

# Does Observation of Others Affect Learning in Strategic Environments? An Experimental Study\*

John Duffy  
Department of Economics  
University of Pittsburgh  
jduffy+@pitt.edu

Nick Feltovich  
Department of Economics  
University of Pittsburgh  
niklas+@pitt.edu

May 1997

## Abstract

This paper presents experimental results from an analysis of two similar games, the repeated ultimatum game and the repeated best-shot game. The experiment examines whether the amount and content of information given to players affects the evolution of play in the two games. In one experimental treatment, subjects in both games observe not only their own actions and payoffs, but also those of one randomly chosen pair of players in the just-completed round of play. In the other treatment, subjects in both games observe only their own actions and payoffs. We present evidence suggesting that observation of other players' actions and payoffs may affect the evolution of play relative to the case of no observation.

*Journal of Economic Literature* Classification Nos: D83, C72, C73.

---

\*Duffy's research was supported by a grant from the University of Pittsburgh Office of Research. Feltovich acknowledges research support from an Andrew Mellon fellowship. We thank David Cooper, Ido Erev, Ron Harstad, John Kagel, Jack Ochs, Jean-François Richard, Al Roth, an associate editor and two anonymous referees for helpful comments and suggestions.

# 1 Introduction

There is now a large experimental literature that examines learning behavior in repeated non-cooperative games. The kind of learning behavior that has been examined in these experiments has been mainly limited to cases where individuals learn on the basis of their own past experience. A more general analysis of learning behavior would consider whether providing players with additional information affects the evolution of play over time. In particular, one might ask whether changes in the content and quantity of information available to players affects the rate at which they learn to play equilibrium strategies. This paper takes a first step toward addressing these kinds of questions.

We examine the effect of additional information on learning in an experiment involving two similar games, the repeated “ultimatum game” and the repeated “best-shot game.” An experimental comparison of these two games was previously conducted by Prasnikar and Roth (1992). Our experiment consists of two treatments. In the first treatment, subjects are informed of only their own actions and payoffs. This treatment is similar to one found in the Prasnikar and Roth (1992) study as well as in many other experimental implementations of repeated non-cooperative games. The second treatment is more novel. In this treatment, subjects are informed of their own actions and payoffs as well as the actions and payoffs of *one* randomly chosen pair of players in the just-completed round of play. We use these two treatments to examine whether observation of other players’ actions and payoffs affects the evolution of play relative to the (standard) case where observation of others is not possible.

From a game-theoretic perspective, allowing for observation of how another pair of players play the game is a step in the direction of satisfying the “mutual knowledge of the strategy choices” condition, which, given that players play rationally according to the payoff structure of the game, is sufficient to ensure that players play according to the equilibrium predictions of game theory (Aumann and Brandenburger (1995)). Thus, one might suppose that when players are allowed to observe the actions and payoffs of others, they will learn more quickly to play in accordance with the game’s equilibrium predictions relative to the case where players do not get to observe the actions and payoffs of others.

The notion that individuals may learn by observing others is one that economists have only recently begun to recognize.<sup>1</sup> Anthropologists, however, have long noted that in many cultures,

---

<sup>1</sup>See for example, Banerjee (1992), Banerjee and Fudenberg (1995), Bikhchandani, Hirshleifer and Welch (1992), Conlisk (1980), Ellison and Fudenberg (1993, 1995), Jackson and Kalai (1995) and Vives (1996) among others.

observation is the *primary* method through which individuals learn. Reichard (1938), for example, has observed that in many languages “the word for ‘teach’ is the same as the word for ‘show,’ and the synonymy is literal.”<sup>2</sup> Among behavioral psychologists, the hypothesis that individuals learn through observation of others is also well established. Bandura and Walters (1963), summarizing a large body of experimental research, conclude that there are three main effects of observation on learning behavior. First, the observer may learn through observation to play new strategies that were previously unknown. Second, the observer may learn through observation to become less phobic about playing certain inhibitory strategies. Third, observation of others may facilitate the play of strategies the observer already knows, but has a low probability of playing.

In this study, we focus on this third effect of observation. We imagine that when players observe strategies resulting in high payoffs, they are more likely to adopt the observed strategy in subsequent rounds of play. We develop an adaptive learning algorithm based on one proposed by Roth and Erev (1995) that incorporates our assumption about how players use the additional information they receive about other players. We use this algorithm to simulate how play might evolve in our experimental treatments with and without observation. The simulations suggest that the effect of observation on learning behavior is *different* in the two games. In the ultimatum game, players who observe the actions and payoffs of other players tend to move *further away from* the subgame perfect equilibrium prediction over time, relative to players who do not observe the actions and payoffs of others. In the best-shot game, players who observe the actions and payoffs of others tend to *move closer toward* the subgame perfect equilibrium prediction over time relative to players who do not observe the actions and payoffs of others.

The predictions of our adaptive learning model are tested in an experiment with human subjects. We find that in the ultimatum game, players who observe the actions and payoffs of other players do indeed tend to move further away from the subgame perfect equilibrium relative to players who do not observe the actions and payoffs of others. By contrast, in the best-shot game, players who observe the actions and payoffs of other players are only slightly, though not significantly, closer to the subgame perfect equilibrium prediction relative to players who do not observe the actions and payoffs of others. Thus our results concerning the effect of observation on learning are mixed and appear to be dependent upon the game played.

---

<sup>2</sup>Reichard (1938) p. 47, as quoted in Bandura and Walters (1963).

## 2 The Games

The first game we consider is the ultimatum bargaining game in which two players attempt to divide \$10. Player 1 moves first by proposing to keep a nonnegative integer amount  $\$y < \$10$  for herself, leaving Player 2 the residual amount  $\$10 - y$ .<sup>3</sup> Player 2 must then decide whether to accept or reject Player 1's offer. If Player 2 *accepts* Player 1's offer, Player 2 receives the proposed offer of  $\$10 - y$ , while Player 1 receives  $\$y$ . If Player 2 *rejects* Player 1's offer, both players receive zero. While this game has many Nash equilibria, the *unique* subgame perfect equilibrium is for Player 1 to demand \$9 and for Player 2 to accept this demand, thereby earning \$1.

The second game we consider is the “best-shot” public good game of Harrison and Hirshleifer (1989). Player 1 moves first by choosing an integer investment amount  $i_1 \in [0, 9]$ .<sup>4</sup> Player 2 observes the investment amount chosen by Player 1, and then also chooses an integer investment amount  $i_2 \in [0, 9]$ . The project level (the amount of the public good actually furnished) is  $Q = \max(i_1, i_2)$ . Once the project level is determined, each player  $k$  receives payoff  $Q(41 - Q)/20 - 1.65 \times i_k$ .<sup>5</sup> Note that Player 2's best response to a *zero* investment by Player 1 is an investment of 4, and Player 2's best response to a *nonzero* investment by Player 1 is an investment of 0. There are two Nash equilibrium outcomes: a subgame perfect equilibrium outcome in which Player 1 chooses to invest 0 and Player 2 chooses to invest 4, and a non-subgame perfect equilibrium outcome in which Player 1 chooses to invest 4 and Player 2 chooses to invest 0.

---

<sup>3</sup>Note that these rules differ from the rules of other ultimatum game experiments in two respects. First, Player 1 proposals must be in discrete, even-dollar amounts. Second, Player 1s are not allowed to demand the entire \$10 prize. We restrict Player 1s to making even-dollar demands as a means of narrowing the strategy space and making it comparable to the discrete strategy space of the best-shot game; the same restricted strategy space was used in our simulations. We do not allow Player 1's to demand the entire \$10 prize to ensure that the subgame perfect equilibrium is unique.

<sup>4</sup>In Harrison and Hirshleifer (1989), and Prasnikar and Roth (1992) the integer interval for investment choices was  $[0, 21]$ . We chose a shorter investment interval to make the Player 1 strategy space as similar as possible to that in our treatment of the ultimatum game and to reduce the likelihood of negative payoffs in the stage game. In both the Harrison-Hirshleifer payoff tables and those we used, investment amounts above 8 are very likely to result in negative payoffs. Since both Harrison-Hirshleifer and Prasnikar-Roth report few, if any, investment choices above 8, we conjectured that removing some of these choices would not significantly alter the play of this game.

<sup>5</sup>These parameters differ somewhat from those used by Harrison and Hirshleifer (1989); for given investment levels, payoffs are roughly twice what they would be in Harrison and Hirshleifer. We chose these numbers so that equilibrium payoffs in the ultimatum game and best-shot game are of approximately the same magnitude. The numbers we use preserve both the Nash equilibria and the best-response correspondence of Harrison and Hirshleifer.

## 2.1 Previous Experimental Results

In a number of controlled laboratory experiments of the ultimatum game, the subgame perfect equilibrium prediction has been soundly rejected.<sup>6</sup> A typical finding is that Player 1s propose to split the money prize nearly equally and that Player 2s accept this proposal. Player 1s who propose to keep the bulk of the money prize for themselves find that their proposals are often rejected by Player 2s, even though Player 2s always receive a higher monetary payoff by accepting any positive offer from Player 1s. One explanation for the inconsistency of the experimental results with the theory is that players begin play with a prior disposition to play fairly; given some experience and/or additional information they might learn to play the subgame perfect outcome. This explanation, while appealing, is not completely satisfying; it does not offer much guidance as to what kind of experience or knowledge players must acquire or how long it will take players to become experienced or knowledgeable enough to play the subgame perfect equilibrium strategy. Harrison and McCabe (1992), for example, have noted that experimental tests of the ultimatum game ignore the theoretical assumption that players have “common knowledge” of the beliefs, motives and strategies of all other players. However, their implementation of the common knowledge assumption comes at the cost of substantially altering the experimental design. Evidence from other experiments suggests that greater or more varied experience in playing the ultimatum game does not greatly affect players’ tendencies to deviate from the subgame perfect equilibrium (see, *e.g.* Güth, Schmittberger, and Schwarz (1982), Bolton (1991), or Roth *et al.* (1991)).

In contrast, the Harrison and Hirshleifer (1989) best-shot game experiment yielded results that were largely consistent with the subgame perfect outcome (*cf.* Harrison and Hirshleifer’s experiment SQ-3). Prasnikar and Roth (1992) repeated the Harrison–Hirshleifer best-shot game experiment, but gave players “full information” about each other’s payoff functions, in contrast to Harrison and Hirshleifer, who did not explicitly tell subjects that they shared the same payoff function. Prasnikar and Roth found that this change in available information did not significantly affect the outcome of the game; players with full information also chose actions that were very close to the subgame perfect equilibrium prediction.

Prasnikar and Roth concluded that the difference in outcomes between the best-shot and ultimatum games could be attributed to differences in *off-the-equilibrium-path incentives*. They

---

<sup>6</sup>See Roth (1995) or Güth and Tietz (1990) for surveys.

observed that Player 1s in the ultimatum game do better in terms of expected payoff (given the observed distribution of Player 2 responses to particular Player 1 actions) by deviating *further* from the subgame perfect equilibrium. In contrast, in the best-shot game, Player 1s do worse in terms of expected payoff by deviating further from the subgame perfect equilibrium. Hence, convergence to equilibrium is much more likely in the best-shot game than in the ultimatum game, regardless of whether players in the best-shot game have full information about payoffs.

### 3 Observation of Others

With the exception of the study by Harrison and McCabe (1992), the experimental studies that we have discussed all envision that players learn using only their own past history of play. We now consider the possibility that players can observe the actions and payoffs of another pair of players. In particular, we want to consider versions of the ultimatum and best-shot games where, after seeing their own actions and payoffs, each matched pair of players is able to observe the actions and payoffs of one randomly chosen pair of players in the just-completed round of play. We refer to these versions of the ultimatum and best-shot games as versions *with observation*. Games in which players observe only their own history of play will be referred to as versions *without observation*.

The restriction on observation to a single pair of players is intended to minimize the potential for strategic actions aimed at reputation-building. In particular, Player 2s in both games may seek to form “tough” reputations as a means of obtaining more favorable treatment from Player 1s, who move first. In the ultimatum game, Player 2s build reputations by rejecting proposals that give them low payoffs. In the best-shot game, Player 2s build reputations by responding to low Player 1 investment amounts with low investment amounts of their own. Reputation-building in both games is a consequence of the repeated nature of the games and of the information that players receive. While reputation-building is possible in both treatments, it may be more prevalent in sessions with observation since in these sessions, Player 2s know that their actions can be observed by *two* Player 1s rather than only one.

We begin our inquiry into the effect of observation on learning behavior by conducting simulations using a version of Roth and Erev’s (1995) reinforcement learning model in versions of the ultimatum and best-shot games with and without observation. The learning algorithm was modified so that it was capable of processing information on *observed* actions and payoffs as well as on

player's own actions and payoffs. We use the simulation results from this model to formulate some hypotheses that we later test in our experiment. We chose to use the Roth–Erev model because it has been successful in tracking the qualitative features of the behavior of subjects in experimental versions of the ultimatum and best-shot games without observation (Roth and Erev (1995)) as well as for other games (Erev and Roth (1996), Feltovich (1996)), and also because it is the only learning model that has thus far been used for tracking behavior in these two games.

### 3.1 The Learning Model

The learning model can be described as follows. In round  $t$ , player  $i$  has a propensity  $q_{ij}(t) \geq 0$  to play her  $j^{\text{th}}$  pure strategy. The *probability* that she plays the  $j^{\text{th}}$  pure strategy in round  $t \geq 1$ ,  $p_{ij}(t)$ , is determined by these propensities:

$$p_{ij}(t) = \frac{q_{ij}(t)}{\sum_{j=1}^n q_{ij}(t)}.$$

In every round  $t$ , each player is randomly matched with a player of the opposite type and chooses a strategy according to his or her probabilities. When observation of other players is allowed, each pair of players is randomly matched with another pair of players, with all pairings being equally likely. The strategies of a player and her opponent determine their payoffs, and these payoffs are used to update the players' propensities for round  $t + 1$ . When there is observation, players also update their propensities to play the strategies of the players they have observed.

Suppose in round  $t$ , the  $i$ th Player 1 plays strategy  $j$  resulting in payoff  $\omega_i(t)$  and observes that the  $i'$ th Player 1 plays strategy  $j'$  resulting in payoff  $\omega_{i'}(t)$ . Her propensities are updated as follows:

$$q_{ij}(t+1) = q_{ij}(t) + \omega_i(t) \quad \text{and} \quad q_{ij'}(t+1) = q_{ij'}(t) + \beta\omega_{i'}(t),$$

where  $\beta \geq 0$  is the weight placed on observed payoffs. Thus Player 1s treat observed outcomes similarly to own outcomes, but give them  $\beta$  times as much weight. This is meant to capture the idea that firsthand experience usually takes precedence over secondhand experience. There is no updating of propensities to play strategies that were neither chosen nor observed in round  $t$ . Initial ( $t = 0$ ) propensities are assumed to be exogenously given.

In modeling how Player 2s update their propensities, we must make assumptions as to how strategies are inferred from observed actions. When a Player 2 observes the action  $a_2$  of another Player 2, he also observes the action  $a_1$  of the Player 1 matched to that Player 2. Thus, each Player

2 knows the set of strategies that could possibly have been played by the observed Player 2, given the observed action profile  $a = (a_1, a_2)$ . Let  $\{j'_1, j'_2, \dots, j'_n\}$  be the set of possible strategies given  $a$ . We assume that

$$\Pr(j'_m|a) = \frac{p_{ij'_m}(t)}{\sum_{\mu=1}^n p_{ij'_\mu}(t)}.$$

This inference technique is consistent with Bayes' rule combined with the assumption that simulated Player 2s act as though they all have the same mixed strategy.

Now suppose that in round  $t$  the  $i$ th Player 2 plays strategy  $j$  resulting in payoff  $\omega_i(t)$  and observes that the action profile  $a$  results in the payoff  $\omega_{i'}(t)$ . His propensities are updated as follows:

$$q_{ij}(t+1) = q_{ij}(t) + \omega_i(t) \quad \text{and} \quad q_{ij'}(t+1) = q_{ij'}(t) + \beta\omega_{i'}(t) \Pr(j'|a)$$

for all  $j'$  that are possible given  $a$ .

Setting  $\beta = 0$  gives us the learning model for the no-observation case, which is the one originally proposed by Roth and Erev (1995).

### 3.2 Application to the Ultimatum and Best-Shot Games

For our ultimatum game simulations, we use the set  $\{0, \dots, 9\}$  of possible demands as the pure strategy set for Player 1s. Theoretically, the Player 2 strategy set in the ultimatum game is the set of all functions mapping the set of demands into the set of responses  $\{Accept, Reject\}$ , but we follow Roth and Erev (1995) in limiting the Player 2 strategy set to “threshold” strategies of the form: “accept demands of  $j$  or less, and reject demands of more than  $j$ ,” for some integer  $j \in \{0, \dots, 9\}$ . The subgame perfect equilibrium is  $(9, 9)$ . Given an observed outcome of  $(j, Accept)$ , the set of possible Player 2 strategies is  $\{k | k \geq j\}$ , and given an observed outcome of  $(j, Reject)$ , the set of possible Player 2 strategies is  $\{k | k < j\}$ .

For our best-shot game simulations, we use the set of investment levels  $\{0, \dots, 4\}$  as the pure strategy set for Player 1s. In our experiment with human subjects we allow Player 1s to choose investment amounts greater than 4. However, we noted that in both the Prasnikar-Roth and Harrison-Hirshleifer best-shot experiments, players very rarely chose investment amounts greater than 4. Because the Player 2 strategy set is large, even after we reduce the number of Player 1 strategies to just 5, we reduce the number of Player 2 strategies by restricting investment choices to the set  $\{0, \dots, 4\}$  and by further restricting Player 2 strategies so that a pure strategy is



characterized by how it responds to a Player 1 investment of 0 and how it responds to Player 1 investment amounts that are greater than 0. This results in Player 2s having 25 available pure strategies of the form “Respond to 0 with  $k_1$  and to 1, 2, 3, or 4 with  $k_2$ .” Given an observed outcome of  $(0, m)$ , the set of Player 2 strategies is  $\{(m, k) \mid k \in \{0, \dots, 4\}\}$ , and given an observed outcome of  $(j, m)$  for  $j \geq 1$ , the set of Player 2 strategies is  $\{(k, m) \mid k \in \{0, \dots, 4\}\}$ .

### 3.3 Simulation Results

We conducted 400 simulations, 100 each of the ultimatum game with and without observation and 100 each of the best-shot game with and without observation. Each simulation was run for 200 rounds and corresponds to an experimental session involving 6 Player 1s and 6 Player 2s. We estimated initial propensities from observed actions in previous experiments: we used the first-round play reported in Prasnikar and Roth’s (1992) best-shot game with complete information for all of our best-shot simulations and we used the first-round play reported in Roth *et al.*’s (1991) (Pittsburgh) ultimatum game for all of our ultimatum game simulations. We set  $\beta = 0.5$  in all of our simulations involving observation.

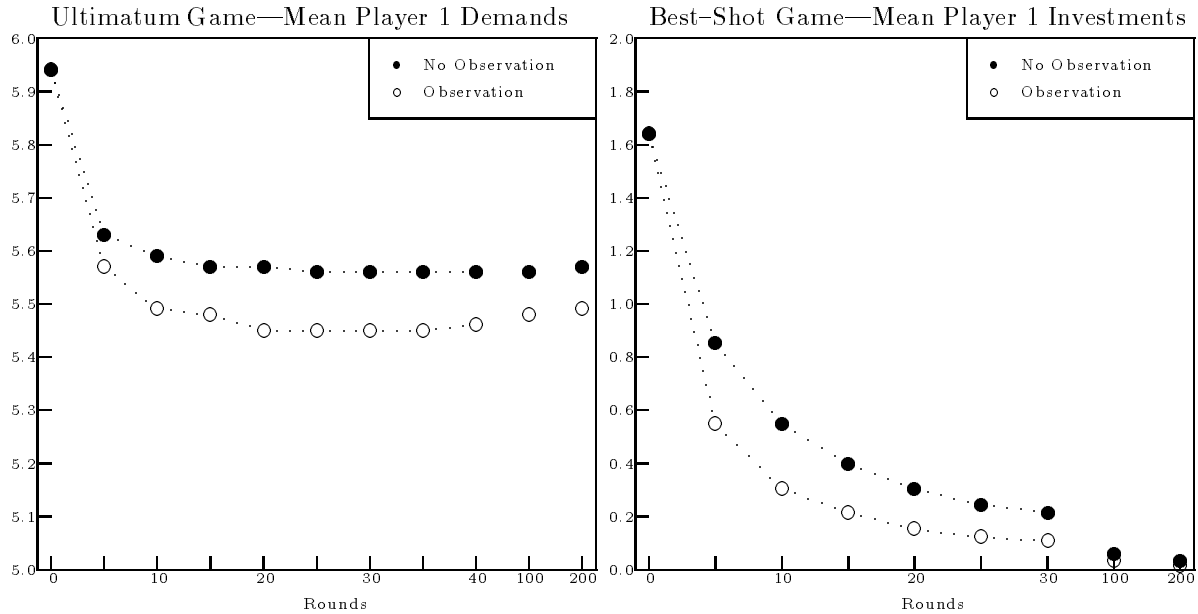
Figure 1 shows mean Player 1 demands in the two ultimatum cells (top panel) and mean Player 1 investment amounts in the two best-shot cells (bottom panel), in both cases averaged over all 100 simulations. We see that observation has opposite effects in the two games. In the ultimatum game, mean Player 1 demands are *further* from the subgame perfect equilibrium with observation than without observation, while in the best-shot game, mean Player 1 investment amounts are *closer* to the subgame perfect equilibrium with observation than without observation. Moreover, the differences in play between the observation and no-observation treatments persist over many rounds, particularly in the ultimatum game.<sup>7</sup>

The difference that we observe between the effect of observation in the ultimatum game simulations and in the best-shot game simulations is due to differences in incentive structures *off the equilibrium path*. Recall that Prasnikar and Roth (1992) found that for “typical” distributions of Player 2 responses in ultimatum game experiments, off-the-equilibrium-path payoffs are such that

---

<sup>7</sup>Our simulation results are robust to changes in the value of our imitation parameter,  $\beta$ , which we set equal to .5. We can think of our simulation results for games *without* observation as corresponding to a choice of  $\beta = 0$ . In some further simulations (not reported above) we considered 10 additional values of  $\beta$ ; 0.1, 0.2, 0.3, 0.4, 0.6, 0.7, 0.8, 0.9, 1.0, and 1.5. We found that in the ultimatum game, as  $\beta$  increases, mean demands move away from the subgame perfect equilibrium more quickly. Similarly, for the best-shot game, as  $\beta$  increases, mean Player 1 investments move closer to the subgame perfect equilibrium more quickly. Thus we conclude that our simulation predictions are robust to changes in the value of  $\beta$ , the imitation parameter.

Figure 1: Simulation Predictions



Player 1s earn higher payoffs by deviating further from subgame perfect play, while in best-shot game experiments, off-the-equilibrium-path payoffs are such that Player 1s earn higher payoffs by playing more in line with subgame perfect play. Since we initialized the learning algorithm using “typical” Player 2 play, the off-the-equilibrium-path incentive structure is maintained in our simulations, and thus deviations from subgame perfect play result in higher payoffs (and more reinforcement) in the ultimatum game than subgame perfect play, while they result in lower payoffs (and less reinforcement) in the best-shot game than subgame perfect play. In ultimatum game simulations with observation, deviations from subgame perfect play are rewarded not only for the player who deviated, but also for the observing player. Thus, the probability and extent of deviation from subgame perfect play become higher even more quickly when we allow for observation of others. In best-shot game simulations with observation, deviations from subgame perfect play are punished not only for the player who deviated, but also for the observing player. Thus, the probability and extent of deviation from subgame perfect play become lower even more quickly when we allow for observation of others.

### 3.4 Hypotheses

Based on our simulations, we hypothesize that allowing observation of the actions and payoffs of another pair of players will have opposite effects on behavior over time in the ultimatum and best-shot games. We focus on the behavior of Player 1s, as their strategies are readily observable from their actions. We hypothesize that: (1) Player 1 demands in the ultimatum game with observation will be further from subgame perfect equilibrium demands (*i.e.*, lower) than Player 1 demands in the ultimatum game without observation; (2) Player 1 investment amounts in the best-shot game with observation will be closer to subgame perfect equilibrium investment amounts (*i.e.*, lower) than Player 1 investment amounts in the best-shot game without observation.

Our first hypothesis runs counter to the claim that deviations from subgame perfect equilibrium behavior are due to players' lack of information. Here we predict that giving players in the ultimatum game more information—a step in the direction of satisfying the mutual knowledge assumption discussed in the introduction—actually increases the degree of deviation from the subgame perfect equilibrium prediction! On the other hand, our second hypothesis reflects the common (game-theoretic) wisdom that giving players more information in the best-shot game leads to play closer to the subgame perfect equilibrium prediction.

## 4 Experimental Design

We used a  $2 \times 2$  experimental design. The two treatment variables were 1) the game played, ultimatum or best-shot, and 2) observation or no observation. We ran three sessions of each of the four experimental cells. Subjects were recruited from the undergraduate population at the University of Pittsburgh. Copies of the instructions used in the experiment are available from the authors upon request.

At the beginning of an experimental session, subjects were randomly assigned the role of Player 1 or Player 2; roles remained unchanged for the duration of the session. In each round of a game, each player was randomly and anonymously paired with a player of the opposite type. After players were paired, Player 1s made and verified their moves. Player 2s observed the moves of their opponents and then made and verified their moves. Then all players were shown their own moves, their opponents' moves, their own payoffs, and their opponents' payoffs. In cells without observation, the next round would begin after all players had pressed a key to continue. In cells

Session	Game	Observe Others?	Number of Subjects	Session	Game	Observe Others?	Number of Subjects
1	Ultimatum	Yes	16	7	Ultimatum	Yes	10
2	Best-Shot	No	8	8	Ultimatum	No	6
3	Best-Shot	Yes	12	9	Best-Shot	No	14
4	Ultimatum	No	12	10	Ultimatum	Yes	14
5	Best-Shot	No	12	11	Ultimatum	No	14
6	Best-Shot	Yes	12	12	Best-Shot	Yes	18

Table 1: Experimental Session Information

with observation, after all players had pressed the key, each pair of players was shown the actions and payoffs of another pair of players, where the observed pairs were assigned as described in section 3.1. After observing these actions and payoffs, players were prompted to press a key to continue. When all players had pressed this key, the next round of play began.

Sessions consisted of 40 rounds of the ultimatum game, or 30 rounds of the best-shot game. (Best-shot games proceeded more slowly.) Subjects were not told how many rounds would be played, but they did know that the session would not exceed 90 minutes. Following the last round, one round was randomly chosen by the computer program, and subjects were paid their earnings in that round in addition to a \$10.00 participation fee. Average total earnings for subjects were roughly \$14.00 in the ultimatum game cells and \$12.00 in the best-shot game cells.

## 5 Experimental Results


### 5.1 Ultimatum Game Results—Player 1 Behavior

Figure 2 shows mean demands of all Player 1s in the ultimatum cells over 5-round intervals. Mean demands in the first 5 rounds are about \$6.00 in both cells.<sup>8</sup> In the no-observation sessions, mean demands increase by about 30 cents from the first 5 to the second 5 rounds and then remain relatively constant while in the observation sessions, mean demands decrease steadily over time to just over \$5.25 in the last 5 rounds. The mean demands illustrated in Figure 2 are also reported in the second and third columns of Table 2. The last two columns of this table report 1-sided robust rank-order test statistics (denoted by  $\hat{U}$ ) and p-values for differences in the distributions of demands between the observation and no-observation sessions for each group of 5 rounds and

---

<sup>8</sup>This mean is similar to mean initial demands found in many ultimatum game experiments; see e.g. Roth (1995).

Figure 2: Ultimatum Game—Mean Player 1 Demands



for all 40 rounds.<sup>9</sup> Looking at the tests, we see that with the exception of the first and fourth 5-round periods, we *can* reject the null hypothesis of no difference between the observation and no-observation sessions in favor of the alternative that Player 1 demands are *lower* in the observation sessions. Thus we find support for our hypothesis regarding the effect of observation on behavior in the ultimatum game.

We further find that we cannot reject the null hypothesis of no change over time in mean demands in the no-observation cell, but we *can* reject this same null hypothesis for mean demands in the observation cell, in favor of the alternative hypothesis that mean demands do change over time (Page test for ordered alternatives,  $p \approx .01$ ).

In addition to examining differences between treatments using pooled *session-level* data, we also made use of the *individual-level* data consisting of the unpooled individual observations on Player 1 demands in all sessions. Using this individual level data from all sessions, we estimated

---

<sup>9</sup>We use the robust rank-order test instead of the more commonly used Mann-Whitney test because the Mann-Whitney test assumes that the samples come from distributions with identical second (dispersion) and higher-order moments, whereas the robust rank-order test makes no such assumption. See Siegel and Castellan (1988). Because the session-level data consist of only six observations for each test, we must use the small-sample distribution of  $\hat{U}$  rather than the normal approximation. We are thus constrained to using only two significance levels—.10 and .05—rather than the near-continuum available for the normal approximation (see Fligner and Policello (1981)). Of course, more session level data would help to alleviate this problem. Our session-level tests can therefore be viewed as rather severe tests for differences between treatments.

Rounds	Mean Demand— Observation Sessions	Mean Demand— No-Observation Sessions	$\hat{U}$	$p$ -value (if $p \leq .10$ )
1–5	5.85	6.09	-0.39	–
6–10	5.85	6.46	$-\infty$	.05
11–15	5.68	6.34	-2.35	.10
16–20	5.66	6.32	-1.13	–
21–25	5.62	6.54	-4.46	.10
26–30	5.60	6.26	-2.35	.10
31–35	5.39	6.40	$-\infty$	.05
36–40	5.27	6.36	$-\infty$	.05
All	5.62	6.35	-2.35	.10

Table 2: Ultimatum Game—Player 1 Demands

an OLS regression model for Player 1 demands.<sup>10</sup> The estimated equation is:

$$\begin{aligned} \text{Player 1 Demand} = & 6.7624 & -0.52354 & (\text{Obs. Dummy}) & +0.00375 & \text{Round} \\ & (0.0801) & (0.1100) & & (0.0030) \\ & -0.0194 & (\text{Round} \times \text{Obs. Dummy}) & + \sum_{i=1}^4 & \text{Session Dummy}_i \\ & (0.0040) & & & \end{aligned}$$

and the adjusted  $R^2$  was 0.2567. The Observation dummy is 1 in sessions with observation and 0 otherwise. Round is the current round number in the session (1–40). The Round  $\times$  Observation Dummy variable is just the product of the round variable and the observation dummy; this variable is intended to capture time-varying effects of allowing for observation of other players. Finally, we included session dummies to control for possible differences across sessions; to save space we have suppressed the coefficient estimates on these session dummies. One can see from the standard errors on the coefficient estimates (given in parentheses) that with the exception of the coefficient on Round, all of the coefficient estimates in the above regression model are significantly different from zero at the .01 level. We conclude that there is support for the notion that observation of others affects the evolution of Player 1 demands.

---

<sup>10</sup>We acknowledge that our use of econometric methods to analyze individual-level data may not be completely valid due to the lack of independence among the individual data; within a particular session, the play of different individuals in a particular round cannot be regarded as consisting of independent observations because the individuals may have interacted with one another before, or may have interacted with the same other players in previous rounds. This problem is further exacerbated in our observation cells. The lack of independence greatly increases the likelihood of Type I error in our hypothesis tests. Thus, our econometric results must be viewed with some caution. We note, however, that in all cases, our regression results using individual-level data serve only to confirm (and do not detract from) the findings of our nonparametric tests based on session-level data which *do* consist of independent observations.

Finally, we examine the effect that observation has on *disaggregated* Player 1 behavior. Recall that in our simulations, we assumed that players updated their propensities for playing observed actions as well as own actions based on the payoffs they received and the payoffs they observed. We now examine whether Player 1s in the experimental sessions used the extra information given them in the observation cell, and if so, how they used it. To assess the effect of observation on individual play, we categorize our sample of Player 1s according to their own outcomes and then examine whether their next-round actions vary with the outcomes they *observed* in the previous round. In particular, we focus on Player 1 actions in rounds following *own* outcomes of (1) an accepted 50–50 split, (2) an accepted 60–40 split, and (3) a rejected 60–40 split. These outcomes account for approximately 87% of all outcomes in all rounds of the ultimatum sessions with observation. For each own outcome, we then examine the extent to which Player 1 actions in the following round are correlated with *observed outcomes*. We consider observed outcomes of (a) an accepted 60–40 split, and (b) a rejected 60–40 split. Table 3a reports, for each possible pairing of these own and observed outcomes, the relative frequency of Player 1s who chose to demand 60% or more of the prize in the next round of play. We see that Player 1s whose demand of 60% was accepted in the current round are substantially more likely to demand 60% or more of the prize in the next round after observing an accepted 60–40 split than after observing a rejected 60–40 split in the current round. That is, Player 1s appear to be taking the observed information into account when choosing actions in later periods. One may wonder whether the apparent correlation between observed outcomes and future actions in Table 3a is due to the additional information players receive, or to some other phenomenon. We sought to determine whether the correlations were spurious by examining similar relative frequencies for the ultimatum game sessions without observation. We randomly assigned all Player 1s in each round of each session *without* observation to another Player 1 in the same round of the same session. We used the same assignment algorithm that we used in our experimental sessions with observation. We then treated the assigned Player 1s as if they were “observed” Player 1s. In Table 3b, we report the mean and standard error of the probabilities (corresponding to those in Table 3a) that we obtained from repeating this exercise 100 times. If the correlations observed in Table 3a are due to some phenomenon unrelated to the presence of observation, we would expect to see similar correlations in Table 3b. Similar correlations are *not* observed; there are no substantial differences between elements in each row of Table 3b. Thus it appears that the correlations seen in Table 3a really are due to Player 1s taking the observed information into account when choosing

Own Outcome	Observed Outcome	
	60 % Accepted	60% Rejected
50 % Accepted	.58	.51
60 % Accepted	.82	.42
60 % Rejected	.44	.32

Table 3a: Use of Observed Outcomes by Player 1s: Probability of Demanding 60% or More in the Next Round Given Own and Observed Outcomes in the Current Round.

Own Outcome	Outcome of Assigned Player 1	
	60 % Accepted	60% Rejected
50 % Accepted	.78 (.05)	.89 (.11)
60 % Accepted	.77 (.02)	.81 (.09)
60 % Rejected	.83 (.07)	.78 (.13)

Table 3b: Player 1s in Sessions Without Observation Randomly Assigned to Other Player 1s: Mean (Std. Error) Probability of Demanding 60% or More in the Next Round Given Own and Assigned Outcomes in the Current Round (100 Simulations).

strategies for the next round.

## 5.2 Ultimatum Game Results—Player 2 Behavior

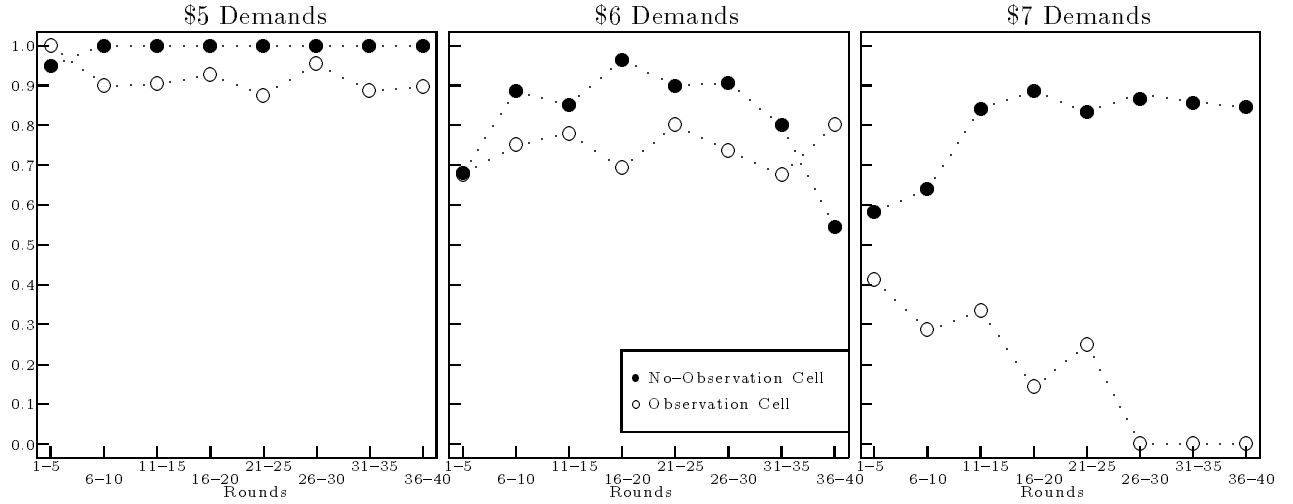
We have suggested that the differences in behavior between Player 1s in the observation and no-observation cells may be explained by the presence of observation in the former. An alternative explanation is that Player 1s are merely reacting to the actions of Player 2s who are behaving differently in the two cells. Many experimenters have sought to explain Player 1 behavior in the ultimatum game by examining how well Player 1 actions correspond to best responses to the historical frequency distribution of Player 2 actions.<sup>11</sup> Figure 3 shows, for both ultimatum game cells, the relative frequencies with which Player 2s accept Player 1 demands of 50%, 60%, and 70% of the prize over 5 round intervals.<sup>12</sup> The principal difference in Player 2 behavior between the two cells appears to be in their responses to proposed 70–30 splits. In the no-observation cell, proposals

<sup>11</sup>See, *e.g.*, Roth *et al.* (1991).

<sup>12</sup>In the bottom panel of Figure 3, frequency values of 0 for a given 5-round period *do not* imply the absence of Player 1 demands of 70%. Rather, they indicate that Player 2s rejected all 70–30 proposals they faced.



Figure 3: Ultimatum Game—Player 2 Frequencies of Acceptance



for 70–30 splits are accepted more than half the time in the first 5 rounds, and the frequency of acceptance increases to over 80% in the last 30 rounds. Nevertheless, using robust rank order tests on frequency data at the session level, we are unable to find statistically significant differences in acceptance rates for Player 2s facing 70–30 proposals between the observation and no–observation treatments. The same finding also holds for Player 2s facing 50–50 and 60–40 proposals, again using robust rank–order tests on frequency data at the session level.

Since the session level data do not reveal any significant differences in Player 2 behavior across treatments, we also examined the individual–level Player 2 data using probit analysis. The dependent variable, *Accept*, is the Player 2’s decision of whether to accept (=1) or reject (=0) a particular Player 1 proposal. We considered Player 2 responses to *all* Player 1 demands as well as to Player 1 demands of 60% and 70%. The model we estimated is:

$$\begin{aligned} \text{Accept} = & \beta_0 + \beta_1 \text{Round} + \beta_2 \text{Observation Dummy} + \beta_3 \text{Round} \times \text{Observation Dummy} \\ & + \sum_{i=1}^4 \beta_{3+i} \text{Session Dummy}_i + \beta_8 \text{Player 1 Demand}. \end{aligned}$$

Here, the Round variable, the Observation dummy, the Round  $\times$  Observation variable and the Session dummies are the same as in the Player 1 regressions reported in section 5.1. The Player 1 Demand variable is the proposal that the Player 2 faced. In those cases where we focus on a particular type of Player 1 proposal, this variable is omitted as a regressor. The regression results

are presented in Table 4.<sup>13</sup>

Regressor	All Player 1 Demands <sup>a</sup>	Player 1 Demands \$6	Player 1 Demands \$7
Constant	8.579*** (0.535)	2.031*** (0.429)	1.106*** (0.282)
Round	-0.007 (0.006)	-0.002 (0.008)	0.016 (0.011)
Observation Dummy	-0.391* (0.228)	-0.012 (0.492)	-0.889* (0.481)
Round $\times$ Obs. Dummy	-0.017* (0.008)	-0.019* (0.011)	-0.075* (0.031)
Player 1 Demand	-1.100*** (0.069)	—	—
Pr $> \chi^2$	0.000	0.000	0.000
Number of Observations	1,440	636	236

Table 4: Probit Estimates: Probability that Player 2 Accepts the Demand of Player 1.

<sup>a</sup>Standard errors in parentheses. Coefficient estimates with a \* (\*\*, \*\*\*) are significantly different from zero at the .10 (.01,.001) level of significance.

We see that in all cases, the coefficient estimate on the Round  $\times$  Observation Dummy term is negative and significant at the .10 level suggesting that over time, the probability that Player 2s accept Player 1 demands is decreasing in the observation treatments as compared with the no observation treatments. We also see that the coefficient estimate on the observation dummy by itself is also negative and significant at the .10 level for the regressions involving all Player 1 demands and Player 1 demands of 70%. Thus we find some evidence that the probability a Player 2 accepts a given Player 1 demand is *lower* in certain cases in the observation treatment compared with the no-observation treatment.

This lower probability of acceptance in the observation treatment may have affected Player 1 demands. Given the frequencies with which Player 1 proposals are accepted by Player 2s the *optimal* Player 1 demand—the demand that maximizes expected monetary payoff—is typically *lower* in the observation cells than in the no-observation cells.<sup>14</sup> Since there is essentially no difference

<sup>13</sup>To save space we have again suppressed the coefficient estimates on the dummy terms.

<sup>14</sup>In two of the three ultimatum sessions with observation, the optimal demand is 5 in all 5-round periods; in the

between Player 2 *initial* behavior (in the first 5 rounds) in the observation and no-observation cells, we conclude that changes over time in Player 2 behavior differ between the observation and no-observation cells, and that this difference may help explain the differences in Player 1 learning behavior.<sup>15</sup>

### 5.3 Best-Shot Game—Player 1 Results

Figure 4 shows mean Player 1 investment amounts in the best-shot game cells over 5-round periods. We see that these investment amounts follow similar paths, both starting above 3, dropping off quickly over the first 10 rounds, and then remaining approximately constant over the last 15 rounds, near the subgame perfect equilibrium prediction of 0.<sup>16</sup> In fact, the modal Player 1 investment amount is 0 in all 5-round periods of the observation cell, and in all 5-round periods except the first in the no-observation cell. We also observe that in every 5-round period, Player 1 investment amounts are somewhat *closer* to the subgame perfect equilibrium prediction of 0 in the observation cell than in the no-observation cell, consistent with our hypothesis regarding behavior in the best-shot game.

However, according to robust rank-order tests reported in Table 5, the observed differences in investment amounts between the two treatments are not significant ( $p > .10$ ), except in the first 5 rounds. The observed decrease over time in mean Player 1 investment amounts in the no-observation cell is significant (Page test for ordered alternatives,  $p \approx .005$ ); however, it is only marginally significant in the observation cell (Page test for ordered alternatives,  $p \approx .1$ ). We conclude that Player 1s in *both* cells quickly learn to play close to the subgame perfect equilibrium and that there is little support for our hypothesis that observation leads Player 1s in the best-shot


---

other session, the optimal demand is 6 in all 5-round periods. In two of the three sessions without observation, the optimal demand varies between 5 and 6; in the other session, the optimal demand is 6 in the first two 5-round periods and 7 in the last six 5-round periods.

<sup>15</sup>It is also of interest to consider whether Player 2's in our ultimatum game experiments are foregoing significant payoff amounts in their rejection of positive offers by Player 1s. If Player 2s perceive that the monetary stakes are low, they may (rationally) refuse to employ much effort in assessing how to play the game. (For a further discussion of this subject, see e.g. Harstad and Marrese (1982) for public good games, or Hoffman, McCabe and Smith (1996) for ultimatum games.) In our ultimatum games *without* observation, the total amount foregone each round by Player 2s averaged 15.7% of expected earnings per round (given Player 1 proposals) while in sessions *with* observation the total amount foregone each round averaged 19% of expected earnings per round (given Player 1 proposals).

<sup>16</sup>These initial investment levels are almost twice as high as those found by Harrison and Hirschleifer (1989) and by Prasnikar and Roth (1992). This difference may be due to the repeated nature of our game versus the one-shot nature of the earlier experiments. It is also worth noting that Harrison and Hirschleifer conducted only one sequential best-shot session and Prasnikar and Roth conducted only two; therefore, it may not yet be clear what initial conditions one should expect to find in a best-shot experiment.

Figure 4: Best-Shot Game—Mean Player 1 Investments



game to adopt strategies closer to the subgame perfect prediction.

An OLS regression serves to confirm that observation of others does not significantly affect Player 1 behavior. Our estimated model of Player 1 investment is:

$$\begin{aligned}
 \text{Player 1 Investment} = & 4.2568 & -0.6662 (\text{Obs. Dummy}) & -0.3456 (\text{Round}) \\
 & (0.3101) & (0.4163) & (0.0432) \\
 & -0.0042 (\text{Round} \times \text{Obs. Dummy}) & +0.0082 \text{Round}^2 \\
 & (0.0581) & (0.0014) \\
 & 0.0008 (\text{Round}^2 \times \text{Obs. Dummy}) & + \sum_{i=1}^4 \text{SessionDummy}_i \\
 & (0.0040) & & 
 \end{aligned}$$

Rounds	Mean Investment— Observation Sessions	Mean Investment— No-Observation Sessions	$\hat{U}$	$p$ -value (if $p \leq .10$ )
1-5	3.14	3.94	$-\infty$	.05
6-10	1.43	1.81	-0.18	—
11-15	0.77	1.47	-1.13	—
16-20	0.92	1.19	-1.13	—
21-25	1.21	1.31	-0.18	—
26-30	0.88	1.14	-0.53	—
All	1.39	1.81	-1.13	—

Table 5: Best-Shot Game—Player 1 Investment Amounts

The adjusted  $R^2$  is 0.1893. Here the Round and Observation variables are as in the ultimatum game regressions. We have also included variables based on a quadratic trend term, Round<sup>2</sup> to account for the possibility of nonlinear changes in Player 1 investment levels over time. We see from the standard errors on the coefficient estimates (given in parentheses), that the coefficients on variables involving the observation dummy are never significantly different from zero. On the other hand, the Round and Round<sup>2</sup> terms are significantly different from zero; the negative coefficient on Round and the positive coefficient on Round<sup>2</sup> indicate that Player 1 investment levels decrease over time at a decreasing rate. We conclude that there is no support for our hypothesis that observation affects the evolution of Player 1 investments. Player 1s in both treatments learn at about the same rate to play near the subgame perfect strategy.

Furthermore, Player 1s in the best-shot game do not appear to be influenced by differences in payoffs between own and observed actions. In contrast to the ultimatum game, we find no strong correlation between best-shot game payoffs (own or observed) and next-round actions using an analysis similar to the one we conducted for the ultimatum game (as reported in Table 3a–3b).<sup>17</sup> Since we assumed in our simulations that Player 1s are influenced by such differences in payoffs, it is not surprising that the predictions of our best-shot simulations were not borne out in the experimental data.

#### 5.4 Best-Shot Game – Player 2 Behavior

Figure 5 illustrates the evolution of Player 2 behavior in the best-shot game over 5-round periods. The top panel shows Player 2s' mean response to *nonzero* investments by Player 1s. While Player 2s do not initially recognize that 0 is a best response to any nonzero investment by Player 1s, they begin to recognize this strategy after the first few rounds, as the mean response to nonzero investment drops sharply from rounds 1–5 to rounds 6–10, and stays low throughout the last 20 rounds.<sup>18</sup> In the bottom panel, we see that Player 2s initially respond to a Player 1 investment of *zero* with 3 on average in both cells. Then, instead of increasing toward the monetary best response of 4, Player 2 responses actually decline slightly in both cells.


Robust rank-order tests show no significant differences between cells in Player 2 responses to either zero or nonzero investment levels by Player 1s using session level data. Therefore, we again

---

<sup>17</sup>We have therefore chosen not to report this analysis here. See Duffy and Feltovich (1996) for details.

<sup>18</sup>The apparent increase in the last few rounds of the sessions with observation is the result of a few outliers; the median and modal responses are 0 in every 5-round period but the first in both treatments.

Figure 5: Best-Shot Game—Mean Player 2 Responses  
 (Player 1 Nonzero Investment) (Player 1 Zero Investment)



turned to an examination of the individual level Player 2 data using probit analysis. The dependent variable, Player 2 Investment, is the Player 2's decision of whether to invest a positive amount ( $=1$ ) or zero ( $=0$ ) in response to a particular Player 1 investment level. As in the aggregate data analysis, we considered responses to all Player 1 investment levels and to Player 1 investment levels of zero and nonzero amounts. The model we estimated is given as:

$$\begin{aligned}
 \text{Player 2 Investment} = & \beta_0 + \beta_1 \text{Round} + \beta_2 + \beta_3 \text{Round}^2 + \beta_4 \text{Observation Dummy} \\
 & + \beta_5 \text{Round} \times \text{Observation Dummy} + \beta_6 \text{Round}^2 \times \text{Observation Dummy} \\
 & + \sum_{i=1}^4 \beta_{6+i} \text{Session Dummy}_i + \beta_{11} \text{Player 1 Investment}.
 \end{aligned}$$

The round and observation variables are similar to those used in the ultimatum game probit analysis. Note that in these regressions, we also include variables involving the square of the round number,  $\text{Round}^2$ ; the purpose is to capture the possibility of very rapid, nonlinear changes in behavior over time. The Player 1 investment variable is the amount invested by the Player 1 that is matched with the Player 2. The regression results are presented in Table 5.<sup>19</sup> We see from this table that in all three cases, the variables involving the observation dummies are never significant; this finding is consistent with our tests for differences in Player 2 behavior using the session-level

<sup>19</sup>To save space we have again suppressed the coefficient estimates on the dummy terms.

Regressor	All Player 1 Investments <sup>a</sup>	Player 1 Investment = 0	Player 1 Investment > 0.
Constant	1.657*** (0.224)	2.703*** (0.551)	0.504* (0.298)
Observation	0.053 (0.273)	-0.007 (0.701)	-0.490 (0.339)
Round	-0.134*** (0.029)	-0.193** (0.621)	-0.190*** (0.044)
Round <sup>2</sup>	0.004*** (0.001)	0.005** (0.002)	0.005*** (0.002)
Round × Obs. Dummy	0.029 (0.038)	0.057 (0.080)	0.048 (0.059)
Round <sup>2</sup> × Obs. Dummy	-0.001 (0.001)	-0.002 (0.002)	-0.001 (0.002)
Player 1 Investment	-0.175*** (0.021)	–	0.067* (0.033)
Pr > $\chi^2$	0.000	0.000	0.000
Number of Observations	1,140	663	477

Table 5: Probit Estimates: Probability that Player 2 Invests a Positive Amount, Given Player 1 Investment

<sup>a</sup>Standard errors in parentheses. Coefficient estimates with a \* (\*\*, \*\*\*) are significantly different from zero at the .10 (.01,.001) level of significance.

data. Like the Player 1 regressions, we also see that the only significant coefficient estimates are associated with the round and round<sup>2</sup> terms; the negative sign on the round coefficient and the positive sign on the round<sup>2</sup> coefficient imply that Player 2’s probability of investing a positive amount decreases over time at a decreasing rate. The deviation from subgame perfect play (in response to zero Player 1 investments) might be explained by reputation–building on the part of Player 2s.<sup>20</sup>

We note that our Player 2 responses to zero investment amounts differ from the game–theoretic best response of 4 and also from the mean responses observed by Harrison and Hirshleifer (1989) and Prasnikar and Roth (1992) that were close to 4. These differences might be explained by the repeated–game nature of our experiment and the resulting incentives for reputation–building that are not present in the experiments of Harrison–Hirshleifer and Prasnikar–Roth. However, while

---

<sup>20</sup>See section 5.5.

Player 2s do not always respond optimally (in terms of monetary payoff) to Player 1 actions, they do play close enough to their monetary best responses that the optimal Player 1 investment in every 5-round period is zero. Given this Player 2 behavior, a Player 1's payoff from investing zero is always much greater than her payoff from investing a positive amount. Thus, even in the early rounds, Player 1s face a strong incentive to play the subgame perfect strategy.<sup>21</sup>

## 5.5 Reputation-Building

We have previously mentioned that the effect of increased information on Player 1 behavior may be complicated by the fact that Player 2s know that this information is being provided, and may recognize and act upon the increased incentive to establish a reputation of "toughness" in treatments with observation relative to those without observation. In Duffy and Feltovich (1996), we look for evidence of such reputation building by examining certain types of behavior by Player 2s in observation and no observation cells. While we find evidence that suggests Player 2s seek to build tough reputations in *both* cells, the evidence also suggests that there is not much of a difference in the incidence of reputation-building by Player 2s in the observation cells as compared with the no-observation cells. This finding is not too surprising, since the incentives for reputation building in the observation cells are not all that much greater than in the no-observation cells; players in the observation cells only get to observe the actions and payoffs of *one* pair of players.

## 6 Conclusion

In this paper we have examined the effect of additional information on learning in the ultimatum and best-shot games. In one treatment we allowed pairs of players to observe, prior to choosing their own actions, the actions and payoffs of one other pair of players in the just completed round of play. We have compared and contrasted the results from this treatment with results from the standard treatment in which observation of other players is not allowed, and found that there were some differences in the results.

Our main finding is that increasing the *quantity* of information available to players need not hasten the rate at which players learn to play subgame perfect equilibrium strategies. In the ultimatum game, giving players additional information leads to play that is significantly further

---

<sup>21</sup>We note again (*cf.* footnote 15) the amounts foregone on average each round by Player 2s were 21.7% of expected earnings per round in sessions without observation and 24% of expected earnings per round in sessions with observation (given Player 1 investment levels).



from the subgame perfect equilibrium prediction relative to the case where players do not receive this additional information. In contrast, in the best-shot game, giving players additional information does not significantly affect the rate at which players learn to play subgame perfect equilibrium strategies. The reason for these different results lies in differences in the *content* of the additional information that players receive. In the case of the ultimatum game with observation, Player 1s see that other Player 1s get higher payoffs by deviating further from the subgame perfect strategy, and hence observation only reinforces this type of behavior. In the case of the best-shot game, the off-the-equilibrium path incentives are such that players do better by moving closer to subgame perfect equilibrium strategies. The incentive to play the equilibrium strategy in both treatments is so strong that players quickly react to this incentive on their own, so that additional information on how other players play the game is not as useful as it is in the ultimatum game, and does not significantly affect the manner in which players play the best-shot game. We conclude that the effect of additional information on the evolution of play cannot be ascertained *a priori*; one must consider the *content* of the additional information as well.

## References

- Aumann, R. and A. Brandenburger (1995), "Epistemic Conditions for Nash Equilibrium," *Econometrica* 63, 1161–1180.
- Bandura, A. and R.H. Walters (1963), *Social Learning and Personality Development*, New York: Holt, Rinehart and Winston.
- Banerjee, A.V. (1992), "A Simple Model of Herd Behavior," *Quarterly Journal of Economics* 107, 797–817.
- Banerjee, A.V. and D. Fudenberg (1995), "A Simple Model of Bayesian Word-of-Mouth Learning," mimeo.
- Bikhchandani, S. D. Hirshleifer, and I. Welch (1992), "A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades," *Journal of Political Economy* 100, 992–1026
- Bolton, G.E. (1991), "A Comparative Model of Bargaining: Theory and Evidence" *American Economic Review* 81, 1096–1136.
- Bush, R. and F. Mosteller (1955), *Stochastic Models for Learning*, New York: Wiley.
- Conlisk, J. (1980), "Costly Optimizers Versus Cheap Imitators," *Journal of Economic Behavior and Organization* 1, 275–93.
- Duffy, J. and N. Feltovich (1996), "The Effect of Information on Learning in Strategic Environments: An Experimental Study," working paper, University of Pittsburgh.
- Ellison, G. and D. Fudenberg (1993), "Rules of Thumb for Social Learning," *Journal of Political Economy* 101, 612–643.
- Ellison, G. and D. Fudenberg (1995), "Word-of-Mouth Communication and Social Learning," *Quarterly Journal of Economics* 110, 93–125.
- Erev, I. and A.E. Roth (1996), "On the Need for Low Game Theory: Reinforcement Learning in Experimental Games With Unique, Mixed-Strategy Equilibria." working paper.
- Feltovich, N. (1996), "Learning and Equilibrium in a Simple Asymmetric-Information Game: An Experimental Study," mimeo, University of Pittsburgh.
- Fligner, M.A. and G.E. Policello (1981), "Robust Rank Procedures for the Behrens-Fisher Problem," *Journal of the American Statistical Association* 76, 162–68.
- Gale, J., K.G. Binmore, and L. Samuelson (1995), "Learning to be Imperfect: The Ultimatum Game," *Games and Economic Behavior* 8, 56–90.
- Güth, W., R. Schmittberger, and B. Schwarz (1982), "An Experimental Analysis of Ultimatum Bargaining," *Journal of Economic Behavior and Organization* 3, 367–388.
- Güth, W. and R. Tietz (1990), "Ultimatum bargaining behavior: A Survey and Comparison of Experimental Results," *Journal of Economic Psychology* 11, 417–449.
- Harrison, G.W. and J. Hirshleifer (1989), "An Experimental Evaluation of Weakest Link/Best Shot Models of Public Goods," *Journal of Political Economy* 97, 201–225.
- Harrison, G.W. and K.A. McCabe (1992), "Expectations and Fairness in a Simple Bargaining Experiment," working paper B-92-10, University of South Carolina.

- Harstad, R.M. and M. Marrese (1982), "Behavioral Explanations of Efficient Public Good Allocations," *Journal of Public Economics* 19, 367–383.
- Hoffman, E., K.A. McCabe, and V.L. Smith (1996), "On Expectations and the Monetary Stakes in Ultimatum Games," forthcoming, *International Journal of Game Theory*.
- Jackson, M. and E. Kalai (1995), "Social Learning in Recurring Games," working paper, Northwestern University.
- Prasnikar, V. and A.E. Roth (1992), "Considerations of Fairness and Strategy: Experimental Data from Sequential Games," *Quarterly Journal of Economics* 107, 865–888.
- Reichard, G.A. (1938), "Social Life," in F. Boas, ed. *General Anthropology* Boston: Heath.
- Roth, A.E. (1995), "Bargaining Experiments," in: J. Kagel and A.E. Roth, eds. *Handbook of Experimental Economics* Princeton: Princeton University Press, 253–348.
- Roth, A.E. and I. Erev (1995), "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior* 8, 164–212.
- Roth, A.E., V. Prasnikar, M. Okuno-Fujiwara and S. Zamir (1991), "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study," *American Economic Review* 81, 1068–1095.
- Siegel S. and N.J. Castellan, Jr. (1988), *Nonparametric Statistics for the Behavioral Sciences* 2d ed., New York: McGraw Hill.
- Vives, X. "Social Learning and Rational Expectations," *European Economic Review* 40, 589–601.