

Testing Subgame Perfection Apart From Fairness in Ultimatum Games*

James Andreoni
Department of Economics
University of Wisconsin
Madison, WI 53706
email: andreoni@wisc.edu

Emily Blanchard
Department of Economics
University of Virginia
Charlottesville, VA 22904

June 2002

Current version July 26, 2005

forthcoming in *Experimental Economics*

Abstract

We present an experiment designed to separate the two commonplace explanations for behavior in ultimatum games—subjects’ concern for fairness versus the failure of subgame perfection as an equilibrium refinement. We employ a tournament structure of the bargaining interaction to eliminate the potential for fairness to influence behavior. Comparing the results of the tournament game with two control treatments affords us a clean test of subgame perfection as well as a measure fairness-induced play. We find after 10 iterations of play that about half of all non-subgame-perfect demands are due to fairness, and the rest to imperfect learning. However, as suggested by models of learning, we also confirm that the ultimatum game presents an especially difficult environment for learning subgame perfection.

*We are grateful to Larry Samuelson, Gary Bolton, Ragan Petrie, Richard Zeckhauser, Klaus Abbink, and two anonymous referees helpful comments. Andreoni thanks the National Science Foundation and Russell Sage Foundation for financial support.

1. Introduction

The ultimatum game (Guth et. al., 1982) has come to symbolize the power of subgame perfection in game theory and its utter failure in practice. In this bargaining game, a ‘proposer’ makes a take-it-or-leave-it offer to the ‘responder,’ who subsequently accepts or rejects the offer. The subgame perfect equilibrium prediction if players care only about their own monetary payoffs is that the proposer offers nothing (or almost nothing) to the responder, who accepts any positive offer. The myriad tests of the ultimatum game find proposers’ offers well in excess of the subgame perfect equilibrium prediction. This paper examines the two most often suggested explanations for this discrepancy between experimental outcomes and the subgame perfect equilibrium prediction: fairness and learning.

Given that some subjects’ preferences may extend beyond their own monetary rewards to include a notion of fairness—a view well accepted by economists—observed experimental outcomes are not inconsistent with the subgame perfect equilibrium prediction. Whether fairness can account for *all* of the deviation from predicted equilibrium play, however, has never been tested explicitly in the ultimatum game setting. Indeed, whether and to what extent imperfect learning also explains experimental outcomes is of considerable interest. Experiential learning models put forward by Camerer and Ho (1999), Erev and Roth (1995), Fudenberg and Levine (1997), and Gale, Binmore, and Samuelson (1995) can explain the deviation from subgame perfect play in the absence of fairness concerns. An intriguing question is then just how much of what looks like fairness is actually behavior by people who have simply not yet adopted the logic of subgame perfection.

This paper reports a test for the subgame perfect prediction using an ultimatum game experiment designed to subtract away fairness concerns. Using a tournament structure of payoffs, we remove the incentives for fairness and evaluate whether subjects’ behavior under this treatment approaches the predicted subgame perfect play. This structure allows us not only to test for the existence of fairness concerns, but also to quantify the proportion of the deviation from the subgame perfect prediction attributable to fairness versus imperfect learning over a game of “standard” length.

The tournament version of the ultimatum game mitigates fairness concerns by paying subjects by their rank relative to the other players of their own type (either proposers or responders), rather than their absolute earnings.¹ Participants are

¹This method has also been employed by Andreoni (1995). We were delighted to learn

ranked based on their total earnings from playing *all* players of the other type in each round. The experiment provides evidence that a great deal of the deviation from the subgame perfect prediction can be attributed to fairness; the tournament induces significantly lower offers by proposers and demands by responders. Yet we find that subjects do immediately adopt subgame perfect strategies, suggesting that the ultimatum game environment is a particularly difficult one for learning. This leads us to believe that future explorations of learning in ultimatum games should focus on tests of much longer term experiments, using dozens or perhaps hundreds of games. We also find that preferences for fairness are robust; even in the face of learning when subjects come to understand the dominance of sub-game perfection, they continue to reject unfair offers.

Several papers are particularly relevant to this work. First, an important paper by Bolton (1991) hypothesizes that individuals have “comparative preferences” and tests this by applying a tournament structure to the multi-period bargaining game of Ochs and Roth (1989). Bolton finds mixed evidence of increased subgame perfect play under the tournament structure, but he does show that the tournament structure substantially reduces the frequency of disadvantageous counter proposals observed by Ochs and Roth. A second paper by Abbink, Bolton, Sadrieh, and Tang (2001) tests for adaptive learning and fairness in a modified version of the ultimatum game. Focusing on the interaction between fairness and learning, the authors suggest that subjects’ utilities from punishment or retribution keep the rejections above the sub-game perfect levels. Finally, a paper by Goeree and Holt (2000) presents an experiment designed to test for fairness in deviations from players’ behavior from the subgame perfect prediction in a two-stage bargaining game with alternating offers. By introducing asymmetric fixed payments to the players, Goeree and Holt accentuate the earnings inequality that arises in the outcome predicted by subgame perfection. The monetary payments are cleverly selected to induce a perfectly *negative* correlation between the “fair” initial proposer offer (the offer that equalizes net payments between the two players) and remaining pie size, whereas the strategically optimal initial proposer offer is perfectly *positively* correlated with remaining pie size. The experimental data reveals a clear negative relationship between initial proposer offers and remaining pie size, indicating a significant role for fairness in deviation from the subgame perfect prediction.

recently that the first application of a tournament condition in an economics experiment was in the senior honors thesis of Richard Zeckhauser, “Collaboration and Composite Games,” chapter 3, presented for graduation from Harvard College in 1962.

The remainder of the paper is structured as follows. The next section reviews our experimental methodology and outlines in detail the structure of each treatment of the ultimatum game. Section 3 presents the results of the experimental sessions, examining in turn the behavior of responders and proposers. The proportion of “imperfect” behavior attributable to learning versus fairness is estimated in section 4. Section 5 concludes.

2. Experimental Design

We employ a tournament-style variation of the basic ultimatum game to examine the relative importance of fairness and learning in explaining deviations from predicted subgame perfect equilibrium behavior. Two additional treatments are included as controls, which we call the ‘standard’ and ‘round robin’ games.

A game of any condition involves 20 subjects, 10 as proposers and 10 responders. The three conditions employ the same basic bargaining interaction. In each interaction a pair of subjects bargains over the division of 100 chips. Our design asks subjects to record their strategies simultaneously; the proposers are asked to make a ‘take it or leave it’ offer while responders are simultaneously asked to make a “demand”, marking the offers below which she will reject the proposer’s offer and thereby leave both subjects empty handed. We interpret outcomes with the agreed division of 100-0 (proposer-responder) or 99-1 as corresponding to a subgame perfect equilibrium when players care only about their monetary rewards. In each game, subjects participate in a series of 10 interactions. We chose the 10-round repetition to make our game consistent with others in the literature, such as Slonim and Roth (1998), and Roth, Prasnikar, Okunofujiwara, and Zamir (1991).

In the standard game, proposers and responders are paired randomly each round. In a given round, each proposer makes an offer to his paired responder, who simultaneously makes a demand. If the offer is greater than the demand, then it is accepted and both players earn the payoffs specified by the offer. If the demand exceeds the offer, the offer is rejected and both players earn nothing. Players are then randomly matched for the subsequent round, making sure never to play the same person twice. Subjects are paid the sum of their earnings over the 10 rounds. Notice two important features about this simple version of the game. First, a subject’s payoff in each round depends only on his own and his partner’s decisions and therefore his earnings from playing any given strategy may vary depending on the person with whom he is paired. Second, in this form of

the game, subjects' preferences for fairness will influence behavior.

The round robin treatment serves as an intermediate control. Here each subject plays *all* of the players of the other role in every round. A proposer's offer is matched against each of the 10 responders' demands for that round, and the proposer is paid the average earnings from all 10 interactions. Similarly, each responder's demand is matched against every one of the 10 proposers' offers, earning him the average payoff from those 10 interactions. Players learn their own average earnings, but do not see the choices of their opponents.² The round robin control is included not because we are interested in the potential effect of modifying the standard game, but because we want to isolate the effect of removing fairness from the bargaining game. Thus, comparing the round robin and tournament results will provide precisely the desired framework for analysis.

The tournament ultimatum game proceeds just as the round robin version, with the exception that subjects are paid based on their earnings *rank* rather than their absolute earnings. Subjects are ranked among their *own* type based on score, where scores are equivalent to the earnings determined in the round robin game. For instance, a proposer's score is the average payoff from matching her offer to all 10 responders' demands, just as in the round robin game. The 10 proposers are then ranked by score and paid accordingly, with the highest score earning the greatest monetary payoff. Similarly, responders are ranked according to their average payoffs from playing all 10 proposers in each round, and are paid according to their ranks among the set of responders. Hence, subjects' earnings are ultimately determined in a constant-sum game among players of their own type.

The important aspect of the tournament structure is that it mitigates fairness vis-a-vis players of the other type. Intuitively, once a responder realizes that the way to make the most money is to accept all offers, then "selfish" offers by proposers will not make the responder worse off or cost him any money, thus the responder has no reason to take exception to them. Stated in terms of some recent

²Averaging the returns from playing multiple players has several implications for the game. The process of averaging means that some of the riskiness of a given strategy may be eliminated in this treatment relative to the standard game. It might also be that playing a bit-part against a large group may change somewhat the fairness incentives in the round robin treatment relative to the standard game. In addition, this format may add to any confusion already present. Finally, playing every player in each round introduces the possibility of a kind of diffused reputation building, whereby a player may seek to influence play over the course of the game. Despite these possibilities, the between-role fairness concerns remain, since payoffs are still determined by the division of a monetary pie.

theories of fairness, the tournament structure aligns concerns for both absolute and relative payoffs.³ The set of relative payoffs is fixed, so the only objective left is to secure a higher rank for one's self. In fact, none of the models of fairness discussed in the literature would predict a role for fairness within roles or between proposers and responders in the tournament game.⁴ It is conceivable that social preferences nonetheless could continue to play a role in some responders' behavior, since any given responder has the ability to improve the rank of more "generous" proposers by rejecting low offers. Yet it is difficult to envision exactly how this sort of social preference might manifest in players' behavior— not only because the consequences of rejecting a low offer are no different from rejecting a high offer (any rejection weakly reduces the responder's own rank but does not reduce the size of the payoff shared among proposers), but also because in the tournament setting there is no sense in which a high offer indicates any greater generosity on the part of proposers to their responder counterparts.

Since the tournament condition removes a role for fairness, any deviation from the subgame perfect prediction may be attributed to imperfect learning considerations. Taking into consideration the potential for learning over time, the theoretical prediction is that the outcomes in the tournament game should tend closer to the subgame perfect prediction than in the other two treatments. To the extent that the tournament ranking-scheme introduces additional confusion to the game, however, the learning process may be somewhat slower in the tournament game. As such, any observed trend towards subgame perfect behavior in the tournament game may *understate* the importance of fairness concerns in the ultimatum game setting.

We ran two sessions of the experiment. Each session required 60 subjects, hence we report data on 120 subjects. Subjects were recruited from the University of Wisconsin-Milwaukee. They were invited to log on to an Internet web site. Upon registering for the game they were assigned at random to one of the three conditions. Participants received instructions for the game, and each made his first decision. They were then instructed to log on again each day for nine more days and to make another decision each day. If decisions had not been recorded

³See Fehr and Schmidt (1999), Bolton and Ockenfels (2000), Rabin (1993), Rabin and Char-ness (2002), and Li (2004).

⁴Is it theoretically possible that tournament players with a strong aversion to inequality could attempt to enforce equality of payoffs by colluding to tie each round, and thus each earn the average. Yet the only cheat-proof opportunity for "fair" collusion is the subgame perfect prediction in which all responders accept all offers and all proposers offer nothing. Hence, the one possible notion of fairness we can imagine will also be expressed as subgame perfection.

by 4:00 P.M. a reminder e-mail was sent. Each subject could amend his decision throughout the day, and could view a “history” of his play and payoffs in previous rounds at any time. The decision on record at 2:00 A.M. was carried out and the results were reported to the subjects via e-mail. Participation was completely anonymous, except to the extent that participants were paid by check at the conclusion of the experimental session.

The format of one-decision-per-day was chosen to allow greater opportunity for learning in an effort to give the subgame perfect prediction its best chance. By giving people time to think about the game over the course of a day, rather than over only a minute or two as in the typical laboratory experiment, we hoped to allow subjects more time for reasoning and reflection, and thus give us greater confidence in the event that we find evidence against the subgame perfect prediction. The move to a decision-a-day format is not without trade-offs, however. First, by switching from the conventional single session laboratory setting, we lose direct comparability to previous ultimatum game experiments. Additionally, the multiple-day testing introduces the possibility of communication among subjects between rounds. In recognition of this potential problem, we took additional precautions in subject recruitment. By recruiting subjects from the UW-Milwaukee campus (rather than Madison, where the study was based) we reduced the probability that subjects knew one another, since the majority of UW-Milwaukee students are commuters and part-time enrollees. Second, as subjects registered for the study they were randomly assigned to one of the three games. Even if two friends enrolled in the experiment together, there was only a one-in-three chance that they would be placed in the same treatment. Finally, we confirmed that no subjects in any session shared the same address.

All interactions were to divide 100 tokens. In the standard and round robin games, tokens were worth \$0.05 each, for a \$5 daily pie, and total stakes of \$50 over the ten rounds. Rank-based payoffs in the tournament treatment were commensurate. Payoffs ranged from \$5 for the top rank down to \$0 for the last rank.⁵ In the event that a subject failed to log in for a day (and thus a round), his prior decision was implemented again and he was charged a penalty of \$2, which was deducted from his total earnings at the close of the game. The threat of this monetary penalty may explain the high participation rate throughout the games.⁶

⁵The 2nd rank earned \$4.25 and the 3rd rank earned \$3.50. After that, payments fell by \$0.50 for each rank, until the 10th rank earned \$0.

⁶In the first session, 12.5 percent of all decisions were the result of a failure to log on. In session 2 the penalty for failing to log on was increased to \$5 per day, and 8 percent of all

We also conducted careful analysis of whether the failure to log on affected our results. Since there were no appreciable or systematic effects in this regard, we have included all these decisions in the data we present here. Subjects earned an average of \$20 in the standard game, \$22.53 in the round robin, and \$23.25 in the tournament, before penalties and the \$5 participation bonus. Instructions for all three sessions are available from the authors.⁷

3. Results

Since the responders' behavior is most central to our hypothesis, we begin there. Section 3.2 then examines the behavior of the proposers.

3.1. Responders


Figures 1*a* and 1*b* depict the responders' average demands by round for sessions 1 and 2. The effect of removing fairness from the ultimatum game is immediate upon comparing the series of average responder demands for each treatment.⁸

First note in Figure 1 that responders' behavior in the standard game is qualitatively equivalent to that found by others, despite our different elicitation method. In both session 1 and session 2, the average demand hovers around twenty percent of the 'pie,' and does not appear to change dramatically over the 10 rounds.


decisions in session 2 were the result of failure to log on.

⁷See <http://www.ssc.wisc.edu/~andreoni/> for downloads of instructions.

⁸We also analyzed median (rather than mean) responder demands and proposer offers and found the results to be equivalent.



(a) Session 1



(b) Session 2

Figure 1: Average Responder Demand

Turning to the round robin sessions, we again see that in both sessions the average responder demand hovers around twenty percent, and no systematic change occurs over the 10 rounds. The theoretical prediction of whether responders should behave differently in the standard and round robin treatments is unclear, since the possible round robin effects of risk-reduction, increased group-size, and reputation building work simultaneously. Indeed, we find no consistent or statistically significant difference between the two treatments,⁹ suggesting that standard ultimatum game play appears robust to changes introduced by the round robin format.

The central finding of this paper is immediate upon examination of responders' average demands in the tournament game. Switching to the tournament setting and thereby eliminating fairness concerns from the ultimatum bargaining game hastens responders' approach to the subgame perfect prediction. To see this, compare the series of tournament-game average demands with the two control series. Allowing for learning to take place over the first several rounds, we find that responders' demands were substantially lower under the tournament game than the two control treatments. In both sessions 1 and 2 the average responder demand for the tournament game fell below that of the other two treatments by the sixth round. Moreover, the average responder demand in both sessions reached ten percent by the last round, half that of the standard and round robin games. These differences are statistically significant.¹⁰

Additional evidence is presented in Figures 2*a* and 2*b*. The frequency of subgame perfect demands by responders is compared across the three treatments for the two sessions. Subgame perfect play is defined as a demand of either zero or one token (the smallest unit of exchange in this model). The evidence from session 1 is not overwhelming; the number of subgame perfect demands is relatively higher for the tournament game, but not significantly so—in most rounds only one tournament game responder made a subgame perfect demand, whereas the

⁹Demands are lower than the standard treatment in session 1, but higher in session 2. Neither difference is statistically significant. For session 1, a heteroskedastic t -test reveals $t = 1.1$ ($\alpha \leq 0.28$), and for session 2 $t = -0.14$ ($\alpha \leq 0.89$). We recognize that observed average demands across rounds are not independent due to subjects' behavior in earlier rounds, but argue that subject pool dependency makes it more, not less, likely that average responder demands would differ across treatments over the course of 10 rounds. Thus, these heteroskedastic t -tests are, if anything, biased against finding that the observed standard and round robin behavior is statistically indistinguishable. The subsequent statistical tests mitigate (but do not eliminate) dependency contamination by restricting attention to within round comparisons across games.

¹⁰For session 1, comparing the round robin to the tournament, we see $t = 2.18$ ($\alpha \leq 0.048$), while for session 2 $t = 1.78$ ($\alpha \leq 0.092$).

other two games typically saw no such demands. Figure 2*b* is more compelling, however. Here, the number of subgame perfect demands is not only substantially higher for the tournament game, but also increases systematically over the 10 rounds.¹¹ Nonetheless, the proportion of responders making subgame perfect demands reaches a maximum at just 40 percent.

These results support the hypothesis that removing fairness from the ultimatum bargaining game has a substantial impact on responders' demands, hastening convergence to the subgame perfect prediction. The importance of allowing for learning is also quite clear; although no meaningful differences in behavior among the three treatments can be seen in the first several rounds of each session, the impact of changing the games' structure is apparent by the latter half of each session. Still, learning is not complete among tournament players, as a majority do not choose subgame perfect moves by the end of the study.

3.2. Proposers


The behavior of proposers in the tournament game differs substantially between sessions 1 and 2. We therefore begin with a discussion of session 1, which has results as expected, before turning to an examination of session 2.

Figure 3*a* presents the average proposer offers by round for each game in session 1. Note first that the results from the standard game are again consistent with findings by Slonim and Roth (1998) and others; the average proposer offer hovers around 40 percent, and no systematic changes appear over the course of the 10 rounds.


Turning to the round robin session, there now may be some reason to expect round robin offers to be below the offers of the standard game. For a given distribution of priors on responders' demands, an offer in the round robin game carries a convexified distribution of ex post payoffs. That is, given identical distributions of responses and an identical offer, the standard and round robin games have the same expected payoff but the round robin game has lower risk. Hence, if risk aversion is significant, the round robin proposers may be willing to make more aggressive offers.¹² Figure 3*a* shows that average round robin proposer behavior is consistent with this hypothesis. The difference between these two series is nominal, however, and so this point should not be overstated.

¹¹A test for equal mean rates of subgame perfect play in the last round between the round robin and tournament games can be rejected only at the 34 percent level in session 1 ($t = -1.0$), but at the 3.7 percent level for session 2 ($t = -2.45$).

¹²See Andreoni, Castillo and Petrie (2003, 2004) for evidence of this effect in a ultimatum games.



(a) Session 1



(b) Session 2

Figure 2: Subgame Perfect Responder Demands by Round


Looking next at the tournament condition, session 1 proposers' behavior is compatible with both the responders' behavior and theoretical predictions. Again focusing on the later rounds of each game, it may be seen clearly that the average proposer offer reaches a significantly lower level under the tournament game than for the two control treatments.¹³ This suggests that proposers in the first tournament game learned that responders' demands declined over the course of the session, and that they decreased their offers accordingly. That is, proposers' behavior in session 1 follows that of the responders, as we would expect, and hence approaches more closely the subgame perfect prediction by the later stages of the game. Indeed, the average offer in the final round of the tournament game fell to 19.8 percent, well below that of the standard and round robin treatments at 40.3 and 39.6 percent, respectively.

Turning to session 2, the story becomes more complicated as shown by Figure 3b. First notice that the results from the standard and round robin games are qualitatively the same as those from session 1; average proposer offers hover around the 40 percent level in both conditions, and do not vary dramatically over the course of the session. Recall also that responders' behavior was comparable across the two sessions for all games. Thus the discrepancy between the two sessions lies only in the average offers observed in the tournament game.


Unlike session 1, the average proposer offer for the session 2 tournament game does not decline as the game progresses. In fact, the average proposer offer is higher for the tournament game than for either of the control treatments for the duration of the session. The last round average proposer offer in session 2 is 49.8 percent, compared to 35.4 and 39.2 percent for the standard and round robin treatments, respectively. At the same time, the average responder demand for session 2 tournament falls to just 10.4 percent.

Why did the proposers fail to decrease their offers in session 2, as they did in session 1? An intriguing possibility suggested by our data is a potential difference in the learning environment across sessions. For example, Gale, Binmore, and Samuelson (1995) present a model of evolutionary learning in the ultimatum game where high variance in the behavior of responders makes learning more difficult for ultimatum game proposers— the more noisy the demands of the responders, the longer it takes for proposers to adopt a subgame perfect strategy. Thus, even in the absence of fairness concerns, subgame perfect predicted play may not be observed until the 'ultra long run,' a horizon potentially encompassing hundreds or thousands of rounds. Models of reinforcement learning would generate similar predictions.

¹³In round 10, comparing the round robin with the tournament, for instance, we find $t = 4.8$ ($\alpha \leq 0.00012$).



(a) Session 1



(b) Session 2

Figure 3: Average Proposer Offer by Round

To the extent that responders' demands were noisier in session 2 than in session 1, the Gale-Binmore-Samuelson theory of imperfect learning may account for the high offers observed in the session 2 tournament game. First we compare the characteristics of responders' demands for the two sessions statistically. Comparing responders' demands across rounds for the two tournament game sessions, we find strong evidence that the conditions for learning were significantly worse in session 2. Though the average responder demand was roughly the same in session 1 (16.4 percent) and session 2 (17.3 percent), the standard deviation increased dramatically from session 1 (9.82 percent) to session 2 (17.53 percent). This increase in variance is statistically significant.¹⁴ Comparing the higher moments of the two distributions simply reinforces that session 2 responders behaved more noisily than their session 1 counterparts.


How the higher variance in responder demands translates to a difficult learning environment is demonstrated in the next figure. Here we calculate the best response offer given the exact distribution of responder demands for each round for the tournament games in the two sessions. Figures 4*a* and 4*b* chart the best response, observed average offer, and the offer by the top ranked proposer for each round of the tournament game in sessions 1 and 2, respectively.

In studying the figures, first notice that the best response function is considerably more variable in session 2, reflecting the relatively difficult learning environment. Next notice that the first-ranked decision in session 1 matches the best response function closely, suggesting that at least one proposer each round in session 1 was able to select the optimal strategy.¹⁵ In contrast, the first-ranked offer in session 2 deviates widely from the best response function, perhaps because the best response function itself seems to shift dramatically. As a result, it appears that session 2 proposers were unable to hone their strategies as effectively as their counterparts in session 1.

The distribution of responders' demands in the session 2 tournament game may thus explain why the average proposer offer failed to decline towards the expected subgame perfect prediction within 10 rounds. Indeed, the finding that proposers reacted more slowly to responders' declining average demands under higher variance conditions seems to endorse the Gale-Binmore-Samuelson story of imperfect learning.

¹⁴The one-tailed F-test for equal variance in responders' demands for the two sessions overwhelmingly rejects the null in favor of greater session 2 variance, with $F = 3.18$ ($\alpha \leq 1.07 \times 10^{-8}$).

¹⁵We, of course, do not know whether this was accidental or the result of savvy play. Nonetheless, it does create an environment that helps all players find the optimum more readily than had the best response been missed each round.



(a) Session 1



(b) Session 2

Figure 4: Tournament Treatment Best Response and First Ranked Offer by Round

4. Learning vs. Fairness

Though the learning process was certainly left unfinished by the last round of our experiment, we can nonetheless use the results to approximate the proportion of non-subgame perfect play attributable to fairness versus incomplete learning within the relatively short experimental session. To avoid the strategic complications of proposers' behavior discussed in the previous section, we focus on responders' demands, which are made in a subgame. The optimal responder demand is independent of proposers' offers, and thus provides us a simple metric by which to measure the deviation of players' demands from the subgame perfect prediction of zero. Additionally, we combine the spring and fall session data, since the responders' demands are very similar across the two sessions.

Figure 5 depicts the average responder demand by round, for the standard, round robin, and tournament games. Recall that the influence of fairness concerns is isolated by comparing the round robin and tournament treatments, since the only difference between the two is that subjects are paid absolute earnings in the former and by within-type rank in the latter. Thus, the difference between responders' demands in the round robin and tournament games provides a rough measure of how much fairness is in the system.




Figure 5: Average Responder Demands, Sessions 1 and 2 Combined

It is clear from Figure 5 that in the first 5 rounds imperfect learning, or confusion, swamps any effect of fairness concerns on responders' play; the average

responder demands were roughly equal across treatments as subjects learned the mechanics of the ultimatum game, and tournament subjects additionally learned about the tournament. In the latter rounds, however, the tournament demands separate from those of the control treatments, suggesting that fairness is responsible for about 40 percent of round robin responders' demands over the last 5 rounds. By the final round, fairness accounts for about half of responders' deviation from subgame perfect play, with imperfect learning accounting for the rest.

These results coincide with the findings from Andreoni (1995), in which altruism and confusion were found equally responsible for deviations from dominant strategy play in a linear public goods game, played over 10 periods. In his study, however, confusion (or imperfect learning) declined to about 13 percent by the tenth iteration. That the learning process appears slower in our experimental session than in Andreoni (1995) underscores further the point of Gale, Binmore, and Samuelson that the ultimatum game presents a more challenging learning environment than do other games.

There is one very important way that the data in Figure 5 differ from the public goods data studied by Andreoni (1995). His data showed a slow but distinct tendency in the standard game toward dominant strategy play. He concluded that fairness concerns were giving way to free riding as subjects became less confused.¹⁶ Notice that neither the standard nor round robin games show any tendency toward more sub-game perfect play as time goes forward. If subjects in the standard and round-robin games are learning sub-game perfection at the same pace as those in the tournament game, then clearly they are not using this knowledge to behave more rationally. Instead, fairness is stubbornly dominating their actions.

5. Conclusion

We present an experiment designed to separate the two commonplace explanations for behavior in ultimatum games—subjects' concern for fairness versus the failure of subgame perfection as an equilibrium refinement. We employ a tournament structure that subtracts the potential for fairness to influence behavior. Comparing the results of the tournament game with two control treatments affords us a clean test of subgame perfection as well as a measure of fairness-induced play in ultimatum game experiments over relatively short horizons.

¹⁶See also Houser and Kurzban (2002).

We find greater support for the subgame perfect prediction in the tournament condition than in the standard ultimatum game. Even so, there was still substantial room for further learning by responders after 10 rounds of play. The data also suggest that the learning environment for proposers was more difficult in our session 2 than in session 1. It appears that the noisy distribution of responder demands in the session 2 tournament game made learning more difficult, such that the horizon necessary for learning to become effective surely extends beyond the 10 rounds tested in our sessions. Still, we are able to identify that imperfect learning generates offers that are about 50–60% of the level seen in the standard and round-robin games, with a clear trajectory toward lower offers. In the mean time, offers in the two control treatments show no tendency toward lower offers, indicating that when subjects learn sub-game perfection, fairness still dominates their choices.

We interpret the findings above as evidence that learning the sub-game perfect strategy in a standard ten-period ultimatum game is extremely difficult. In this way, our results confirm the theoretical predictions from models of learning in this game. More data on 10-period games would, we expect, continue to confirm this view. Moreover, this data is sufficient for us to answer the major question of this research—how important is learning vis-à-vis fairness. We have been able to confirm that indeed learning is a major component of choices in 10-period ultimatum games, and roughly equal in significance to fairness. Moreover, we are able to conclude that, even if subjects in ultimatum games understand sub-game perfection, their choices are nonetheless governed by concerns for fairness rather than rational adherence to game theory.

Our research design also raises an important methodological question. We used a “remote” or one-decision-per-day design, with subjects entering decisions on a web-site. This was done with the belief that it could foster better learning by giving subjects more time to consider their decisions. It came at the cost of losing some control over the subjects, despite several safeguards meant to minimize these concerns. We found that subjects in our game behaved remarkably similarly to participants of 10-period games played in a single lab sitting, suggesting that neither the extra time nor weakened controls made a difference. Without direct controls on the setting, however, the effect of the remote design remains an open question. We view this as an interesting topic for further research.

What is the next avenue for research on the question learning in ultimatum games? We think that evidence presented in this paper, along with results from Slonim and Roth (1998) and Roth, et.al.(1991), indicates that more data on

10-period ultimatum games may yield little additional information on learning. Instead, research should next turn to pushing the boundaries of learning and fairness. In particular, just how much time and exposure is needed for learning to be completed in this game and when, if ever, will subjects come to understand sub-game perfection in the ultimatum game? Given that learning will eventually take hold, will fairness still reign over selfish sub-game perfection as is does in the shorter run games we and others have explored? To answer these questions will require all new studies, and in particular studies that experiment with the very long run, allowing dozens or perhaps hundreds of rounds.

References

- Abbink, K., G. Bolton., A. Sadrieh, and F. Tang. "Adaptive Learning versus Punishment in Ultimatum Bargaining," *Games and Economic Behavior*, 37(1), October 2001, 1–25.
- Andreoni, J. "Cooperation in Public-Goods Experiments: Kindness or Confusion?" *American Economic Review*, 85(4), September 1995, 891–904.
- Andreoni, J., M. Castillo, and R. Petrie. "What do Bargainers' Preferences Look Like? Exploring a Convex Ultimatum Game," *American Economic Review*, 93(3), June 2003, 672–685.
- Andreoni, J., M. Castillo, and R. Petrie. "Revealing Preferences for Fairness in Ultimatum Bargaining: Parametric and Non-Parametric Estimates of Preferences," working paper, University of Wisconsin, 2004.
- Bolton, G. "A Comparative Model of Bargaining: Theory and Evidence," *American Economic Review*, 81(5), December 1991, 1096–136.
- Bolton, G. and A. Ockenfels. "ERC: A Theory of Equity, Reciprocity, and Competition," *American Economic Review*, 90(1), March 2000, 166–93.
- Camerer, C. and T. Ho. "Experience-Weighted Attraction Learning in Normal Form Games," *Econometrica*, 67(4), July 1999, 827–74.
- Erev, I. and A. Roth. "Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term," *Games and Economic Behavior* 8(1), January 1995, 164–212.
- Fehr, E., and K. Schmidt. "Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, 114(3), August 1999, 817–68.
- Fudenberg, D. and D. Levine, "Measuring Players' Losses in Experimental Games," *Quarterly Journal of Economics*, 112(2), May 1997, 507–36.
- Gale, J., K. Binmore, and L. Samuelson. "Learning to Be Imperfect: The Ultimatum Game," *Games and Economic Behavior*, 8(1), January 1995, 56–90.
- Goeree, J. and C. Holt. "Asymmetric Inequality Aversion and Noisy Behavior in Alternating-offer Bargaining Games," *European Economic Review*, 44, May 2000, 1079–89.
- Guth, W., Schmittberger, R., and Schwarze, B. "An Experimental Analysis of Ultimatum Bargaining," *Journal of Economic Behavior and Organization* v3, n4, December 1982, 367–88
- Houser, Daniel and Robert Kurzban, "Revisiting Kindness and Confusion in Public Goods Experiments," *American Economic Review*, 94 (4), November 2002, 1062–1069.

- Li, Jing, “A Model of Ethical Behavior,” working paper, University of Wisconsin–Madison, 2004.
- Ochs, J. and A.E. Roth, “An Experimental Study of Sequential Bargaining.” *American Economic Review*, 79(3), June 1989, 355–84.
- Rabin, M. “Incorporating Fairness into Game Theory and Economics,” *American Economic Review*, 83(5), December 1993, 1281–1302.
- Rabin, M. and G. Charness. “Understanding Social Preferences with Simple Tests,” *Quarterly Journal of Economics*, 117, August 2002, 817–861.
- Roth, Alvin E., Prasnikar, V., Okunofujiwara, M. and Zamir, S. “Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study.” *American Economic Review*, December 1991, 81(5), 1068–95.
- Slonim, R. and A. Roth. “Learning in High Stakes Ultimatum Games: An Experiment in the Slovak Republic,” *Econometrica*, 66(3), May 1998, 569–96.