

IBM Research Report

Hide and Seek in Arizona

Robert Rosenthal

Department of Economics
Boston University

Jason Shachat

IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Mark Walker

Department of Economics
University of Arizona



Research Division

Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

Hide and Seek in Arizona*

Robert W. Rosenthal[†], Jason Shachat[‡], and Mark Walker[§]

June 29, 2001

Abstract

Laboratory subjects repeatedly played one of two variations of a simple two-person zero-sum game of pursuit and evasion. Three puzzling departures from the prescriptions of equilibrium theory are found in the data: an asymmetry related to the player's role in the game; an asymmetry across the game variations; and positive serial correlation in subjects' play. Possible explanations for these departures are considered.

*We are grateful to Todd Swarthout for invaluable research assistance and to Kevin Lang, Mark Machina, and John Wooders for helpful comments. Some of the work was done while Rosenthal was visiting the Department of Economics at MIT and the Department of Political Economy at the University of Siena and while Shachat was visiting IBM Research Laboratories. All three authors are grateful to the National Science Foundation for financial support of this research.

[†]Department of Economics, Boston University; deceased February 27, 2002

[‡]IBM TJ Watson Research Laboratory; jshachat@us.ibm.com

[§]Department of Economics, University of Arizona; mwalker@arizona.edu

1 Introduction

Barry O’Neill’s insightful 1987 paper revived interest in the long-dormant idea of using experimental methods to test the predictive power of mixed-strategy equilibrium. Following O’Neill’s experiment, a number of other researchers have conducted new experiments based on two-person zero-sum games that have unique mixed-strategy equilibria.¹ The evidence for and against equilibrium theory has been mixed: While the aggregate data is typically better explained by the simple equilibrium hypothesis than by alternative simple theories, and while the theory seems to perform better than it did in the experiments done prior to O’Neill’s,² there have nevertheless been consistent departures from the implications of equilibrium theory in each of the experiments.

We report here on an experiment designed to shed some light on these empirical departures from the theory’s predictions, and on some related issues. The experiment was based upon two simple games of pursuit and evasion, or “hide and seek.” Since the games our subjects played are seemingly simpler than the games used in most previous experiments, the theory might be expected to perform better here than it had done before.

Each of our subjects played only one of the two games, playing it many times against a single, unchanging opponent (another subject). Although the two games are different, they have the same representation as a 2×2 normal form game in which each of the players, the Pursuer and the Evader, can move either to the Left or to the Right. The 2×2 game has a unique equilibrium, in which the Pursuer and the Evader each use the same mixture: each player plays his Left action with probability $2/3$ and his Right action with probability $1/3$.

The subjects’ play in our experiment resembled, in some respects, the results that have been consistently observed in previous experiments: choice frequencies when aggregated across all subjects were relatively close to the equilibrium frequencies, there was substantially more variation in choice frequencies across subjects than the theory predicts, and there was significant serial correlation in choices. But there were also some departures from the theory that had not been observed in previous experiments. These departures were unexpected and were strikingly systematic. They

¹See, for example, Rapoport and Boebel (1992), Mookherjee and Sopher (1994, 1997), Bloomfield (1994), Ochs (1994), Shachat (forthcoming), and Binmore, et al. (forthcoming). See also Brown and Rosenthal (1990) and O’Neill (1991) for an analysis and discussion of the extent to which the data support the theory.

²For example, Suppes and Atkinson (1960), and Malcolm and Lieberman (1965).

were of three kinds:

- (1) Asymmetric role behavior: Nearly every Evader chose Left more often than his opponent.
- (2) A game asymmetry effect: Both roles played Left more often in one of the games than in the other.
- (3) Strong positive serial correlation: In contrast to previous experiments, in which subjects tended to change their actions too frequently to be consistent with random choice, our subjects switched from one action to the other too rarely.

We describe the experiment in detail in Section 2 and we describe the data in Section 3. In Section 4 we consider three prominent alternative models, and we investigate whether any of these models does better than the standard equilibrium model at organizing the data. The evidence is somewhat mixed. All the alternative models, for example, are able to accommodate asymmetric role behavior, but not at the magnitude observed in our data, and none of the models explains other important features of the data. Although we do not have a complete theory to offer, we have some opinions about what is driving the main features of our data and about why the alternative models fail to explain the data. All of this is discussed in the concluding Section 5.

2 The Experiment³

At the beginning of each experimental session, the subjects were matched into pairs. Each pair of subjects then repeatedly played a simple two-person zero-sum game. The game was described to the subjects as a game of pursuit and evasion. Two games were used in the experiment, but each pair of subjects played only one of the games, playing the same game either 100 times or 200 times. We call the two stage games the Stochastic Game and the Deterministic Game.

In the Stochastic Game, each player has two choices, Left and Right. If the players choose different directions (one Left and one Right), then the Evader has successfully avoided detection and no money changes hands. If the players choose the same direction (both Left or both Right), then the Pursuer has “found” the Evader, and a random device then determines whether the

³All sessions of the experiment took place in 1996 and 1997 at the University of Arizona. The subjects were all Arizona undergraduates recruited from undergraduate economics courses. No subject participated in more than one session. Sample instructions for subjects are contained in an appendix.

Pursuer wins money from the Evader: if both chose Left, then with probability $1/3$ the Pursuer wins a fixed amount of money from the Evader; and if both chose Right, then with probability $2/3$ the Pursuer wins the same amount from the Evader.

In the experiment, the subjects sat side-by-side at a table, separated from each other by an opaque screen, facing a human monitor who recorded all details of play.⁴ The subjects' roles as pursuers or evaders were chosen by lot at the outset, and each subject retained both the same role and the same opponent during the entire course of play. Each subject had two cards, one with the label L on its face, the other with R on its face. In each round of play, each subject selected one of his two cards, placing it face down and sliding it toward the monitor. Both cards were then turned face up within sight of all three individuals. If the cards matched, the monitor rolled a single die; if the die came to rest with one or two spots showing after both subjects had played L , then the Evader transferred to the Pursuer a token with a known monetary value (to be described momentarily); and if both subjects had played R , then the Evader transferred a token when the die showed three or more spots. The Evader was given some working capital of tokens at the beginning of play and the Pursuer started with no tokens.

Half of our subjects played the Stochastic Game just described, and the other half played the Deterministic Game. The Deterministic Game differs from the Stochastic Game in just one respect: in the Deterministic Game when the two players choose the same direction (both Left or both Right) the Pursuer wins an amount from the Evader *for certain*, and he wins twice as much when both have chosen Right as when both have chosen Left.

Note that the (normalized) expected monetary payoffs for the Pursuer in both games can be described by the same payoff matrix

		Evader	
		Left	Right
Pursuer	Left	1	0
	Right	0	2

The utility payoffs associated with these monetary payoffs will be considered below.

⁴Our intent was that each subject see his or her opponent, drawn from the same subject pool, at the beginning of play, so that any thoughts of the opponent's being a stooge or a computer program would be minimized. The opaque screen prevented each player from being able to key on the opponent's movements, other than, perhaps, hesitations, during the course of play.

The monetary transfers in the experiment were effected by using push-pins as tokens. Each subject had a pad into which the pins could be easily pushed, and the pads were demarcated so that the subject could easily see both his current balance of pins and the dollar value of that balance. The pads were redeemed for cash at the end of the session at a monetary rate per pin that was stated at the beginning of the session. The monetary redemption rate for the Stochastic Game was a half-dollar per pin. For the Deterministic Game it was one-sixth of a dollar per pin with one pin being transferred after a Left-Left match and two pins after a Right-Right match. (The redemption rates were selected so that at the equilibrium mixture the expected monetary payoffs were the same in both stage games.) The Evaders began with \$25 worth of tokens in the 100-play sessions and with \$50 worth of tokens in the 200-play sessions. (It was therefore theoretically possible for an Evader to become bankrupt before the nominal end of the experiment, and the players were notified that if that were to happen then the game would terminate at that time. In practice, no Evader's balance ever fell below \$6.50.⁵)

Twenty pairs of subjects played the Stochastic Game repeatedly, and another twenty pairs played the Deterministic Game. In each case, fourteen of the pairs played the game 100 times and the remaining six pairs played 200 times.

The natural prediction from game theory is that the subjects in the experiment will play a Nash equilibrium of the repeated game (100 plays or 200 plays), which might or might not consist of the players playing at each stage the equilibrium of the stage game. Moreover, the equilibrium of the stage game potentially depends on the players' risk attitudes.

O'Neill and others explicitly avoided risk attitudes in their stage games by devising games with only two outcomes (for example, every action profile leads either to one player or the other winning a nickel). The Wooders-Shachat (2001) Theorem ensures that for such two-outcome stage games the only equilibrium of the finitely repeated game consists of the players playing the stage-game equilibrium at every stage. Thus, the two-outcome design introduced by O'Neill provided a much

⁵The reader may verify that having one subject pay the other from an initial endowment enables the experimenter to provide stronger incentives (i.e., larger payoff differences) in each round per dollar paid to the subjects, as compared with a design in which both subjects are paid nonnegative amounts in each round. The disadvantage of this design element is that it raises the bankruptcy problem. But, as the probability of bankruptcy is negligible under any reasonable description of play, we assume that the subjects' decisions were not affected by this possibility.

sounder connection between equilibrium theory and mixed-strategy equilibrium experiments.

But O’Neill’s elimination of theoretical concerns about the subjects’ preferences came at some cost: because the games used only deterministic payoffs, four or more actions for each player were needed in order to produce a game that had only two outcomes and in which the “right” way to play was not transparent.⁶ In our Stochastic Game, each player has only two actions, and the Nash equilibria in both the stage game and the repeated game are still independent of players’ risk attitudes under the additional hypothesis that the preferences satisfy the compound lottery reduction axiom of decision theory. By making the payoff in each cell of the stage game a lottery over just two basic outcomes, instead of a deterministic monetary outcome, we have created a game that seems much simpler and might therefore give the theory its best chance to succeed in predicting behavior accurately.

If we assume that the subjects prefer first-order stochastically dominating lotteries and that the subjects satisfy the compound lottery reduction axiom, then the players’ stage-game utility payoffs in the Stochastic Game are simply their expected monetary payoffs. As noted above, there is a unique equilibrium in this stage game in which both the Pursuer and the Evader play the mixture that places $2/3$ probability on Left and $1/3$ probability on Right. The monetary value of the game to the Pursuer is $1/9$ of a dollar and the value to the Evader is $-1/9$ of a dollar (i.e., $2/9$ of a token; recall that whenever a monetary transfer occurs in the Stochastic Game it is a half-dollar).

The Deterministic Game has three outcomes, not two. However, if the players are risk-neutral, then the monetary payoffs in the stage game can again be taken to be their utilities, and the Deterministic Game is game-theoretically equivalent to the Stochastic Game. For most of our analysis, we will assume that our experimental subjects were indeed risk neutral at each play of the game, since they were playing a succession of 100 or 200 plays of the game for relatively small stakes. In Section 5.1, however, we will take some account of the effect of risk attitudes, and we will find that Prospect Theory may provide a partial explanation of differences in play between the Stochastic and Deterministic Games.

⁶The only mixed-equilibrium games where theory does seem to square with observation are games with so much symmetry that the equilibrium consists of each player mixing uniformly over his actions – games such as “matching pennies” or “rock, paper, scissors.” It is generally agreed that the experimental results in these games may be driven by the symmetric frame and are therefore not useful for testing equilibrium theory in general. Cf. Mookherjee and Sopher (1994).

To the extent that we can ignore risk attitudes in the stage games and to the extent that we can ignore repeated play effects, observed play in the Deterministic Game sessions should, according to equilibrium theory, be the same as in the Stochastic Game sessions: each subject should mix $2/3$ on Left and $1/3$ on Right, independently of all else, at every play of each game.

As already mentioned, the monetary values of the push-pin tokens, described above, made the equilibrium values of both the Stochastic and Deterministic Games equal to one-ninth of a dollar for the Pursuer. Under equilibrium play, then, the Pursuer’s expected earnings were a little more than \$11 in the 100-play experiments and a little more than \$22 in the 200-play sessions, and the Evader could therefore expect on average to keep about \$14 of the \$25 he started with in the 100-play experiments and about \$28 of the \$50 he started with in the 200-play sessions. Each experimental session lasted for about one hour.

3 The Data

Individual subjects’ play, aggregated over all 100 or 200 of the subjects’ action choices, is presented in Tables 1(a) and 1(b) and in Figure 1. Each row of Table 1 describes the play of one pair of subjects. The first twenty pairs of subjects played the Stochastic Game: the pairs labeled #1 to #14 played 100 repetitions each (*i.e.*, there were 100 “plays,” indicated in the second column of Table 1); and the six remaining pairs (#15 to #20) played 200 repetitions each. The twenty pairs labeled #21 to #40 played the Deterministic Game; again, fourteen of these pairs played 100 repetitions and six pairs played 200 repetitions. The four columns labeled “Action Profile” in Table 1 report the frequencies with which each of the four joint action profiles LL, LR, RL, and RR were played by each pair. The next two columns, “Pur L” and “Eva L,” report the frequencies with which the Pursuer subject or the Evader subject played his Left action (*i.e.*, the Pur L entry in any row is the sum of the LL and LR entries, and the Eva L entry is the sum of the LL and RL entries). Figure 1 depicts these Pur L and Eva L frequencies for each pair of subjects in a scatter diagram: each of the forty points in Figure 1 represents the frequencies of Left play by a single pair of subjects.

We describe several noteworthy features of the data, organized around the following questions: Did the subjects play Left and Right in the equilibrium proportions? Did Pursuers and Evaders play

the same way, as the symmetric equilibrium predicts? Was play the same in both the Stochastic and Deterministic Games, as their identical normal form representation predicts? And were the subjects playing randomly, with no serial correlation, as equilibrium theory predicts?

3.1 Equilibrium Frequencies

The first thing one sees in the data in Table 1 and in Figure 1 is that the theoretical equilibrium has some qualitative saliency, but that the equilibrium's predicted mixtures do not accurately describe the subjects' play. The equilibrium mixture for each player in each game calls for Left with probability $2/3$, so we should see substantially more Left play than Right. And indeed, every Evader and almost every Pursuer did play Left more often than Right, and most of them played Left substantially more often. But it is clear in Figure 1 that the Evaders, on average, played Left significantly more often than the equilibrium rate of two-thirds, and further that there was a great deal of subject-to-subject variation. And the Pursuers, while indeed playing Left about two-thirds of the time on average, also displayed substantial variation from subject to subject. Column #9 in Table 1 presents the p -value of the chi-square test that compares the joint action frequencies (columns #3 to #6 in the table) with the Nash equilibrium frequencies. The equilibrium frequencies are rejected at the 5% significance level for 14 of the 20 pairs in each game.

However, since the games in our experiment are zero-sum, if only one player adopts his part of the mixed-strategy equilibrium, both players still receive their equilibrium (or minimax) payoffs. Given this, it is natural to ask whether either of the players in a fixed pair is playing the equilibrium strategy. If neither is, then we have a stronger rejection of the theory: neither player is ensuring his minimax value and both players' play is exploitable. We assess whether each subject is playing the equilibrium proportion of Left play with a simple binomial hypothesis test. Column #10 in Table 1 reports the value of $Pr(x; n, p)$, where $Pr(\cdot; \cdot, \cdot)$ is the cumulative distribution function for the binomial distribution, x is the number of periods the Pursuer chose Left, n is the total number of stage games, and p is the Nash equilibrium mixture probability, $2/3$. Column #11 reports the same calculation for the Evaders. If we conduct a two-tailed test at the 5% level of significance, we reject equilibrium play for probabilities less than .025 and greater than .975. The Nash equilibrium frequencies are rejected at the 5% level for 45 of our 80 subjects. In 12 of the 40 pairs, we reject that *either* of the subjects' play is consistent with equilibrium.

Since our subjects all played at least 100 times, and some 200 times, the question arises whether their behavior changed over time, and in particular whether their play could have been converging to an equilibrium. Figure 2 provides some evidence on this question: it shows the frequency of Left play in successive time blocks of twenty periods each. The figure depicts in separate panels the subjects who played 100 periods and those who played 200; in the top half of the figure is the play in the Deterministic Game, and in the bottom half the play in the Stochastic Game. In none of the four panels is there any indication that play is moving toward equilibrium (the dotted line, for both player roles). Even in the 200-play panels, on the right-hand side of Figure 2, the sum of the distances of the two roles' frequencies from the equilibrium frequencies is larger in the final two blocks than it is in two of the three immediately preceding blocks.

3.2 Symmetry of Player Roles

Another feature of the data that stands out sharply is that nearly every one of the points in Figure 1 lies above the 45°-line: the subjects who played the Pursuer role and those who played the Evader role were not playing the same, despite the fact that the equilibrium mixtures for Pursuer and Evader are the same in each game.⁷ In 35 of the 40 pairs the Evader played Left more often than the Pursuer, in many cases much more often. In only 3 pairs (#16, 19, and 22) did the Pursuer play Left more often than the Evader, and in each case it was only slightly more often. Aggregated across all forty pairs, the Pursuers played Left 64% of the time, the Evaders 74% of the time. This is a striking departure from the equilibrium prediction: while the equilibrium predicts stochastic play, so that even with identical equilibrium mixtures we should not expect the realized play of any pair of subjects to be exactly the same, nevertheless the two players' play should be *distributed* identically. But the Pursuers and the Evaders were clearly playing differently. We refer to this feature of the data, in which the two players' observed behavior was distinctly asymmetric, as *asymmetric role behavior*.

For a formal test whether Pursuers and Evaders were playing the same (the null hypothesis) or whether Evaders were playing Left with higher probability (the alternative), we conduct a

⁷In Binmore, et. al. (forthcoming) subjects played a game with the same normal form, under considerably different conditions. A similar role asymmetry appears in their data. A similar phenomenon appears in Rubinstein (1999), but there the behavior is what one would expect from the unusual informational conditions.

nonparametric sign test. Under the null hypothesis the probability that a given Evader’s realized number of Left choices will exceed the Left choices by his Pursuer opponent is one-half. (In fact, it is slightly less than one half, since there is a non-zero probability they will play Left the same number of times.) Thus, the number of pairs in which the Evader will play left more often than the Pursuer has a binomial distribution $Pr(x; n, p)$, where n is the number of pairs in which the Evader and Pursuer play Left an unequal number of times and $p = 1/2$ (the probability under the null hypothesis). In our data 35 of the 40 Evaders played Left more than their opponents; the probability of 35 or more such “successes” is 6.9×10^{-7} under the null hypothesis, a very strong rejection. It is clear that asymmetric role behavior is a statistically significant feature of our data.

3.3 Behavior Across the Two Games

Comparing play in the Stochastic Game with play in the Deterministic Game reveals another systematic feature of the subjects’ behavior that is not accounted for by equilibrium theory. Here we see a *game asymmetry effect*: despite the fact that both games have the same profile of equilibrium mixtures, both the Pursuers and the Evaders played Left (on average) more often in the Stochastic Game than in the Deterministic Game. Table 2 treats the two games separately but aggregates across all pairs in each game. According to Table 2, the Pursuers played Left 61% of the time in the Deterministic Game and more than 67% of the time in the Stochastic Game. The Evaders played Left less than 72% of the time in the Deterministic Game and more than 76% of the time in the Stochastic Game.

The game asymmetry effect is not as striking as the asymmetric role behavior described above, nor as uniform, but it is nevertheless statistically convincing. Using the data in Table 2, we can test whether the same multinomial process generated the data in both games – *i.e.*, whether the joint distribution that governed the subjects’ play in the Stochastic Game was the same as the distribution that governed the subjects’ play in the Deterministic Game. The homogeneity test statistic is distributed chi-square with three degrees of freedom under the hypothesis that the two games’ distributions of play were the same. The statistic’s value for the data in Table 2 is 36.437, for which the p -value is less than 10^{-7} .

3.4 Serial Correlation

In many previous experiments with mixed-equilibrium games,⁸ the subjects have exhibited significant negative serial correlation: they tended to change their actions too often to be consistent with truly random play. A similar behavioral phenomenon – subjects switching too often – has been consistently observed by experimental psychologists when they have asked subjects to simulate random sequences.⁹ But the subjects in our experiment exhibited the opposite behavior: they tended to switch between Left and Right much less than they would have done if they had been behaving randomly.

In order to evaluate whether the experimental subjects were changing their actions from play to play either too much or too little to be consistent with random play, we conduct a runs test for serial independence on each subject’s sequence of choices. The test’s null hypothesis assumes that a subject’s choices are generated by a fixed mixed strategy which is independent across time. Larger realized values of the test statistic (and thus larger p -values) are produced by sequences with more (and thus shorter) runs – *i.e.*, sequences with more switching of choices.

Table 3 reports the test statistic’s p -value for each of our subjects. With a two-sided hypothesis test and a 5% significance level, we reject the assumption of serial independence for 31 of our 80 subjects: these are the bold-face entries in Table 3. There are 27 subjects for whom rejection is the product of too little switching and only four for whom rejection is the result of too much switching. (Note that if all the subjects had truly been playing randomly, the expected number of subjects for whom this 5% test would yield a rejection is four – two subjects on average would appear to switch too often and two too rarely.) Overall, then, play for many of our subjects was characterized overwhelmingly by positive serial correlation.

⁸For example, O’Neill (1987), Mookherjee and Sopher (1994), Rapoport and Boebel (1992).

⁹See, for example, Tune (1964), Wagenaar (1972) and Lopes (1982). Rapoport and Budescu (1992) found, however, that negative serial correlation was less pronounced in subjects’ action choices in a mixed-equilibrium game than in parallel simulation tasks.

4 Alternative Models

Many previous studies¹⁰ have demonstrated the inability of mixed strategy Nash equilibrium to provide a comprehensive description of human behavior in laboratory settings. Alternative models have therefore been proposed to replace Nash equilibrium as positive theories of game playing. We use our experimental data to estimate three prominent alternative models, and then to evaluate the models' success in explaining the anomalies in our data.

4.1 Quantal Response Equilibrium

The Quantal Response Equilibrium (QRE) introduced by McKelvey and Palfrey (1995) has successfully organized the aggregate empirical frequencies of action choices in many experimental games.¹¹ Like Nash equilibrium, it is a static concept, or more specifically, a parametrized family of static equilibrium models.

The key assumption of the QRE model is that each player calculates the expected value of each of his available actions given his (correct, in equilibrium) belief about his opponent's choice probabilities, and that he attempts to respond optimally, but he makes random errors in the process. In our games this assumption implies that the Pursuer calculates the expected utility of playing Left and Right as follows, assuming that the Evader plays Left with probability q :

$$E_L(q) = q + \varepsilon_{PL} \quad \text{and} \quad E_R(q) = 2(1 - q) + \varepsilon_{PR},$$

where ε_{ij} is a random payoff disturbance for player i and action j . If the $\{\varepsilon_{ij}\}$ are independent random variables having identical Log-Weibull distributions with variance parameter $1/\lambda$, which is the hypothesis of the Logistic QRE, then the Pursuer will play Left with probability

$$p = \frac{e^{\lambda q}}{e^{\lambda q} + e^{2\lambda(1-q)}}.$$

Similar calculations yield the Evader's probability of Left:

$$q = \frac{e^{-\lambda p}}{e^{-\lambda p} + e^{-2\lambda(1-p)}}.$$

¹⁰For example, Brown and Rosenthal (1990), Rapoport and Boebel (1992).

¹¹For example, McKelvey, et.al. (2000), Anderson, et. al. (1998), Capra, et. al. (1999).

For any fixed value of λ , a Logistic QRE (*i.e.*, an LQRE) for the Pursuer-Evader game is a mixed strategy pair $(p^*(\lambda), q^*(\lambda))$ satisfying

$$p^*(\lambda) = \frac{e^{\lambda q^*(\lambda)}}{e^{\lambda q^*(\lambda)} + e^{2\lambda(1-q^*(\lambda))}} \quad \text{and} \quad q^*(\lambda) = \frac{e^{-\lambda p^*(\lambda)}}{e^{-\lambda p^*(\lambda)} + e^{-2\lambda(1-p^*(\lambda))}}.$$

Notice that as λ goes to zero the variance of the errors goes to infinity, and the LQRE has each player playing each action with equal probability; and as λ goes to infinity, the variance of the errors goes to zero, and the LQRE coincides with the Nash equilibrium.

Figure 3 depicts the graphs of the $p^*(\cdot)$ and $q^*(\cdot)$ functions that uniquely solve the pair of equations above, *i.e.*, it depicts the QRE for each value of λ . Note that $p^*(\lambda) < q^*(\lambda)$ for all $\lambda \in (0, \infty)$: the $p^*(\cdot)$ graph lies significantly below the $q^*(\cdot)$ graph, so the LQRE robustly (and strikingly) predicts the role asymmetry effect.

Next we estimate separately, for each of our two games, the value of λ that maximizes the likelihood of the data from that game.¹² These estimated λ -values, along with their corresponding LQRE choice probabilities of Left play – *i.e.*, the model’s predictions – are presented in Table 4. All four predictions are consistent with two of the features we’ve identified in the data: each player in each game is predicted to play Left substantially more often than he plays Right, as the subjects did (on average), and as Nash equilibrium also predicts; and the Evader is predicted to play Left substantially more than the Pursuer, as nearly all of our subject pairs did, and which the Nash equilibrium does not predict.

Note, however, that all four predictions of Left play are nevertheless significantly lower than the corresponding actual play of Left by our subjects: .568 predicted vs. .675 actual for Pursuers in the Stochastic Game, and .698 predicted vs. .762 actual for the Evaders; and .543 predicted vs. .610 actual for the Pursuers in the Deterministic Game, and .692 predicted vs. .717 actual for the Evaders. For each game, the multinomial distribution over joint actions predicted by the maximum likelihood LQRE is overwhelmingly rejected by the chi-square goodness-of-fit test, as reported in Table 4.

Table 4 also presents the log-likelihood values for our data, for each game, under the maximum-likelihood QRE and under the Nash equilibrium. The Nash equilibrium model is nested within the QRE model (the Nash model corresponds to $\lambda = \infty$), so one could conduct a log-likelihood test of

¹²In each game’s data set, each of the $14 \times 100 + 6 \times 200$ observations was weighted equally.

one model against the other. Given the difference in their log-likelihood values, such a test would clearly favor the QRE model. But the statistical validity of this test would be questionable: the Nash model lies on the boundary of the QRE parameter space.

A QRE explanation of the game asymmetry effect would require that a single value of the parameter λ yield different play in our two games, but that cannot happen: the QRE is defined in terms of a game's normal form, and our two games have the same normal form. The QRE also cannot account for the serial correlation in our data, because it is a static equilibrium concept.

4.2 Learning Models

Another class of alternative models assumes adaptive behavior by individuals as they repeatedly play a game. The motivation for these models comes partly from the results of laboratory experiments, in which behavior is typically not stationary. (The presence of nonstationarities in our own data is suggested, for example, by Figure 2, which depicts the data in blocks of twenty stages: in seven of the eight series depicted, there are upward trends; these are most pronounced in the two left panels, which account for most of the data.) We consider two such learning models, the Reinforcement Learning model introduced by Erev and Roth (1998) and the Experience Weighted Attraction (EWA) model introduced by Camerer and Ho (1999).

4.2.1 Reinforcement Learning

Erev and Roth introduced a simple three-parameter reinforcement-learning dynamic which they used to predict play in games that have a unique mixed-strategy equilibrium. In their model each player maintains a score for each of his actions as the game is played repeatedly. At each play, the player selects actions probabilistically; actions with a higher current score are chosen with higher probability. Specifically, let $R_{ij}(t)$ denote player i 's score for his j th action prior to the game at stage t ; let $\sigma_{ij}(t)$ denote the probability that i chooses j at stage t ; and let X_i denote the set of player i 's possible stage-game payoffs. The two initial conditions for the dynamical system are (a) that at the initial stage, each of a player's actions has the same probability of being selected (thus, in our setting, $\sigma_{ij}(1) = .5$ for each player i and each action j) and (b) that

$$R_{ij}(1) = \sigma_{ij}(1)S(1)\overline{X}_i,$$

where $S(1)$ is an unobservable strength parameter and \overline{X}_i is the average of player i 's payoffs across all action profiles.

After the play at each stage, each score is updated according to the rule

$$R_{ij}(t+1) = (1 - \phi)R_{ij}(t) + ((1 - \varepsilon)I_{(a_i(t)=j)} + \frac{\varepsilon}{2})(\pi_i(j, a_{-i}(t)) - \min\{X_i\}),$$

where ϕ is an unobservable parameter that discounts past scores; where $I_{(a_i(t)=j)}$ is an indicator function for the event that player i selected action j in period t ; where the unobservable parameter ε determines the relative impacts on the scores of the selected vs. the unselected action; and where $\pi_i(j, a_{-i}(t))$ is i 's payoff when he plays action j against the opponent's stage- t action $a_{-i}(t)$. (Player i 's minimum possible payoff for any action profile, $\min\{X_i\}$, is subtracted from $\pi_i(j, a_{-i}(t))$ for normalization purposes and to avoid negative scores.) The model is completed by assuming that

$$\sigma_{ij}(t) = \frac{R_{ij}(t)}{\sum_k R_{ik}(t)}.$$

Notice that the range of X_i values varies across the two games even for fixed i : the different monetary rewards and the move by nature both enter explicitly in the transition function for the Stochastic Game.

We use the same estimation procedure used by Erev and Roth (1998): we estimate the values of the three unobservable parameters $S(1)$, ϕ , and ε by minimizing the mean square error of the model's predicted proportions of Left play in 20-period blocks, as depicted in Figure 4.¹³ More specifically, for each fixed triple of parameter values from a discrete grid, we proceeded as follows: For each of the two games we simulated the play of 500 pairs playing 200 stage games, and then we calculated separately the frequency of Left play by the 500 Pursuers and by the 500 Evaders in each 20-period block. These frequencies are the model's predictions for that triple of parameter values. The true values of the parameters should of course not depend upon which game the subjects are playing,¹⁴ so the parameters were estimated using the data from both games together. Within the three-dimensional parameter-value grid, the triple of parameter values that minimizes the mean

¹³Because there were only six pairs in each of the later time blocks, and twenty pairs in each of the early blocks, we weighted each mean square error by the square root of the number of pairs observed for that block.

¹⁴But because the realized monetary payoffs that are used to update the scores R_{ij} are generated differently in our two games, the same parameter values will typically generate different predictions in the two games.

square error, given the data, turns out to be

$$(S(1), \phi, \varepsilon) = (3, .47, .09).$$

In the center panels of Figure 4 we present the average of 10,000 simulations using these estimated parameter values. It is evident in Figure 4 that the model does predict that the Evaders will play Left more often than the Pursuers, as our subjects did. Comparison of the left and center panels of Figure 4 (the data, and the model’s predictions) reveals that the model underpredicts the magnitude of the asymmetry, just as the QRE model does. This is particularly so in the first 5 time blocks, which contain most of the data.

The model does not account for the game asymmetry effect: indeed, the model actually makes the opposite prediction, that Left will be played more frequently in the Deterministic Game than in the Stochastic Game. With regard to the data’s dynamic features, comparison of the left and center panels of Figure 4 suggests that the model might be consistent with the upward trend of Left play in the Deterministic Game, but it does not capture the (stronger) upward trend in the Stochastic Game.

Using the parameter values we have estimated, the Erev and Roth model generates positive serial correlation in action choices: whenever an action is chosen its score is increased by more than the unchosen action’s score, and its probability of being selected in the next stage game therefore increases. This prediction by the model loses some of its attractiveness, however, when we note that the model also predicts positive serial correlation in many mixed-equilibrium games in which subjects actually exhibit negative serial correlation.

4.2.2 The Experience Weighted Attraction Model

The second dynamic model we consider is EWA. In this model subjects choose stage-game actions probabilistically according to the expression

$$\sigma_{ij}(t) = \frac{e^{\lambda R_{ij}(t)}}{\sum_k e^{\lambda R_{ik}(t)}},$$

where at stage t player i chooses action j with probability $\sigma_{ij}(t)$; where λ is a parameter that plays the same role here as in the Logistic QRE model; and where $R_{ij}(t)$ is a scoring function, as in the reinforcement model, but is defined (i.e., updated) differently than in the reinforcement model. The updating of $R_{ij}(t)$ involves a “discounting” factor $N(t)$, which is updated according to

$$N(t+1) = \rho N(t) + 1 \text{ for } t \geq 1,$$

where ρ is an unobservable discount parameter and $N(1)$ is an unobservable parameter, interpreted as the strength of experience prior to the beginning of play. The score $R_{ij}(t)$ is then updated as follows:

$$R_{ij}(t+1) = \frac{N(t)\phi R_{ij}(t) + ((1-\varepsilon)I_{(a_i(t)=j)} + \frac{\varepsilon}{2})\pi_i(j, a_{-i}(t))}{N(t+1)},$$

where $\pi_i(j, a_{-i}(t))$, ϕ , and ε have the same interpretations as in the model of Erev and Roth. The initial scores, $R_{ij}(1)$ for each i and j , are additional unobservable parameters.

The EWA is a flexible specification that includes several other models as special cases. For example, simple reinforcement learning is generated when $N(1) = 0$, $\varepsilon = 0$, and $\rho = 0$; and probabilistic fictitious play is generated when $\varepsilon = \rho = \phi = 1$.¹⁵

As we did in our reinforcement-learning estimation, we estimated the EWA model from the entire data set, a process in which given parameter values generate distinct predictions for the Stochastic and Deterministic Games. The data from our experiment yield the following maximum-likelihood estimate of the nine-tuple of unobservable EWA parameters:¹⁶

$$\begin{aligned} & (N(1), R_{PL}(1), R_{PR}(1), R_{EL}(1), R_{ER}(1), \rho, \phi, \varepsilon, \lambda) \\ & = (.833, .657, 0, -1.863, 0, .993, .998, 1, .578). \end{aligned}$$

Note that the estimated model is very close to probabilistic fictitious play.

We used our estimated EWA model to simulate 10,000 pairs playing 200 periods of each game. The rightmost panels in Figure 4 report the simulations' averages in each block of twenty periods. First note that the model's predictions quickly move toward the Nash equilibrium, but that the asymmetric role behavior does appear: the predictions consistently have the Pursuer playing Left slightly less than the Evader. A comparison of the upper and lower right-hand panels reveals that the model does not predict the game asymmetry effect.

¹⁵We refer the reader to Camerer and Ho (1999) for more discussion of how EWA can emulate various models and for a more complete interpretation of the parameters.

¹⁶ $R_{PR}(0)$ and $R_{ER}(0)$ are arbitrarily fixed at zero. This is a normalization.

As in the Erev and Roth model, with our estimated parameter values the EWA model generates positive serial correlation in the players' choices, but the same caveat applies here as in the Erev and Roth model: the EWA model predicts positive serial correlation even in many games in which subjects actually exhibit negative serial correlation.

Summarizing, each of the three alternative models generates, at least qualitatively, the asymmetric role behavior observed in our data: the Evaders consistently playing Left more often than the Pursuers. In the QRE model this effect is present for all possible values of the model's parameters; in the learning models the effect is less pronounced than in the QRE model, and its presence depends upon the values of the models' parameters. In all three models the magnitude of the effect is considerably less than what is observed in the data. None of the models captures the game asymmetry effect, and while the two learning models do predict positive serial correlation, we have indicated why this prediction is not very convincing.

5 Discussion

If the three alternative models we have considered aren't able to rationalize the data generated by our experiment, then how do we explain the behavior we have observed? We provide several suggestions, none of which is a fully articulated theory or model, but each of which may be playing some kind of role in our experiment.

5.1 Risk Attitudes

Perhaps the most important insight that led O'Neill to reject prior mixed-strategy experiments was that the experiments used games with three or more possible outcomes, and the results of the experiments could therefore have been affected by the subjects' risk attitudes. Like O'Neill's game, our Stochastic Game has only two possible outcomes for each player, and play should therefore be unaffected by risk attitudes. The Deterministic Game, on the other hand, has three outcomes, so risk attitudes might play a role in that game.

Incorporating risk attitudes into the Deterministic Game is straightforward. Since for each player there are only three possible outcomes, a player's preferences over lotteries among the outcomes can be completely described by his relative utility for the intermediate outcome (i.e., the

outcome +1 for the Pursuer and -1 for the Evader). To make things precise, normalize each player's utility so that his worst outcome (0 for Pursuer, -2 for Evader) gives him utility 0 and his best outcome (+2 for Pursuer, 0 for Evader) gives him utility 2. Let π_P denote the utility the Pursuer obtains from his intermediate outcome, +1; and let π_E denote the utility the Evader obtains from his intermediate outcome, -1 . Then a player is risk-neutral over his three possible outcomes if $\pi = 1$; he is risk averse if $\pi > 1$; he is risk-preferring if $\pi < 1$; and in general he is more risk averse as the value of his π is larger.

It is easy to verify that the Pursuer's and Evader's Nash equilibrium mixtures place probability

$$\sigma_P = \frac{2}{4 - \pi_E} \quad \text{and} \quad \sigma_E = \frac{2}{2 + \pi_P}$$

on play of Left. If each player is risk-neutral, then we have $\sigma_P = \sigma_E = 2/3$, as we already know. If both are risk averse, then we have $\sigma_E < 2/3 < \sigma_P$. But in virtually every one of our subject pairs the Evader played Left more often than the Pursuer. So risk aversion is simply not consistent with our subjects' behavior. Indeed, the data suggest that $\sigma_P < \sigma_E$.

Prospect Theory (Kahneman and Tversky (1979)) holds that while people are indeed risk averse for gains, they are instead *risk preferring* when facing losses, or "loss averse." This would mean that the π -values of players in the Deterministic Game would satisfy $\pi_P > 1$ and $\pi_E < 1$, and the equilibrium in the Deterministic Game would therefore satisfy the inequalities $\sigma_E < 2/3$ and $\sigma_P < 2/3$. Since the equilibrium in the Stochastic Game is unaffected by risk attitudes, each player in that game uses a mixture that places $2/3$ probability on playing Left. Thus, if people's risk attitudes are correctly described by Prospect Theory, then both Pursuers and Evaders should (in equilibrium) play Left less often in the Deterministic Game than in the Stochastic Game – precisely the game asymmetry effect that appears in our data. However, average play of Left was *more* than $2/3$ by subjects playing both roles in both games.

Prospect Theory holds additionally that loss aversion is stronger than risk aversion in people's risk attitudes. For the payoffs in the Deterministic Game this is equivalent to the inequality $1 - \pi_E > \pi_P - 1$, which is in turn equivalent to $\pi_P + \pi_E < 2$, and the latter inequality is equivalent to $\sigma_P < \sigma_E$. Prospect Theory therefore provides, at least qualitatively, an equilibrium explanation of the observed asymmetric role behavior in the Deterministic Game. Furthermore, its explanatory power could be tested by modifying our design so that both Pursuer and Evader are always paid nonnegative amounts (cf. footnote 5).

Risk attitudes and Prospect Theory are clearly not the whole story, however. They do nothing to account for the strong asymmetric role behavior in the Stochastic Game. Further, a significant number of the Evader subjects in the Deterministic Game played Left on substantially more than 2/3 of their plays, which requires that $\pi_P < 1$ – i.e., it requires that these Evaders believed (correctly) that their opponents were risk preferring for gains, while the other Evader subjects believed (also correctly) that their opponents were risk averse for gains. This seems unlikely. So while risk attitudes may have contributed to the observed behavior, they seem to fall short of fully accounting for the behavior.

5.2 Objective vs. Subjective Uncertainty

In the Introduction we indicated that since the Stochastic Game is simpler than the games typically used in previous two-outcome mixed strategy experiments (each player has only two pure strategies instead of four or more), the theory might be expected to perform better here than in the previous experiments. But the Stochastic Game might not actually be so simple. In two of the four cells of the payoff matrix the outcome is a lottery, an uncertain prospect. So while the game has far fewer strategy combinations than a 4×4 or larger game, the reduction in the strategy sets might have been more than offset by replacing outcomes that are certain with ones that are uncertain.

For example, in some of our experimental sessions we asked the subjects, upon completing the experiment, to write down for us any strategy(s) they were using as they played the game. Several of the subjects who played the Stochastic Game indicated that they were paying more attention to the behavior of the die than to their opponent’s behavior. One subject (an Evader) wrote “I tried to guess the probability of when a 1 or 2 was ‘due’ up on the dice, and sneak a right card in. It is not very easy to figure out when the other person is going to change their ‘favorable’ left or right card.” A Pursuer subject wrote “Sometimes I would try to guess what the die roll would be, and didn’t even think about what the evader was going to do.” Altogether, nearly half of the Stochastic Game subjects indicated that on at least some of their plays they were trying to predict the outcome of the die roll.

These comments suggest that some of the subjects did not view the die rolls as a serially independent random process. While it’s not likely that they treated the die, or Mother Nature, as literally a third player in the game, they may have viewed the opponent’s actions and the die’s

outcomes as similar processes, each generating unpredictable “actions” that affected their own payoff.

How would this view of the lottery outcomes affect behavior? We know from the Ellsberg Paradox, and similar examples, that people are typically more comfortable with objective risks than with subjective ones. Of the two kinds of uncertainty faced by a subject in the Stochastic Game, uncertainty about what one’s opponent will do is surely the more subjective, and uncertainty about the die roll more objective. A subject might therefore view the die-uncertainty as more “focal” and devote more of his attention to “playing against the die” than to playing against his opponent – exactly the approach described by the subjects quoted above.

Consider the extreme situation in which the players in the Stochastic Game are playing *only* against the die, and placing no weight on the opponent’s likely action. Then a Pursuer would play Right, and an Evader would play Left. This is of course not an equilibrium (at least not a Nash equilibrium): while the Pursuer is using his best response to the die’s “mixture,” he is using his *worst* response to the Evader’s choice of Left. The Evader’s play of exclusively Left, on the other hand, is a best response to both the die’s “mixture” *and* his opponent’s play.

Because the Pursuer does so badly here, this pattern of play could not be expected to persist, and indeed we observe every Stochastic Game Pursuer play Left on at least 40% of his plays. But the additional uncertainty of the die in the Stochastic Game, which is more objective than the strategic uncertainty about one’s opponent, could be expected to reinforce a tendency toward Left play by Evaders. Moreover, greater Left play by Evaders in the Stochastic Game should eventually bring forth greater Left play by Pursuers as well, as in the game asymmetry effect that our data exhibits.

5.3 Initial Beliefs and the Dynamics of Play

Each of the subjects in our experiment was surely a novice at the game he or she was playing: none of the subjects had seen or played either game before. So there is no reason to expect equilibrium play at the outset of a match. This leads us to consider how play might progress as players play the game repeatedly. We have already discussed reinforcement and EWA learning models, each of which predicts a path of play that depends partly upon assumptions about the model’s initial conditions. Here we focus directly on the idea that players respond to their beliefs about how their

opponents will play, and we ask what players' initial beliefs might have been and how they might have influenced observed play.

In belief-based learning models¹⁷ a player has a *belief* about how his opponent is going to play, and chooses actions for himself that he thinks will do well against his opponent's play. As play progresses, a player observes his opponent's actual play and uses that information to update or revise his belief, which produces a corresponding change in his own play in order to do well against the revised belief.¹⁸

The left panels of Figure 4 show the time series of play, aggregated across all subjects, and aggregated across time into twenty-play blocks. In the earliest stages of play the Evader subjects, on average, played Left somewhat more often than the equilibrium prescribes, while the Pursuer subjects played Left less often than called for in equilibrium. This pattern of play at the outset would be consistent with beliefs by Evader and Pursuer subjects alike that their opponents would play Left less often than the 2/3 rate prescribed by equilibrium.

Can we say anything directly about the beliefs our subjects actually did hold at the outset of an experimental session? We did not ask our subjects for any information prior to their play. But the beliefs of another, similar set of subjects provide some evidence on this question. One of the authors presented the Stochastic Game instructions to the 44 students in his game theory course. This was done at the first meeting of the course, before the students had learned anything about game theory; thus, they were drawn from roughly the same subject pool as the subjects in our experiment. The students were asked to predict the average number of Left plays, out of 100 plays, by people who had played this game in an experiment. (A five dollar reward was promised for the prediction that came closest to the actual average of play.) The average prediction about Evader play was 58 Left plays out of 100; the average prediction for Pursuer play was 49 Left plays out of 100. Only five of the 44 student respondents predicted more than 67 Left plays by Pursuers, and only eight respondents predicted more than 67 Left plays by Evaders.

Thus, it seems very likely that the predominant belief of our experimental subjects was that

¹⁷For example, Milgrom and Roberts (1991) and Fudenberg & Levine (1998).

¹⁸The belief might be a simple forecast of his opponent's action, or a mixture he expects his opponent to use, or even a belief about the opponent's characteristics, or "type," from which one would infer what action the opponent is likely to take. Playing well against the opponent's expected play might consist of best responding to it, or it might instead be some sort of stochastic play.

one's opponent would play Left less than $2/3$ of the time. As we have described above, this would lead, in a belief-based model of play, to play in the early stages that would be similar to our data.

Subsequent belief-based adjustment by the players would lead Evaders to increase their rate of Left play (because Pursuers are playing Left too little), and it would lead Pursuers to increase their rate of Left play as well (because Evaders are playing Left too often). Figure 2 shows that play did progress in this way. In the Stochastic Game, Evaders quickly increased their rate of Left play to levels well above $2/3$, and Pursuers followed by increasing Left play to about $2/3$ and then above $2/3$. This should have induced Evaders to reduce their rate of Left play, and indeed Figure 3 shows that they did. In the Deterministic Game, Evaders' rate of Left play increased more slowly, and Pursuers' Left play also increased more slowly, remaining below $2/3$ for much longer than in the Stochastic Game. But eventually in the Deterministic Game as well, Pursuer play reached and even exceeded $2/3$, and Evader play then decreased somewhat.

Thus, the progress of play, given the subjects' initial beliefs about their opponents' likely play, is at least qualitatively consistent with belief-based learning, even if this is not captured very well by the belief-based component of the EWA model.¹⁹ And this path of play exhibits the role asymmetry effect observed in our data, albeit as a transitory effect: as Pursuers come to play Left more often than $2/3$ of the time, Evaders should reduce their Left play below $2/3$. This explanation, however, mainly serves to focus attention on the following question: why are subjects' initial beliefs about their opponents' play skewed toward equally likely play of Left and Right? We don't have an answer to that question.

This discussion also leads, of course, to the question whether all this belief modification will eventually settle down and produce play that is closer to the equilibrium. That was the reason we ran sessions with 200 rounds of play, and those sessions produced no such evidence. But this is not really a question that can be answered definitively on purely empirical grounds: no matter how many repetitions of play we observe, we cannot be certain that play will exhibit the same characteristics if it continues for further repetitions. The best we could hope for in this regard

¹⁹There is perhaps some additional evidence that our subjects were responding directly to the play of their opponents (rather than responding to it indirectly, via the reinforcement their own earnings provided). About half of the responses to our debriefing question indicated the respondent was trying to anticipate his opponent's play, or to outwit his opponent, or that he had played as he did in order to do well against the play he had observed by his opponent.

would be a theory of disequilibrium play, or learning, for which our data provide support. As we have seen, however, our data lend little support to current theories.

5.4 Offense vs. Defense, and Serial Correlation

When a player randomizes independently at each play in a repeated game, as required in a mixed strategy Nash equilibrium, we might say that he is playing *defensively*: The only reason for a player to randomize in this way, and to use the equilibrium mixture rather than some other mixture over the same pure strategies, is to avoid being exploited. This requires that a player be unpredictable, and that he be unpredictable in a specific way. By contrast, if I believe that my opponent is not playing according to the equilibrium prescription – if he is predictable, for example, or even if he is unpredictable but at the wrong rate – then I can take advantage of that. We could refer to this as playing *offensively*.

In their responses to our debriefing question, only 10 subjects (out of 37) made any reference at all to the need to be unpredictable. More than twice as many of the subjects described trying to outthink or outwit their opponent, or trying to respond to (i.e., take advantage of) the opponent’s play, or trying to predict the die roll and respond to that.

Now consider the one unexpected feature of our data that we haven’t yet addressed: the strong positive serial correlation in most of our subjects’ choices. This is in sharp contrast with evidence from experimental psychology that when asked to simulate random processes (such as tosses of a fair coin), people exhibit a strong tendency to “switch too often” – i.e., they exhibit strong *negative* serial correlation. But perhaps it is not such a contrast with the psychology experiments after all. If most of our subjects were, as they indicated, predominantly playing offensively, instead of trying to be unpredictable, then there is no reason to expect them to exhibit either independence or the classic negative serial correlation of the psychology experiments.

We might then ask why offensive play would exhibit positive serial correlation. It is implicit in our description of offensive play that the player is responding to his belief about his opponent’s play. If this belief does not change appreciably from one play to the next, then the offensive player (who sees little need to be unpredictable) will change his action only rarely. Indeed, some of the subjects who played the Stochastic Game eventually played Left almost exclusively.

This is clearly an incomplete explanation of positive serial correlation, and it offers no clues to

why other experiments, such as O'Neill's, have generated play that is negatively serially correlated. Some additional theorizing along these lines would perhaps be fruitful.

5.5 Concluding Remarks

We have described several puzzling features of the data generated by our experiment, but we have not identified or developed a comprehensive theory that accounts for them all. Perhaps future experiments will provide sharper versions of these phenomena, which will in turn suggest a unified theory. Or perhaps one of our readers will see a pattern in this data set that has eluded us and that will reveal all.

Appendix: Experimental Instructions

The instructions to subjects for the Stochastic Game were as follows. The instructions for the Deterministic Game were nearly identical.

You and the person seated next to you will repeatedly play a simple game of pursuit and evasion, or “hide and seek.” If you are sitting at the right-hand side of your table, you are the pursuer; if you are on the left-hand side you are the evader. You will play to win tokens, and each token will be redeemed for fifty cents at the end of the experiment. Each game will begin with the evader placing one of his tokens in the center of the table. The pursuer then tries to “find” the evader and to “capture” his token. At the beginning of the experiment the evader will be given 50 tokens (worth \$25). During the experiment the pursuer will capture some of the evader’s tokens.

The game proceeds as follows. The evader begins by placing a token in the center of the table. Then each of you will make a private choice to move either Left or Right, by placing either your Left card or your Right card face down in front of you (out of sight of the other player because of the partition).

After both of you have played your cards, the monitor (seated in front of you) will randomly select a number from 1 to 6 by rolling a die. Note that each of the six numbers is equally likely to occur.

Next, in sight of both players, the monitor will turn over the cards both players played. If the cards do not match (i.e., if one is Left and one is Right), then the evader has succeeded: he has not been found by the pursuer, and the evader does not have to give up his token. On the other hand, if the cards DO match (i.e., if both are Left or both are Right), then the pursuer has FOUND the evader, but whether he actually captures the token depends upon the die number that occurred. The result is determined as follows:

	<u>If die is 1 or 2</u>	<u>If die is 3,4,5, or 6</u>
If both moved Left:	Pursuer captures token	Evader keeps token
If both moved Right:	Evader keeps token	Pursuer captures token
If one Left, one Right:	Evader keeps token	Evader keeps token

When play is complete and the token has been either captured or retained, the monitor will record the moves and the die number; the players’ cards will be returned to them; and a new play will begin.

You will play the game 100 times, unless the evader loses all his tokens, in which case all play will be terminated. When all play has ended, the experiment will be over, and you will redeem your tokens for cash – a dollar for every two tokens you have.

References

- Simon P. Anderson, Jacob K. Goeree, and Charles A Holt: “Rent seeking with bounded rationality: An analysis of the All-Pay Auction”, *Journal of Political Economy*, 106 (1998), 828–853.
- Ken Binmore, Chris Proulx, and Joseph Swierzbinski: “Does minimax work? An experimental study”, *Economic Journal*, forthcoming.
- Robert Bloomfield: “Learning a mixed strategy equilibrium in the laboratory”, *Journal of Economic Behavior and Organization*, 25 (1994), 411–436.
- James Brown and Robert Rosenthal: “Testing the Minimax Hypothesis: A re-examination of O’Neill’s game Experiment”, *Econometrica*, 58 (1990), 1065–1081.
- Colin Camerer and Teck-Hua Ho: “Experience-weighted attraction learning in normal form games”, *Econometrica*, 67 (1999), 827–874.
- Monica C. Capra, Jacob K. Goeree, Rosario Gomez, and Charles A Holt: “Anomalous behavior in a Traveler’s Dilemma?”, *American Economic Review*, 89 (1999), 678–690.
- Ido Erev and Alvin E. Roth: “Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria”, *American Economic Review*, 88 (1998), 848–881.
- Drew Fudenberg and David Levine: *Theory of Learning in Games*, MIT Press, 1998.
- Daniel Kahneman and Amos Tversky: “Prospect theory: An analysis of decision under risk”, *Econometrica*, 47 (1979), 263–292.
- L. Lopes.: “Doing the impossible: A note on induction and the experience of randomness”, *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 8 (1982), 626–636.
- Richard D. McKelvey and Thomas R. Palfrey: “Quantal response equilibrium for normal form games”, *Games and Economic Behavior*, 10 (1995), 6–38.
- Richard D. McKelvey, Thomas R. Palfrey, and Roberto A. Weber: “The effects of payoff magnitude and heterogeneity on behavior in 2 x 2 games with unique mixed strategy equilibria”, *Journal of Economic Behavior and Organization*, 42 (2000), 523–548.

David Malcolm and Berhnardt Lieberman: “The behavior of responsive individuals playing a two-person, zero-sum game requiring the use of mixed strategies”, *Psychonomic Science*, 12 (1965), 373-374.

Paul Milgrom and D. John Roberts: “Adaptive and sophisticated learning in repeated normal form games”, *Games and Economic Behavior*, 3 (1991), 82–100.

Dilip Mookherjee and Barry Sopher: “Learning behavior in an experimental matching pennies game”, *Games and Economic Behavior*, 7 (1994), 62–91.

Dilip Mookherjee and Barry Sopher: “Learning and decision costs in experimental constant sum games”, *Games and Economic Behavior*, 19 (1997), 97–132.

Jack Ochs: “Games with unique mixed strategy equilibria: An experimental study”, *Games and Economic Behavior*, 10 (1994), 202–217.

Barry O’Neill: “Nonmetric test of the Minimax Theory of two-person zerosum games”, *Proceedings of the National Academy of Sciences, U.S.A.*, 84 (1987), 2106–2109.

Barry O’Neill: “Comments on Brown and Rosenthal’s reexamination”, *Econometrica*, 59 (1991), 503-507.

A. Rapoport and R. Boebel: “Mixed strategies in strictly competitive games: A further test of the Minimax Hypothesis”, *Games and Economic Behavior*, 4 (1992), 261–283.

Amnon Rapoport and David V. Budescu: “Generation of random series in two-person strictly competitive games”, *Journal of Experimental Psychology, General*, 121 (1992), 352–363.

Ariel Rubinstein: “Experience from a course in game theory: Pre- and post-class problem sets as a didactic device”, *Games and Economic Behavior*, 28 (1999), 155–170.

Jason Shachat: “Mixed strategy play and the Minimax Hypothesis”, *Journal of Economic Theory*, (forthcoming).

Patrick Suppes and Richard C. Atkinson: *Markov Learning Models for Multiperson Interactions*, Stanford University Press, 1960.

G. S. Tune: “Response preferences: A review of some relevant literature”, *Psychological Bulletin*, 61 (1964), 286-302.

W. A. Wagenaar: “Generaton of random sequences by human subjects”, *Psychological Bulletin*, 77 (1972), 65-72.

John Wooders and Jason Shachat: “On the irrelevance of risk attitudes in repeated two-outcome games”, *Games and Economic Behavior*, 34 (2001), 343–363.

Figure 1: Evader vs. Pursuer Left Play

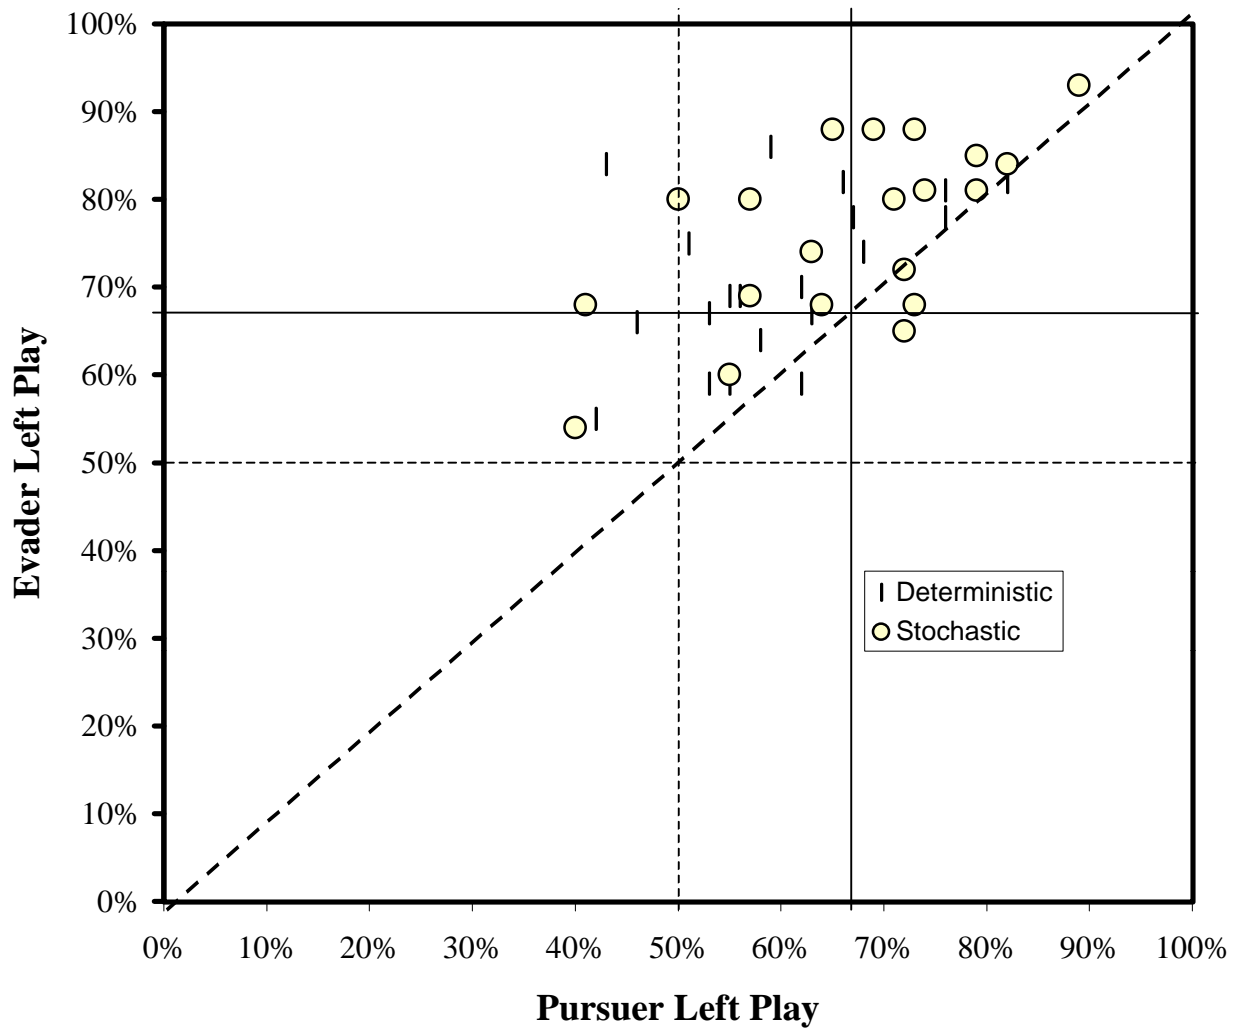

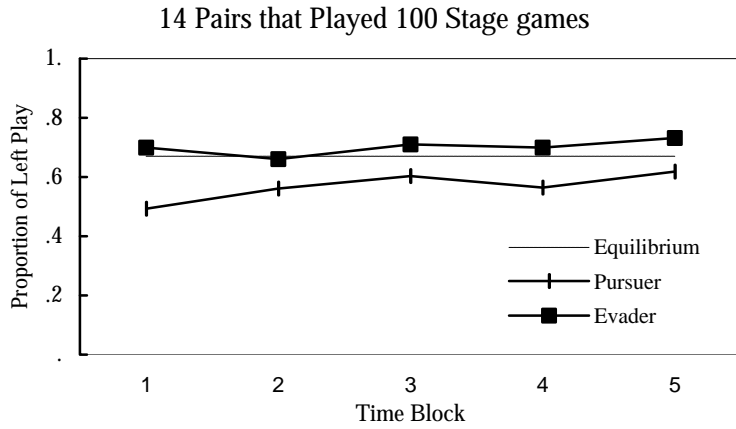


Figure 2
 Times Series of Proportion of Left Play: 100 Periods and 200 Periods
 Deterministic Treatment



Stochastic Treatment





Figure 3: Logistic QRE for Pursuer-Evader Game





Figure 4: Times Series of Proportion of Left Play: Actual and Predicted

Deterministic Treatment



Stochastic Treatment




Table 1
 Pair Level Data Analysis: Stochastic Treatment

Pair #	Stages	Action Profile				NE test		Binomial	Binomial	
		LL	LR	RL	RR	Pur. L	Eva. L	P-value	Pur. L	Eva. L
1	100	.34	.07	.34	.25	.41	.68	.000	.000	.647
2	100	.84	.05	.09	.02	.89	.93	.000	1.000	1.000
3	100	.39	.25	.29	.07	.64	.68	.203	.320	.647
4	100	.34	.16	.46	.04	.50	.80	.000	.000	.999
5	100	.62	.07	.26	.05	.69	.88	.000	.723	1.000
6	100	.71	.08	.14	.07	.79	.85	.000	.998	1.000
7	100	.21	.19	.33	.27	.40	.54	.000	.000	.006
8	100	.30	.25	.30	.15	.55	.60	.028	.010	.097
9	100	.47	.10	.33	.10	.57	.80	.007	.028	.999
10	100	.52	.20	.20	.08	.72	.72	.457	.893	.893
11	100	.41	.16	.28	.15	.57	.69	.181	.028	.723
12	100	.59	.06	.29	.06	.65	.88	.000	.398	1.000
13	100	.71	.11	.13	.05	.82	.84	.000	1.000	1.000
14	100	.64	.10	.17	.09	.74	.81	.001	.954	1.000
15	200	.64	.15	.17	.04	.79	.81	.000	1.000	1.000
16	200	.47	.25	.18	.11	.72	.65	.382	.955	.281
17	200	.46	.17	.28	.10	.63	.74	.123	.121	.985
18	200	.56	.16	.25	.05	.71	.80	.000	.917	1.000
19	200	.48	.25	.20	.07	.73	.68	.191	.977	.680
20	200	.62	.10	.26	.02	.73	.88	.000	.968	1.000

Table 1
 Pair Level Data Analysis: Deterministic Treatment

Pair #	Stages	Action Profile				(Pursuer, Evader)		NE test	Binomial	Binomial
		LL	LR	RL	RR	Pur. L	Eva. L	P-value	Pur. L	Eva. L
21	100	.36	.07	.48	.09	.43	.84	.000	.000	1.000
22	100	.36	.26	.23	.15	.62	.59	.304	.188	.066
23	100	.32	.23	.27	.18	.55	.59	.032	.010	.066
24	100	.56	.20	.22	.02	.76	.78	.013	.984	.995
25	100	.50	.09	.36	.05	.59	.86	.000	.066	1.000
26	100	.55	.11	.27	.07	.66	.82	.013	.481	1.000
27	100	.36	.17	.31	.16	.53	.67	.038	.003	.566
28	100	.47	.15	.23	.15	.62	.70	.274	.188	.791
29	100	.54	.13	.24	.09	.67	.78	.093	.566	.995
30	100	.27	.26	.32	.15	.53	.59	.004	.003	.066
31	100	.41	.14	.28	.17	.55	.69	.047	.010	.723
32	100	.38	.18	.31	.13	.56	.69	.137	.017	.723
33	100	.23	.19	.32	.26	.42	.55	.000	.000	.010
34	100	.31	.15	.35	.19	.46	.66	.000	.000	.481
35	200	.43	.21	.25	.13	.63	.67	.741	.153	.566
36	200	.66	.16	.16	.03	.82	.82	.000	1.000	1.000
37	200	.37	.21	.27	.16	.58	.64	.037	.004	.190
38	200	.62	.14	.19	.06	.76	.81	.000	.997	1.000
39	200	.38	.13	.38	.12	.51	.75	.000	.000	.996
40	200	.49	.19	.26	.07	.68	.74	.128	.625	.990

Table 2
Contingency Tables for Aggregate Data

Stochastic Game

		<u>Evaders' Choices</u>		Marginal Frequencies for Pursuers
		Left	Right	
Pursuers' Choices	Left	.521 (.444) .010	.154 (.222) .008	.675 (.667) .009
	Right	.241 (.222) .008	.085 (.111) .006	.325 (.333) .009
Marginal Frequencies for Evaders		.762 (.667) .009	.238 (.333) .009	

Deterministic Game

		<u>Evaders' Choices</u>		Marginal Frequency for Pursuers
		Left	Right	
Pursuers' Choices	Left	.441 (.444) .010	.169 (.222) .008	.610 (.667) .009
	Right	.276 (.222) .008	.114 (.111) .006	.390 (.333) .009
Marginal Frequencies for Evaders		.717 (.667) .009	.283 (.333) .009	

Table 3
Runs Tests' P-Values

(Boldface indicates a rejection of independence at the 5% level of significance.)

Stochastic Treatment			Deterministic Treatment		
Pair #	Pursuer P-value	Evader P-value	Pair #	Pursuer P-value	Evader P-value
1	.092	.790	21	.993	.664
2	.031	.209	22	.026	.632
3	.662	.850	23	.180	.009
4	.420	.172	24	.000	.953
5	.521	.294	25	.063	.487
6	.014	.840	26	.001	.694
7	.735	.989	27	.003	.657
8	.064	.799	28	.731	.684
9	.846	.172	29	.050	.689
10	.567	.092	30	.993	.551
11	.205	.338	31	.380	.521
12	.018	.000	32	.046	.006
13	.000	.001	33	.604	.016
14	.181	.000	34	.172	.023
15	.057	.007	35	.000	.023
16	.001	.012	36	.000	.045
17	.940	.002	37	.005	.979
18	.004	.588	38	.222	.013
19	.000	.685	39	.002	.935
20	.378	.000	40	.000	.992

Table 4
Logistic QRE estimates and Hypothesis Test

Treatment	MLE λ	Evader's Left Play		Pursuer's Left Play		QRE	NE	chi-square	p-value
		Predicted	Observed	Predicted	Observed	Ln Likelihood	Ln Likelihood	g.o.f stat	
Stochastic	2.84	.698	.762	.568	.675	-2934.25	-3124.57	380.64	0.00
Deterministic	2.19	.692	.717	.543	.610	-3023.91	-3320.73	593.65	0.00