that:

> 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.[2] 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 as to 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 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…")

[2] In 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.

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 using Lo's (1996) game theoretic adaptation of the multiple prior 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 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 3x3 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 more than 2/3rds 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? It appears to us that this contention is rejected. 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.

In light of the fact that the criticisms of the Nash equilibrium prediction seem to be borne out in the data, we ask whether other models of subject play offer better predictions. We estimate the quantal response equilibrium (QRE) model of McKelvey and Palfrey (1995) and find that, in the class of static models considered, the QRE outperforms all other models in predictive power, according to standard scoring rules applied to out-of-sample predictions. We also consider the performance of learning models for capturing dynamic patterns in the data. To this end, we estimate the experience-weighted attraction (EWA) model of Camerer and Ho (1999). This model includes both choice-reinforcement (Roth and Erev, 1995) and belief-based learning models (Fudenberg and Levine, 1998) as special cases. The dynamic EWA model outperforms all the static models by successfully predicting the heterogeneity of strategies observed in the data. Of all the models we examine, the two *worst* at predicting behavior are the pure Nash and the Maxmin hypotheses.

To 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 2x2 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 using 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. Our experiments differ from the extant literature along several dimensions, two of which are especially noteworthy. 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.

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 static models, in section 4. In Section 5 we estimate the EWA model and measure its predictive power to that of the static models. Section 6 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 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 3x3 games with 3 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.

Proposition 1: There does not exist any 2x2 game with distinct Nash and Maxmin profiles in pure strategies.[3]

Thus, we are forced to have three or more pure strategies to meet our criteria. The following proposition shows that with symmetric 3x3 games, we must have more than two payoff levels as well.

Proposition 2. There does not exist a symmetric 3x3 unprofitable game with two payoff levels and distinct Nash and Maxmin strategies.

Our Baseline Game, shown in Figure 2, establishes that there is a 3x3 unprofitable game with three payoff levels that does meet our criteria.

Figure 2. Baseline Game

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

[3] See Appendix A for a proof of this and other propositions in this Section.

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. 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 player with probability .4 and C with probability .6 as well as two asymmetric mixed strategy equilibria.[4]

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 3x3 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 3x3 unprofitable game with distinct Nash and Maxmin profiles in pure strategies has a symmetric Nash equilibrium in mixed strategies.

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 to over a set of pure strategies that includes the Maxmin strategy.[5] It is hard to see why such a randomization would be willingly

[4] In 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.

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

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.[6] Further, even in laboratory experiments where a unique mixed strategy equilibrium corresponded with Maxmin play, these predictions have received, at best, limited support.[7] 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 Figure 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.[8]

Figure 3. Upside Game

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

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 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

[6] Rubinstein (1991) presents a wide-ranging discussion of these issues.
[7] See for instance Brown and Rosenthal (1990).
[8] In 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.

for a given normal form game. Let $g_i{:}S \to X$ be the outcome function of the game for player *i*. Let $u_i{:}M(X) \to R$ 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

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

This equation says that a strategy σ_i is a best response for player *i* if she can do no better than selecting σ_i given that whatever strategy she chooses, her opponent selects a strategy $p_i \in B_i$ that minimizes *i*'s payoffs. With this definition in mind, the following is immediate:

Remark 1: Suppose that $B_i = M(S_j)$, then a best response for *i* is a Maxmin strategy.

Remark 1 states that if a player's priors are completely diffuse, then she can do no better than to play a Maxmin strategy.

To assess the relative riskiness of Nash strategies in the Baseline versus the Upside game using 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

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

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

For the Upside Game,

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

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 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 be 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. Thus, once again, we have that the Baseline game is riskier than the Upside game.

3. Experimental Design and Procedures

The experiment consists of 12 sessions, six sessions conducted at the University of Newcastle in spring 1998 and six at Penn State University 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 US and UK sessions allow us to investigate robustness of results across subject pools, since as far as was possible the same procedures were used in each country. For the Newcastle 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 Penn State 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 (UK) or 25¢ per 40 points (US).[9]

The following procedures were common to all sessions. At the beginning of a session, the ten 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.[10]

The session then consisted of fifty 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 payoff matrices described in the previous section. At the end of each round subjects 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 less than 50 minutes and average earnings were £7.44 (UK Baseline sessions), £9.26 (UK Upside), $12.33 (US Baseline) and $14.53 (US Upside). This is considerably more than outside earnings opportunities of the subjects. In fact, subjects completed a short post-experimental questionnaire and the question

[9] At the time of the experiments, £1 = $1.65 so that the stakes are quite comparable.
[10] Appendix B contains copies of the instructions and quiz.

"Would you be willing to take part in other experiments of this sort?" received 119 out of 119 affirmative responses.[11]

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 attention to 3x3 symmetric games with three levels of payoff. We regarded this to be 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.[12]

4. Results

In this Section, we examine aggregate behavior for evidence of venue or game treatment effects. Our primary 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

[11] There was one non-respondent.
[12] An 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.*, 1995).

well the various theoretical predictions of the Baseline and Upside game do in describing subject behavior. We find that neither the pure strategy Nash nor the Maxmin predictions do especially well in describing subject behavior. We compare these benchmarks to three other static 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 out-performs both pure Nash and Maxmin. In the Baseline game, the mixed strategy Nash equilibrium does reasonably well according to a Mean Square Deviation criteria, 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 static models considered, overall the best performing is the QRE model. In the remainder of the section, we shall examine these results in more detail.

4.1 Comparing pure Nash and Maxmin predictions

Histograms of subject choices in each of the six Baseline sessions are presented in Figure 4. Although this figure shows some variability in choices *across* sessions (for example the proportion of A's varies between about 5% and about 25%), similar patterns are observed in all sessions. That is, in all Baseline sessions the most frequently observed choice is C and the least frequently observed is A. Thus, at least for this game, it seems that the Maxmin outperforms the pure Nash strategy as a predictor of behavior. It should be noted, however, that although most choices are C's, hence consistent with Maxmin play, there is a substantial amount of B play, 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.

Perhaps subject choices move toward one of the pure strategies as rounds progress, and data from later rounds would be more consistent with Nash or Maxmin play? This turns out not to be the case. Figure 5 plots the proportions of A and C choices by round, and there is no clear trend. For the Baseline game then, the pure strategy hypotheses do not provide good predictors of choices at the session level, nor is there any indication that their performances improve as subjects gain experience in the game.

It is useful to contrast the behavior in the Baseline game with that observed in the Upside game. Figure 6 displays the proportions of subject choices in each of the six Upside sessions. As in the Baseline game, there is some variability in choices *across* sessions, but similar patterns are observed in all sessions. In contrast to the Baseline game, the most frequently observed choice in all the Upside sessions game is A, the pure Nash strategy. The Maxmin strategy (C) and the best response to the Maxmin strategy (B) are both chosen less frequently in all of the Upside sessions than in any of the Baseline sessions. 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. While there is some support for the pure Nash strategy in the Upside sessions, the majority of choices are *not* the pure Nash strategy. As in the Baseline sessions, there is no evidence of convergence to either the Nash or the Maxmin strategies in later rounds (see Figure 7).

As we noted earlier, there is no difference between the Maxmin prediction in the Upside game versus the Baseline game. This is likewise true for the pure Nash equilibrium. Thus, neither the pure Nash nor Maxmin concepts predict differences between our two treatments. Nonetheless, we argued previously that the riskiness of the Nash equilibrium varies across the games and might lead to differences in behavior in the two games. The histograms appear to indicate such differences. In particular, the proportion of A's is lower in the Baseline than in the Upside game, whereas the reverse is true of the proportion of B's and C's. Variation in choices across venues is much less marked.

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.[13] 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 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.[14] 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 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. Choice proportions in first-round data are given in Table 1. The Fisher exact test of the hypothesis that there is no game treatment effect rejects at the 1% level (p-value = .002).

Although equilibrium and Maxmin payoffs are 40 points in both games, Table 2 reveals an interesting difference in the 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

---

[13] Specifically, let Y_s be a summary statistic based on session *s* data. We conclude 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 $\Pr\left\{\text{Max}\{Y_{US\_1}, Y_{US\_2}, Y_{US\_3}\} < \text{Min}\{Y_{UK\_1}, Y_{UK\_2}, Y_{UK\_3}\}\right\} = 3!3!/6! = 0.05$, so this is a 10% significance test.

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

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 6 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 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 case 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 making 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. Further, 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 subject behavior, nor their questionnaire responses, indicate the degree of pessimism about 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 by making the Upside 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 forego 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 sub-section we focus on measuring the applicability of these solution concepts relative to some alternative models.

4.2 Assessing static models of behavior

In assessing the performance of the pure Nash and Maxmin hypotheses, we compare them with three alternative *static* 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 static 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 value of λ.[15] The parameter λ can be interpreted as a sensitivity parameter: when $\lambda=0$ each choice is 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.

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

Since the QRE prediction is based on the subject data whereas the other models are not, we follow Camerer and Ho by estimating the model using the first 70% of the data, reserving the last 30% of the data for out-of-sample prediction.[16] The maximum likelihood estimates for the QRE model, estimated from the first 35 rounds of data, are presented in Tables 3 and 4.

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

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

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 $MSD = 0$, while the worst possible score is $MSD = 1$.

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.[17] 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):

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

[16] An alternative employed in Chen and Tang (1998) is to use all observations to calibrate and validate the models. Since all of the static 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.

[17] We thank a reader for bringing this point to our attention.

According to the *MSD* criterion, Maxmin predicts better than pure Nash in the Baseline sessions (Table 5), and *vice versa* for the Upside sessions (Table 6). In this respect, the *MSD* scoring rule quantifies the information contained in Figures 4 and 6 rather well. By comparison, in every session all of the other models presented in Tables 5 and 6 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:

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

where the relation "$\succ_{\text{MSD}}$" 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 4 of the 6 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 5 and 6, 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; thus there is little support for the Maxmin postulate. Of the models considered, QRE does best overall; however, a shortcoming of all of these models is that they impose homogeneity across both individuals and rounds.[18] While Figures 5 and 7 reveal no discernible trends in the data, there is considerable heterogeneity in individual choices.

Figures 8 and 9 summarize the distribution of choices across subjects in the Baseline and Upside sessions respectively. Each point plotted in a panel represents an individual subject's choices averaged over fifty rounds. The (x, y) coordinates give the proportions of A's and B's chosen respectively, so that (1,0) represents 100% A choices, (0,1) represents 100% B choices and the origin represents 100% C choices. Notice that the main source of heterogeneity in the Baseline sessions is in subjects' propensities to play B versus C. To see this note that 20% of the subjects never chose A at all, while only three subjects chose A more than 40% of the time, and none chose A more than 60% of the time. In contrast, the breakdown between B and C choices across individuals shows greater variability across subjects. 27 subjects chose C more than 60% of the time, and only one subject never chose C at all. Alternatively, we may summarize this difference in terms of inter-quartile ranges for the proportions of each choice. The proportion of A's has an inter-quartile range of 19%, whereas the ranges for B and C are much larger: 35% and 49%, respectively. We interpret this as reflecting different responses to the tension between the cautious Maxmin choice and the best response to this cautious choice - perhaps as a result of differing risk preferences. In the Upside sessions there is also considerable heterogeneity: the

[18] In principle the QRE model can be generalized to allow the sensitivity parameter to vary across rounds and/or individuals.

inter-quartile ranges are 24%, 21%, and 32% for the proportions of A, B, and C choices respectively.

There is considerable evidence that dynamic models generally outperform their static counterparts in predicting subject behavior. In the next section, we consider dynamic approaches for modeling behavior.

5 Dynamic Models

We now turn to dynamic explanations of subject behavior based on the EWA model. This model specifies that the probability with which a subject chooses a given strategy depend upon that strategy's "attractiveness" relative to all other strategies. Formally, the probability with which subject *i* chooses strategy *j* in round *t* is given by

$$p_i^j(t) = \frac{\exp\left\{\lambda A_i^j(t-1)\right\}}{\sum_{k\in\{A,B,C\}} \exp\left\{\lambda A_i^k(t-1)\right\}},$$

where $A_i^j(t)$ is the attractiveness of strategy *j* to subject *i* at the end of round *t* and λ is the sensitivity parameter as in the QRE model.

Subjects begin with initial attractions to each strategy, and these attractions then change as subjects observe the results from playing particular strategies in prior rounds. In the EWA model the attractions of each strategy reflects the payoffs subjects actually received against the opponent from the previous round, and the payoffs hypothetical choices *could* have received against the opponent from the previous round. We extend the EWA model to also allow the updating of attractions to reflect the payoffs a subject could have received had their actual choice been matched with hypothetical opponents, and had they made hypothetical choices and been matched with hypothetical opponents.[19]

We use the notation $\pi(j,s)$ to denote the payoff to a player who chooses j in response to s, $s_k(t)$ to denote the choice of subject k in round t and $s_{m(i,t)}(t)$ to denote the choice of subject i's opponent in round t. Then, in our generalized EWA model, the attractions of each strategy are given by

$$A_i^j(t) = \frac{\phi N(t-1)A_i^j(t-1) + \gamma[\delta + (1-\delta)d_i^j(t)]\pi(j, s_{m(i,t)}(t)) + (1-\gamma)[\delta + (1-\delta)d_i^j(t)]\frac{1}{S-1}\sum_{k \neq i}\pi(j, s_k(t))}{N(t)}$$

where

$$N(t) = \rho N(t-1) + 1, t \geq 1,$$

is a sequence of experience weights where more recent events are given greater weight. Notice that when $\gamma = 1$ subjects do not use information on how they may have fared had they been matched with a different opponent when they update attractions. In this case, our model reduces to the version of the EWA model used in Camerer and Ho. If we impose the additional restriction $\delta = 0$ (so that only the attraction associated with the actual strategy used is updated) then the model reduces to a simple choice-reinforcement model. If, on the other hand, γ is unrestricted and $\delta = 1$, the EWA model reduces to weighted fictitious play. Thus, the EWA model encompasses two of the more commonly used learning models as special cases.[20]

---

[19] This extension is irrelevant for the experiments considered in Camerer and Ho, where subjects play against the same opponent(s) in each round and only receive feedback information on the results of interaction with this subject. In our experiment, we deliberately provided subjects with summary feedback information on the choices made by the entire cohort and the payoffs to each choice in the previous round.

[20] Camerer and Ho discuss the relationship between the EWA model and its special cases in detail.

We use the first 35 rounds for estimation and impose an homogeneity restriction $A_i^j(t) = A^j(t)$ for all i, and an identifying restriction $A^A(0) = 0$. Thus, the log-likelihood for the observed sample has eight parameters to be estimated:

$$\ell(N(0), A^B(0), A^C(0), \phi, \delta, \rho, \gamma, \lambda) = \sum_{t=1}^{35} \sum_i \sum_{j \in \{A,B,C\}} d_i^j(t) \ln(p_i^j(t)).$$

We estimated these parameters by maximum likelihood, imposing the same restrictions as Camerer and Ho: $\phi > 0$, $\lambda > 0$, $0 \leq \delta \leq 1$, $0 \leq \rho \leq 1$, $| A^j(0) | \leq \pi(B,A) - \pi(A,B)$, $0 \leq N(0) \leq (1-\rho)^{-1}$. For our additional parameter we imposed the restriction $0 \leq \gamma \leq 1$. The resulting estimates, standard errors, maximized log-likelihood, and out-of-sample scores are tabulated in Tables 7 and 8.

In terms of predictive success, the model improves upon the static models described in the previous section. In every session, the scoring rule gives a better score to the EWA model than to the best of the static models. It is interesting to see why this is the case. Although the model treats all subjects as identical, *ex ante*, and may seem ill-suited for modeling the data presented earlier (which exhibited substantial heterogeneity), the model scores well in part because it captures the heterogeneity of choices across subjects. To illustrate this, we construct a figure for each session showing estimated individual choice probabilities for round 36 using the first 35 rounds of data (Figures 10 and 11). Each plotted point corresponds to a subject, and the resulting clusters of points can be compared with Figures 8 and 9.

In the EWA model, heterogeneity can emerge endogenously as subjects accumulate individual-specific histories. However, the fictitious play element of the choice propensities is similar to all subjects. To see this, notice that the attractiveness of a given choice to different subjects contains a common component (the final term in the numerator of the expression for $A_i^j(t)$). If $\delta=1$ and $\gamma=0$, then this common component of attractions carries maximum weight, naturally reducing the amount of heterogeneity of subject strategies predicted by the model. In contrast, if δ=0 and γ=1, the common component of experiences across subjects carries no weight and heterogeneity will be more pronounced. Given the observations of Figures 8 and 9 regarding differences in the strategies employed by subjects, one may well anticipate that the EWA parameter estimates will lean toward the choice-reinforcement model.

This conjecture is supported in Table 9, which presents the estimated choice-reinforcement model for the Baseline game. In none of the Baseline sessions can the parameter restriction $\delta= \gamma-1=0$ be rejected (based on a comparison of the log-likelihood with the corresponding log-likelihood reported in Table 7). In addition, the predictive performance of the choice-reinforcement model is at least as good as that of the EWA model in four of the six sessions. (And in one of the other two sessions, the choice-reinforcement model out-performs the EWA model according to the *MSD* criterion.) Thus, for the Baseline sessions, individually experienced payoffs, rather than the summary data on payoffs, seem to be more important determinants of choice propensities, allowing the model to predict variability among subject choices within a round. These observations suggest that, at least for the Baseline sessions, a parsimonious characterization of the data generation process may be given by a Choice Reinforcement model, namely the EWA model with $\gamma = 1$ and $\delta = 0$.

For the Upside sessions the situation is not so clear (see Table 10). The parameter restriction is rejected at the 5% level for four of the six sessions, and in none of the sessions does the choice-reinforcement model unambiguously out-perform the EWA model in terms of out-of-sample prediction. Thus, while the choice-reinforcement model does an excellent job of explaining the heterogeneity observed in Baseline sessions, it does not seem able to explain changes in behavior as the riskiness of equilibrium strategies varies.

Another way to capture heterogeneity across subjects is to use models with individual specific parameters. Thus we also investigated a model in which we augmented the basic EWA model by adding individual specific initial choice propensities. Tables 11 and 12 compare the results of the choice-reinforcement, EWA and augmented EWA models in terms of maximized log-likelihood and out-of-sample predictive success. Because the augmented model introduces 18 additional parameters, relative to the EWA model, and 20 additional parameters relative to the Choice Reinforcement model, the maximized log-likelihoods will of necessity (weakly) favor the augmented EWA model. However, the individual-specific parameters are jointly insignificant in four of the sessions, and in only one of the remaining sessions does the addition of these parameters improve the predictive performance.

6. Conclusion

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

Our results strongly support the hypothesis that the Nash prediction improves when Nash play is less risky. Indeed, using 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, by increasing the upside sufficiently, play would 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.

In view of the limitations of the Nash and Maxmin predictions in organizing the data, we examine several prominent alternatives in our context. Of the static models, the QRE model is the best predictor of subject behavior; however, it is incapable of explaining the heterogeneity of subject behavior observed in the data. The dynamic EWA model does better in this regard despite the absence of clear trends in subject choices observed in the data. Its success hinges on the fact that it models play as depending on individual specific rewards. Subjects who have encountered favorable outcomes from choosing non-Maxmin strategies in the past are much more likely to play these strategies than are subjects who have encountering unfavorable outcomes. In the Upside game, the level of reward associated with these favorable outcomes is higher and, as a consequence, the dynamic models (correctly) predict that the proportion of Maxmin play will decline relative to the Baseline game. Thus the EWA model captures the main treatment effect observed in our data quite well. Indeed, in the Baseline sessions, a special case of the EWA model, the Choice Reinforcement model, leads to little decrease in model performance and provides a parsimonious explanation of subject choices. In the Upside game, however, the Choice Reinforcement model does not fare as well: for the majority of sessions in this treatment it is out-performed by EWA.

Yet, this cannot be the whole story. In describing the strategies they employed throughout the game, most subjects mentioned the importance of a forward-looking element in forming their strategies. That is, subjects claimed that their strategies depended upon what they anticipated would happen in the next period. Moreover, their willingness to play strategies other than the Maxmin on the basis of their forecast seemed to depend on the riskiness of non-Maxmin play. How seriously should one take these remarks? One test is to simply eliminate feedback about the choices made by other subjects in our games. If Choice Reinforcement is the correct model, then such a variation should have no effect on play. It is our conjecture, however, that play would be affected, as in Mookherjee and Sopher (1994). In our view a successful model of subject play in unprofitable games must incorporate notions of riskiness and anticipation with the reward-penalty structure of the Choice Reinforcement model. The suggested experiment and the expansion of the theory remain for future research.

References

Aumann, R.J. and M. Maschler (1972). Some thoughts on the minimax principle. *Management Science,* **18**, 54-63.

Aumann, R.J. (1985). On the Non-Transferable Utility Value: A Comment on the Roth-Shaffer Examples. *Econometrica*, **53**, 667-677.

Bernheim, D. (1984). Rationalizable Strategic Behavior. *Econometrica,* **52**, 1007-1028.

Brown, J. and R. Rosenthal (1990). Testing the Minimax Hypothesis: A Reexamination of O'Neill's Game Experiment. *Econometrica*, **58**, 1065-1081.

Camerer, C. and Teck-Hua Ho (1999). Experience-Weighted Attraction Learning in Normal Form Games. *Econometrica*, 67, 827-874.

Chen, Y. and F-F Tang (1998). Learning and Incentive-Compatible Mechanisms for Public Goods Provision: An Experimental Study. *Journal of Political Economy*, **106**, 633-662.

Fudenberg, D. and D. Levine (1998). *The Theory of Learning in Games*. MIT Press.

Gilboa, I. and D. Schmeidler (1989). Maxmin expected Utility with non-unique prior. *Journal of Mathematical Economics*, **18**, 141-153.

Harsanyi, J.C. (1964). A General Solution for Finite Non-Cooperative Games, based on Risk-Dominance. In Dresher, M., L.S. Shapley and A.W. Tucker (Eds) *Advances in Game Theory*. (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 University Press.

Holler, M. and V. Host (1990). Maximin vs. Nash Equilibrium: Theoretical Results and Empirical Evidence. In Quandt, R. and D. Triska (Eds) *Optimal Decisions in Markets and Planned Economies*. (Boulder: Westview Press.)

Lo, K.C. (1996). Equilibrium in Beliefs under Uncertainty. *Journal of Economic Theory*, **71**, 443-484.

McKelvey, R. and T. Palfrey, (1995). Quantal Response Equilibria for Normal Form Games. *Games and Economic Behavior,* **10**, 6-38.

Mookerjee, D. and B. Sopher, (1994). Learning Behavior in an Experimental Matching Pennies Game. *Games and Economic Behavior*, **7**, 62-91.

Ochs, J. (1995). Games with Unique Mixed Strategy Equilibria: An Experimental Study. *Games and Economic Behavior*, **10**, 174-189.

O'Neill, B. (1987). Nonmetric Test of the Minimax Theory of Two-person Zerosum Games. *Proceedings of the National Academy of Science*, USA, **84**, 2106-2109.

Pearce, D. (1984). Rationalizable Strategic Behavior and the Problem of Perfection. *Econometrica*, **52**, 1029-1050.

Roth and Erev (1995). Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term. *Games and Economic Behavior*, **8**, 164-212.

Rubinstein, A. (1991). Comments on the Interpretation of Game Theory, *Econometrica*, **59**, 909-924.

Selten, R., A. Sadrieh, and K. Abbink (1995). Money Does Not Induce Risk Neutral Behavior, but Binary Lotteries Do Even Worse. Universitat Bonn Sonderforschungsbereich 303, Discussion Paper B343.

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 of Operations Research*, **38**, 189-205.

Tables

Table 1. Choice Proportions for Baseline and Upside Games				
	Baseline		Upside	
Choice	First Round	All Rounds	First Round	All Rounds
A	13.33	14.13	41.67	46.5
B	26.67	32.90	16.67	23.03
C	60.00	52.97	41.67	30.47

Table 2. 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

Table 3. Maximum Likelihood Estimates of QRE Models for Baseline Game					
Session	λ	se(λ)	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

Table 4. Maximum Likelihood Estimates of QRE Models for Upside Game					
Session	λ	se(λ)	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

Table 5. Predictive Success of Static 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

Table 6. Predictive Success of Static 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

Table 7. Estimated Experience-Weighted Attraction Model for Baseline Game (standard errors in parentheses)											
Session	$A^B(0)$	$A^C(0)$	ρ	N(0)	λ	φ	δ	γ	$-\ell$	$-\ell_{35+}$	MSD
UK_1	0.289 (0.410)	0.556 (0.470)	1.000 (0.000)	8.287 (8.076)	1.245 (0.655)	0.959 (0.027)	0.000 (0.000)	1.000 (0.000)	311.877	113.787	0.215
UK_2	-5.000 (0.000)	5.000 (0.000)	0.758 (0.143)	0.000 (0.000)	0.244 (0.087)	0.965 (0.033)	0.134 (0.103)	1.000 (0.000)	323.218	138.949	0.288
UK_3	0.410 (0.837)	0.587 (1.196)	0.993 (0.066)	18.412 (45.648)	3.151 (6.143)	0.930 (0.041)	0.476 (0.255)	0.000 (0.017)	282.606	119.939	0.228
US_1	3.561 (0.861)	5.000 (0.014)	0.910 (0.047)	2.378 (1.348)	0.516 (0.104)	0.935 (0.032)	0.000 (0.000)	1.000 (0.000)	298.978	112.290	0.219
US_2	0.282 (0.272)	0.428 (0.270)	1.000 (0.000)	35.023 (14.748)	7.060 (0.480)	0.902 (0.034)	0.425 (0.027)	0.409 (0.027)	213.404	90.262	0.168
US_3	-0.374 (0.693)	0.206 (0.704)	0.900 (0.062)	2.282 (2.172)	0.750 (0.345)	0.931 (0.036)	0.101 (0.143)	1.000 (0.001)	266.093	85.456	0.154

Table 8. Estimated Experience-Weighted Attraction Model for Upside Game (standard errors in parentheses)											
Session	$A^B(0)$	$A^C(0)$	ρ	N(0)	λ	φ	δ	γ	$-\ell$	$-\ell_{35+}$	MSD
UK_1	-6.573 (7.984)	11.000 (0.032)	0.226 (0.392)	0.599 (0.433)	0.097 (0.049)	0.892 (0.029)	0.320 (0.131)	0.204 (0.167)	346.293	133.763	0.268
UK_2	-11.000 (0.006)	-1.295 (2.226)	0.717 (0.119)	0.000 (0.000)	0.277 (0.090)	0.897 (0.035)	0.159 (0.126)	0.448 (0.169)	283.457	87.931	0.161
UK_3	-2.693 (2.254)	-0.427 (1.155)	0.984 (0.044)	0.136 (0.558)	0.575 (0.183)	1.002 (0.035)	0.259 (0.148)	0.341 (0.187)	312.697	157.714	0.323
US_1	-2.010 (3.696)	0.118 (1.701)	0.747 (0.319)	2.128 (3.820)	0.278 (0.323)	0.918 (0.042)	0.317 (0.171)	0.454 (0.250)	321.603	132.937	0.266
US_2	-1.265 (1.154)	1.172 (1.008)	0.743 (0.031)	3.896 (0.474)	0.350 (0.029)	0.873 (0.029)	0.132 (0.122)	0.276 (0.178)	317.606	105.536	0.200
US_3	-2.070 (4.176)	2.609 (3.939)	0.658 (0.159)	0.000 (0.002)	0.196 (0.080)	0.921 (0.030)	0.000 (0.000)	0.413 (0.241)	294.898	109.698	0.208

Table 9. Estimated Choice-Reinforcement Model for Baseline Game (standard errors in parentheses)									
Session	$A^B(0)$	$A^C(0)$	ρ	N(0)	λ	ϕ	$-\ell$	$-\ell_{35+}$	MSD
UK_1	0.289 (1.090)	0.556 (1.395)	1.000 (0.001)	8.286 (35.177)	1.245 (2.958)	0.959 (0.156)	311.877	113.787	0.215
UK_2	-5.000 (0.000)	5.000 (0.000)	0.756 (0.123)	0.000 (0.002)	0.226 (0.072)	0.961 (0.034)	323.653	139.093	0.287
UK_3	0.636 (0.874)	0.769 (1.035)	1.000 (0.000)	44.868 (76.448)	2.936 (4.058)	0.946 (0.029)	283.932	116.825	0.223
US_1	3.561 (0.841)	5.000 (0.006)	0.910 (0.046)	2.378 (1.249)	0.516 (0.099)	0.935 (0.029)	298.978	112.290	0.219
US_2	0.402 (0.501)	0.495 (0.606)	1.000 (0.000)	85.820 (119.281)	7.705 (9.500)	0.914 (0.025)	214.362	87.894	0.165
US_3	-0.347 (0.815)	0.320 (0.716)	0.901 (0.074)	2.477 (2.700)	0.694 (0.368)	0.935 (0.038)	266.281	86.059	0.156

Table 10. Estimated Choice-Reinforcement Model for Upside Game (standard errors in parentheses)									
Session	$A^B(0)$	$A^C(0)$	ρ	N(0)	λ	ϕ	$-\ell$	$-\ell_{35+}$	MSD
UK_1	-8.041 (11.329)	11.000 (0.027)	0.554 (0.240)	0.409 (0.483)	0.097 (0.046)	0.905 (0.026)	349.382	135.112	0.270
UK_2	-11.000 (0.000)	-1.428 (2.375)	0.736 (0.123)	0.000 (0.002)	0.231 (0.075)	0.901 (0.032)	286.770	90.226	0.165
UK_3	-11.000 (0.010)	-1.165 (1.784)	0.943 (0.049)	0.000 (0.002)	0.324 (0.064)	0.992 (0.031)	317.001	156.078	0.323
US_1	-3.776 (6.115)	-0.970 (2.610)	0.809 (0.260)	2.652 (4.570)	0.206 (0.250)	0.944 (0.031)	323.039	132.954	0.266
US_2	-3.108 (3.897)	3.663 (3.040)	0.406 (0.136)	1.684 (0.387)	0.113 (0.027)	0.880 (0.023)	322.533	109.060	0.206
US_3	-3.154 (7.343)	3.971 (5.683)	0.557 (0.256)	0.000 (0.026)	0.129 (0.066)	0.930 (0.023)	297.117	112.567	0.214

Table 11. Comparison of models for Baseline Game [comparison with EWA in parentheses]									
	Choice Reinforcement			EWA			Augmented EWA		
Session	$-\ell$	$-\ell_{35+}$	MSD	$-\ell$	$-\ell_{35+}$	MSD	$-\ell$	$-\ell_{35+}$	MSD
UK_1	311.877 [0.000]	113.787	0.215	311.877	113.787	0.215	278.818 [66.120***]	107.388	0.210
UK_2	323.653 [0.870]	139.093	0.287	323.218	138.949	0.288	293.820 [58.796***]	221.353	0.372
UK_3	283.932 [2.652]	116.825	0.223	282.606	119.939	0.228	259.518 [46.176***]	152.168	0.241
US_1	298.978 [0.000]	112.290	0.219	298.978	112.290	0.219	264.130 [69.696***]	111.940	0.225
US_2	214.362 [1.916]	87.894	0.165	213.404	90.262	0.168	184.766 [57.276***]	95.067	0.180
US_3	266.281 [0.376]	86.059	0.156	266.093	85.456	0.154	263.222 [5.742]	85.324	0.153

Table 12. Comparison of models for Upside Game [comparison with EWA in parentheses]									
	Choice Reinforcement			EWA			Augmented EWA		
Session	$-\ell$	$-\ell_{35+}$	MSD	$-\ell$	$-\ell_{35+}$	MSD	$-\ell$	$-\ell_{35+}$	MSD
UK_1	349.382 [6.178**]	135.112	0.270	346.293	133.763	0.268	318.864 [54.858***]	141.017	0.294
UK_2	286.770 [6.626**]	90.226	0.165	283.457	87.931	0.161	277.091 [12.732]	88.079	0.161
UK_3	317.001 [8.608**]	156.078	0.323	312.697	157.714	0.323	303.335 [18.724]	158.182	0.324
US_1	323.039 [2.872]	132.954	0.266	321.603	132.937	0.266	291.827 [59.552***]	180.694	0.307
US_2	322.533 [9.854***]	109.060	0.206	317.606	105.536	0.200	287.678 [59.856***]	139.194	0.291
US_3	297.117 [4.438]	112.567	0.214	294.898	109.698	0.208	285.147 [19.502]	108.725	0.206

Figures

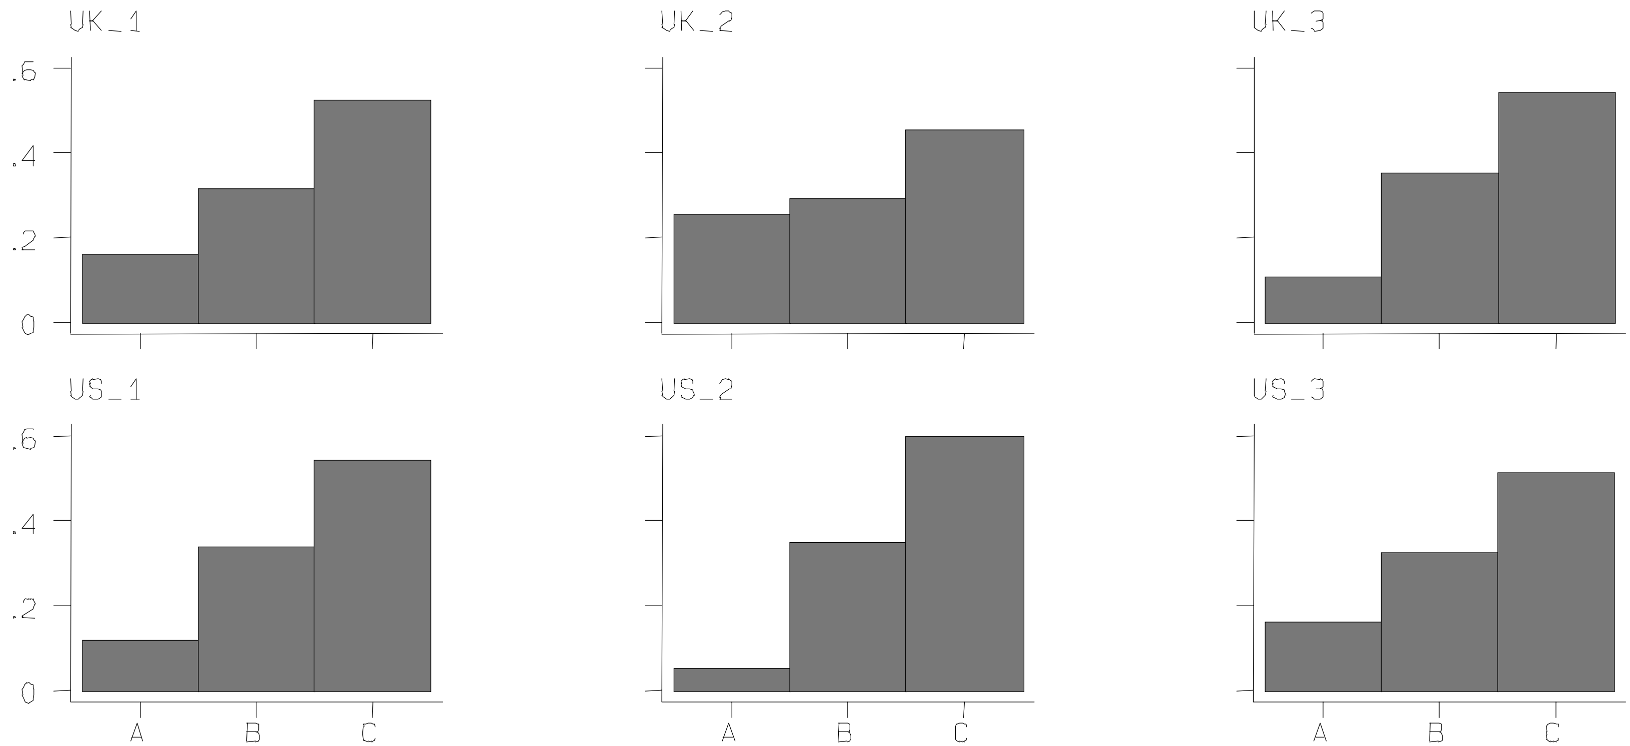

Figure 4. Histograms of Choices in Baseline Sessions.

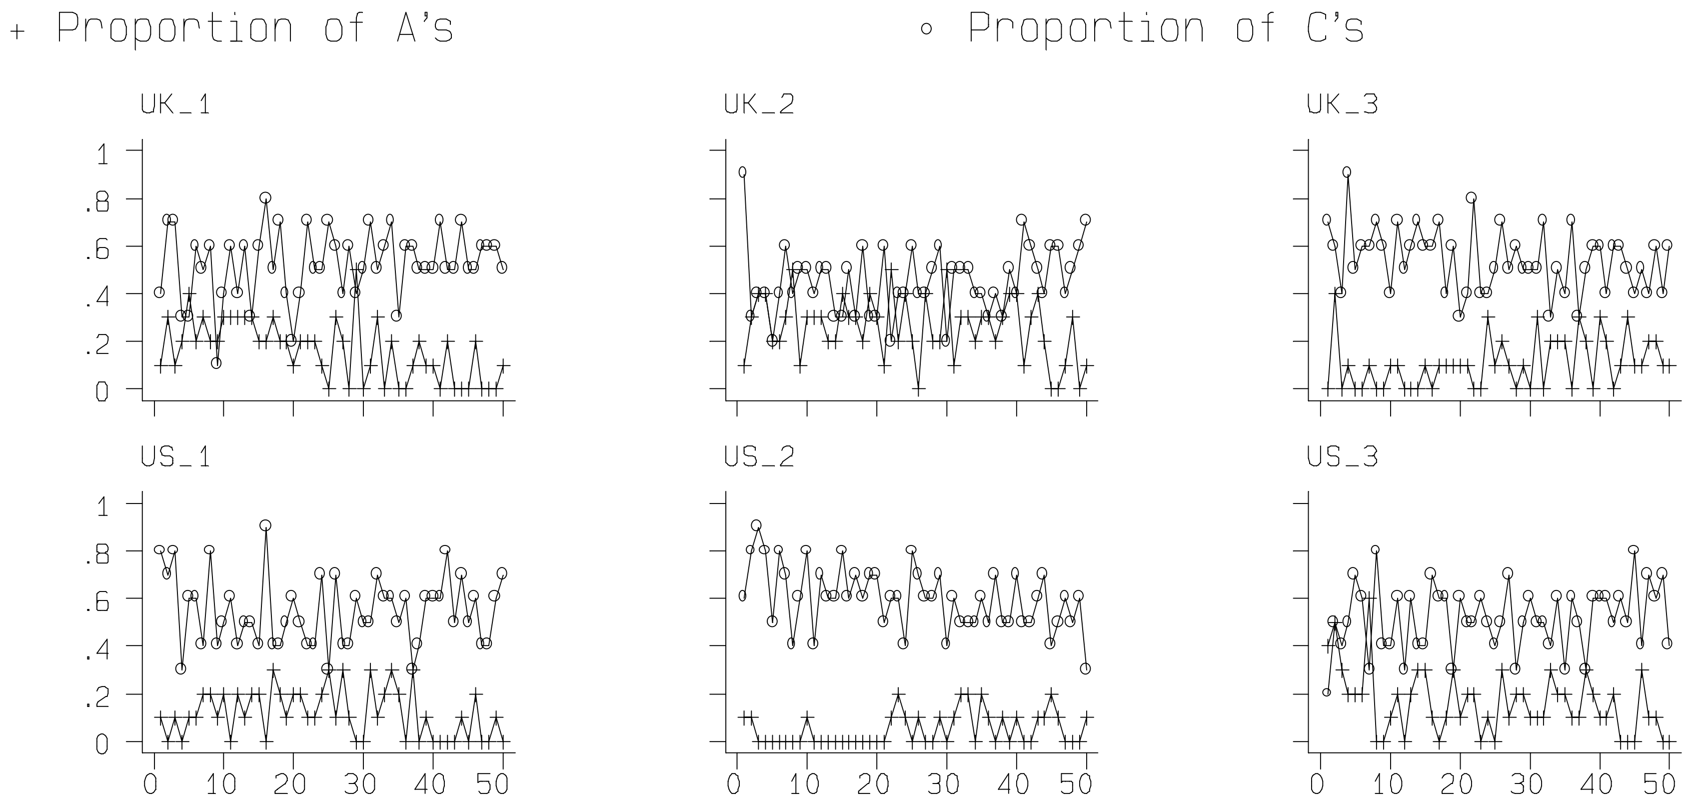

Figure 5. Proportions of Choices Across Rounds in Baseline Sessions.

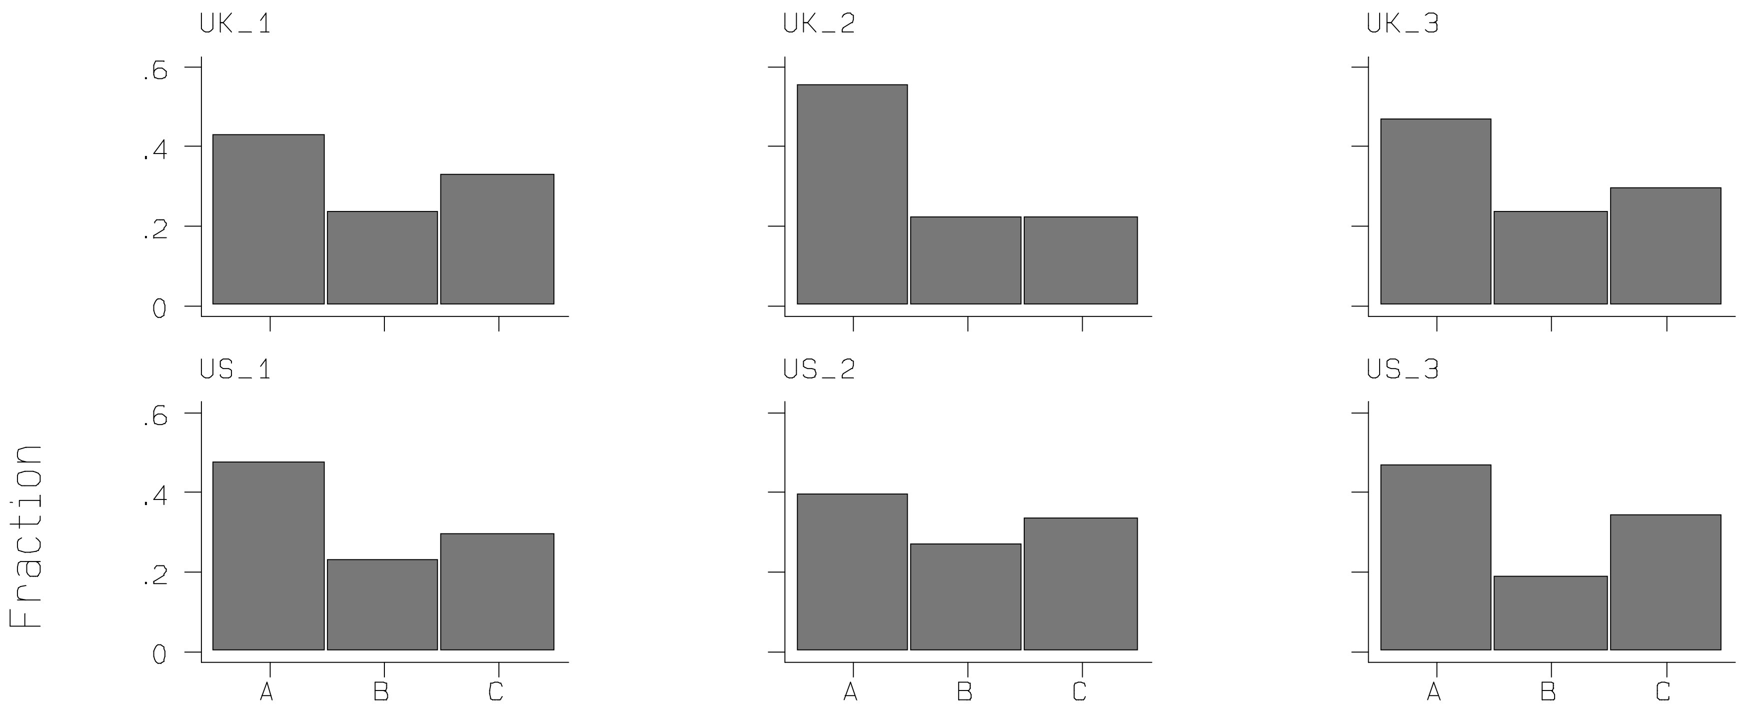

Figure 6. Histograms of Choices in Upside Sessions.

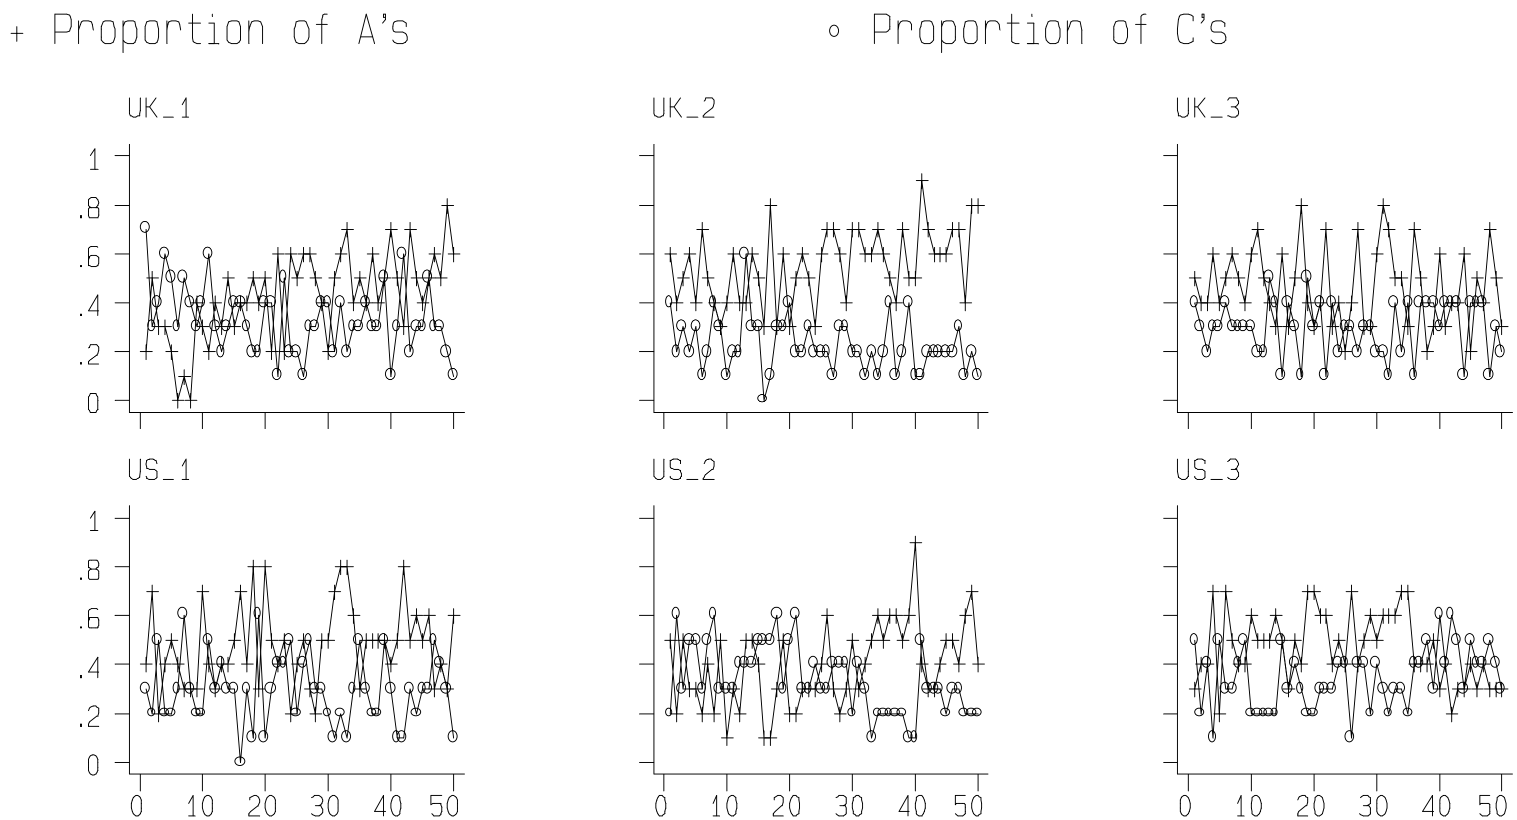

Figure 7. Proportions of Choices Across Rounds in Upside Sessions.

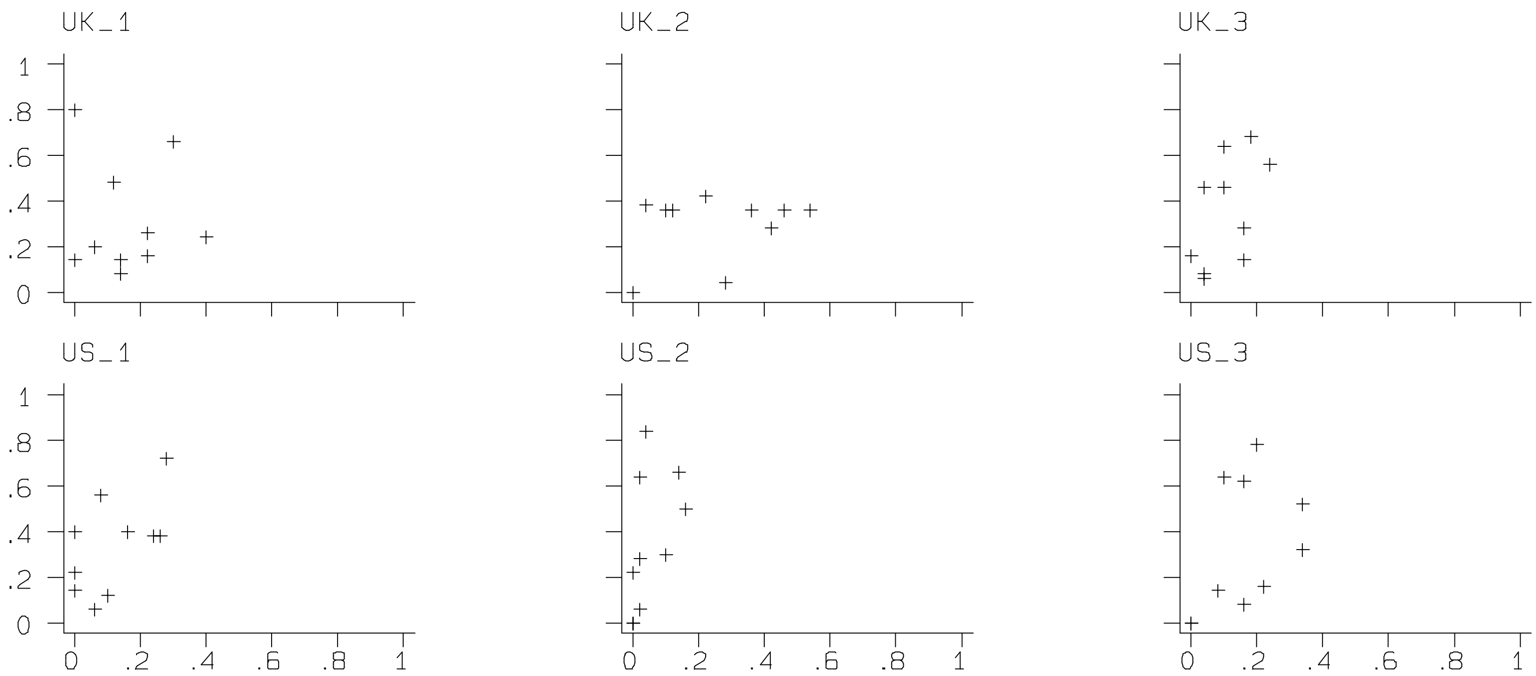

Figure 8. Proportions of Choices Across Subjects in Baseline Sessions.

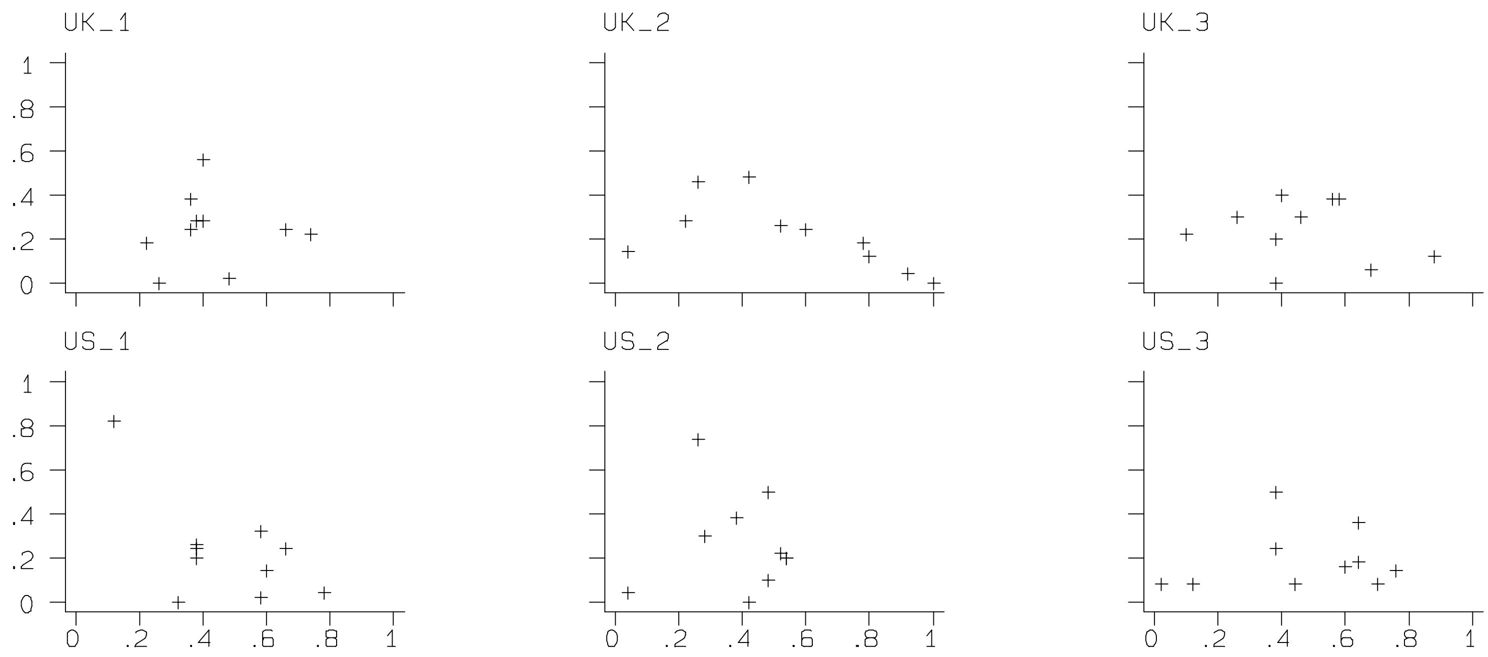

Figure 9. Proportions of Choices Across Subjects in Upside Sessions.

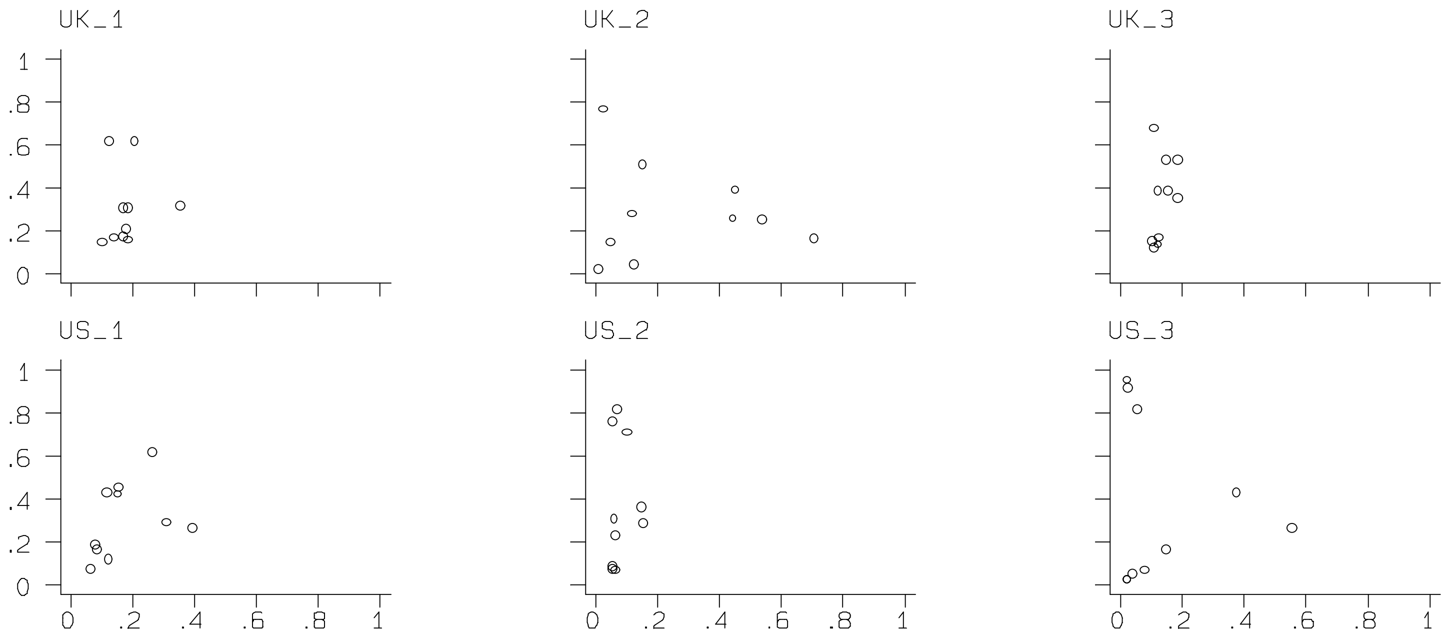

Figure 10. Predicted Choice Triangles for Baseline Game

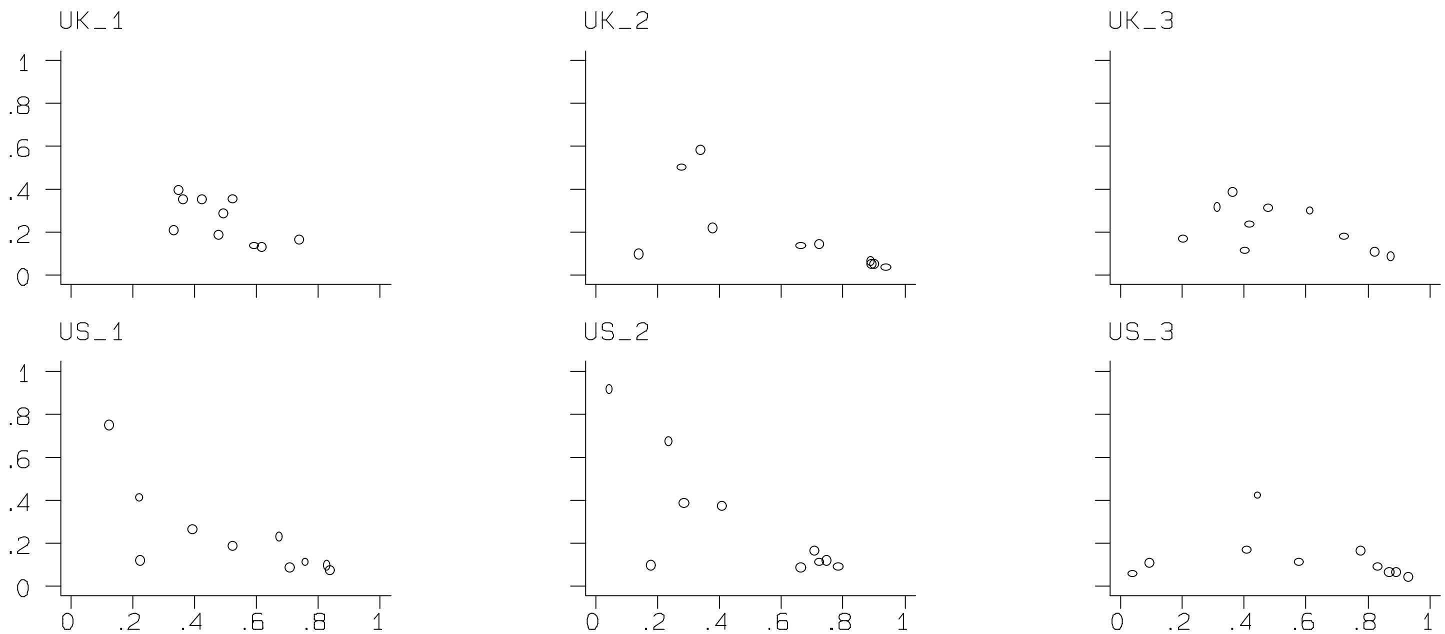

Figure 11. Predicted Choice Triangles for Upside Game

Appendix A. Proofs

Proposition 1. There does not exist any 2x2 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 also a Maxmin strategy.

Proposition 2. There does not exist a symmetric 3x3 unprofitable game with two levels of payoff and distinct Nash and Maxmin strategies.

Proof: Consider the payoffs to row player represented as a 3x3 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 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 3x3 2 outcome symmetric game may be represented by the matrix:

A 3x3 2 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 3x3 unprofitable game with distinct Nash and Maxmin profiles in pure strategies has a symmetric Nash equilibrium in mixed strategies.

Proof: Consider the set of 3x3 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), the Maxmin strategy be C, and normalize the equilibrium payoff strategy to be 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

In order 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

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

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 either 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	3
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 either 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 fifty 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 $ ______.