Revealing Preferences for Fairness in Ultimatum Bargaining*

by

James Andreoni
University of Wisconsin

Marco Castillo
Georgia Institute of Technology

Ragan Petrie
Georgia State University

February 2004
Current version July 2, 2004

Abstract: The ultimatum game has been the primary tool for studying bargaining behavior in recent years. However, not enough information is gathered in the ultimatum game to get a clear picture of responders' utility functions. We analyze a convex ultimatum game in which responders' can "shrink" an offer as well as to accept or reject it. This allows us to observe enough about responders' preferences to estimate utility functions. We then successfully use data collected from convex ultimatum games to predict behavior in standard games. Our analysis reveals that rejections can be "rationalized" with neo-classical preferences over own- and other-payoff that are convex, nonmonotonic, and regular. These findings present a precise benchmark for models of fairness and bargaining—modellers need no longer guess about the underlying shape of responders' preferences.

*Andreoni thanks the National Science Foundation and the Russell Sage Foundation for financial support. We are grateful to Pam Welter and Mark Zeller for expert research assistance, and Gary Bolton and David K. Levine for helpful comments.

1. Introduction

Perhaps the most important tool for studying bargaining over the last 15 years has been the ultimatum game. This game has been explored in hundreds of laboratory experiment, in dozens of countries, and for stakes that range from the trivial to the profound. In each case the results are similar: Fairness in the minds of responders has a significant effect on allocations.[1]

The ultimatum game is a two-stage interaction. In the first stage a proposer offers a division of a pie. In the second stage the responder can accept the proposal or reject it. A rejection means that both players get nothing. Sub-game perfection indicates that all offers should be accepted, and thus the smallest possible offer should be made. Nonetheless, responders tend to reject "unfair" offers, and thus proposers tend not to make them. Recent studies show that behavior remains quite far from the prediction even when using multiple trials and large stakes.[2]

The challenge to economists has been to understand this outcome within the context of an economic model. What kind of preferences or beliefs could result in this type of behavior? While several scholars have proposed models of fairness that could be consistent with the choices of subjects, getting a clear picture of the preferences of bargainers has been elusive. One reason is that an ultimatum game experiment collects too little information from responders to learn about their utility functions. With the dichotomous choice of accept or reject, it is difficult, even over multiple games, to learn enough about the objective functions of bargainers to cleanly test hypotheses about their behavior.

This paper will recover bargainers' preferences by employing a convex ultimatum game.[3] The convex ultimatum game contains the regular ultimatum game as

[1] For important early contributions see Werner Guth et al. (1982), Jack Ochs and Alvin E. Roth (1989), Alvin Roth et al. (1991), and Robert Forsythe et al. (1994). For reviews of the literature on bargaining and ulitmatum games, see Alvin E. Roth (1995) and Colin Camerer (2003). For important recent contributions, see Robert Slonim and Alvin E. Roth (1998) on large stakes, Catherine C. Eckel and Philip Grossman (2001) on gender differences, and William Harbaugh et al. (2000) on bargaining of children. Uri Gneezy, Ernan Haruvy, and Roth (2003) look at ultimatum bargaining under a deadline, where proposers' offers can be revised before they are rejected.

[2] See Slonim and Roth (1998) and Roth, Prasnikar, Okunofujiwara and Zamir (1991). These studies show that while repetition and large stakes both move players in the direction of subgame perfection, few outcomes actually reach the prediction.

[3] The convex ultimatum game was explored by Andreoni, Castillo and Petrie (2003) in a single shot framework. There have been other convex games introduced in the literature. Matthew Rabin (1997) discussed the game theoretically (called the "squishy game"). Ramzi Suleiman

a special case, has the same subgame perfect equilibria, but allows us to measure much more about responders' preferences than the regular ultimatum game. As before, the proposer offers a split of the pie. Then the responder chooses how much to "shrink" the pie. Not shrinking at all is to accept the offer, shrinking to zero is to reject it, but in the convex game shrinking to some intermediate level is also possible. The advantage of the convex game is that one can learn more about the shape of each bargainer's utility function and, by adding some minor structure to the problem, estimate whether well-behaved preferences could have generated behavior. With multiple observations from each subject, one can get a picture of what these preferences look like.

We examine subjects who play a series of 20 regular or convex bargaining games, each with a randomly changing partner. We find that subjects' choices can be captured by preferences that are convex, regular, but not monotonic. Moreover, we are able to estimate utility functions for each of our subjects in the convex games. When we ask whether these estimated utility functions can predict the behavior in regular ultimatum games, we get a strikingly strong fit.

The approach also highlights the great deal of heterogeneity across subjects. Only about 12 percent of respondents act like money-maximizers, and only 25 percent have "linear" preferences that would lead them to accept or reject (but not "shrink") all offers. The rest have strictly convex preferences with indifference curves that, to varying degrees, bend back at the extreme allocations.

Our findings represent the first direct look into the preferences of responders. For theories of fairness, they provide a more precise benchmark that the models must meet. We find that our results can be used to draw sharp distinctions among several approaches to fairness. For models of bargaining, the wide heterogeneity of preferences suggests models of incomplete information, along with fairness, may be necessary to understand ultimatum game behavior.[4] It also suggests a possible role for new bargaining institutions that take advantage of natural tastes for fairness to promote economic efficiency.

The paper is organized as follows. The next section outlines the theoretical implications of the convex ultimatum game. Section 3 describes our experiment. Sections 4, 5 and 6 present the results for both proposers and responders, while section 7 is a conclusion.

(1996) presents a game simliar to ours, but instead of letting the responder choose how much to shrink the pie, that is fixed by the experimenter. Marlies Ahlet et al. (2001) also propose a convex game, where the amout of pie shrinkage is bound from above by the show-up fee and offer amount. So, only offers of 50-50 can be shrunk to zero.

[4]See, e.g., Kennan and Wilson (1993) for a review.

2. Theory and Hypotheses

Figure 1 illustrates the convex and regular ultimatum games. Formally, let M be the total number of dollars to be bargained over. The proposer chooses a division $a \in [0, 1]$, and the responder chooses a number of dollars to divide m such that the payoffs for the proposer and responder are, respectively,

$$\begin{aligned} \pi_p &= (1 - a) \times m \\ \pi_r &= a \times m. \end{aligned}$$

By restricting m to be either 0 or M we have a regular ultimatum game. That is, the proposal is fully accepted or fully rejected. In the convex ultimatum game, by contrast, we allow m to be any number between 0 and M. Thus, the regular ultimatum game is nested within the convex game. Moreover, the standard subgame perfect equilibrium is the same in both—since no offers should be shrunk, minimal offers should be made.

Figure 1 shows how a responder who dislikes inequality may react differently to the same offer in the two games. It shows that a proposal that would be rejected in the ultimatum game will be shrunk in the convex game. Similarly, some proposals that would be accepted in an ultimatum game will also be shrunk in a convex game.[5]

[5] The graph also makes the case for the presence of disadvantageous counterproposals (Ochs and Roth, 1989). Notice that when the responder's preferences are nonlinear, but not necessarily monotonic, the responder would prefer an allocation that gives him a smaller amount of money, provided the proposer's payoff is also reduced.

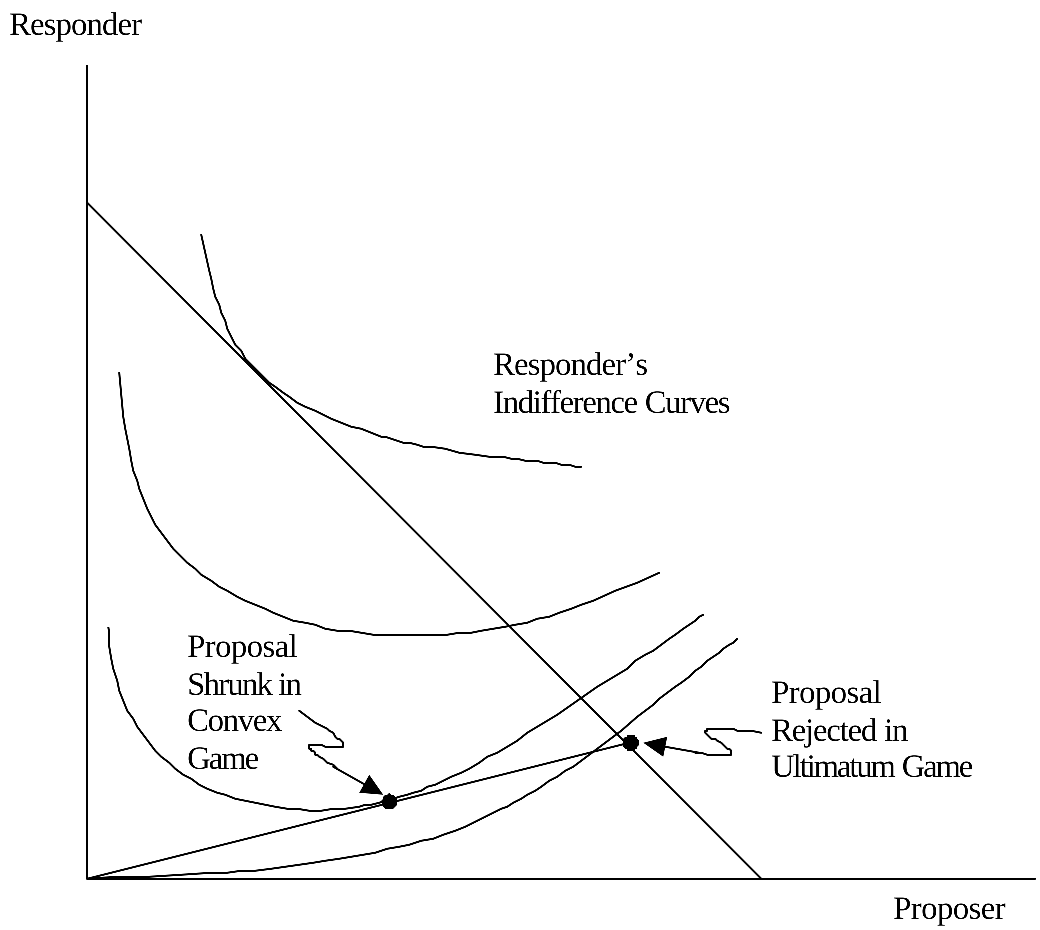

Figure 1: Responder Choices in Convex and Standard Ultimatum Games

There are several clear and systematic differences in the two games. These can be seen in Figure 2. First, we can graph the best reply functions of a responder in the two games. We see first that in the ultimatum game there will be a critical offer at which a responder switches from reject to accept. With the convex game the movement between the two extremes is continuous, with some offers being shrunk. Hence, the best-reply function in the convex game will tend to cross that of the regular game from below, as shown in the figure.

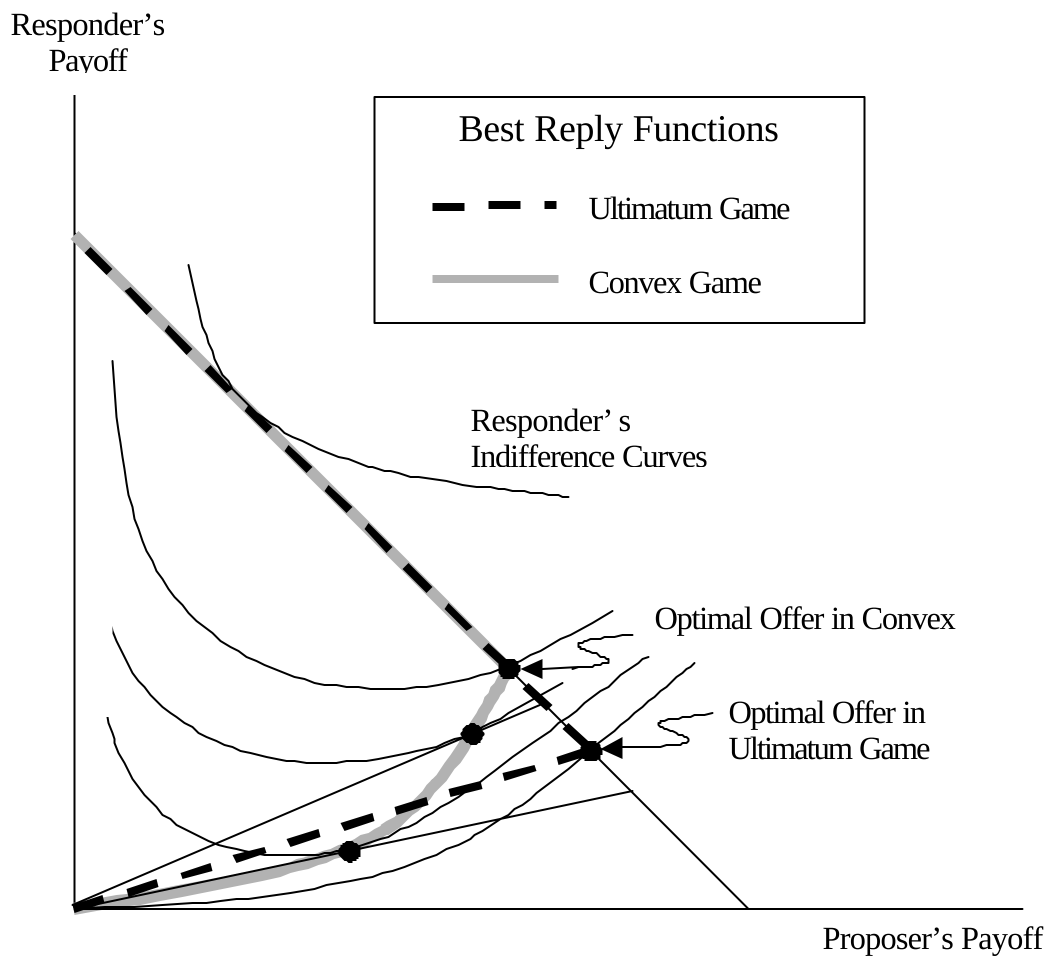

Figure 2: Best Reply Function

Figure 2 also illustrates how the convex game shifts bargaining power to responders. A money-maximizing proposer will tend to make higher offers in the convex game than in the standard game. Notice that, while this means that the responders will always have higher utility in the convex game, they will not necessarily end up with more money. A responder may end up accepting offers in the ultimatum game that she would prefer to shrink were she playing the convex game. Hence, the prediction about relative earnings is ambiguous, even though the prediction on relative utility is not.[6]

[6] See Matthew Rabin (1997) for more formal consideration of this game. It is also interesting

Note as well that there is nothing that precludes bargainers who care about inequality of all kinds from shrinking extra-generous offers. A responder who, for instance, is offered 90 percent of the pie may find it immodest to accept this, and may shrink or reject the offer. As we will see, we actually observe behavior consistent with this for some subjects. Such behavior, while uncommon, is consistent with findings elsewhere.[7]

The most important difference in the two games, however, is the added data that is gathered in the convex game. By observing offers that are shrunk we can gain valuable information about the shape of responders' indifference curves that will be useful in uncovering preferences in the standard ultimatum game.

3. Experimental Design

Our data was collected at the University of Wisconsin. Subjects were volunteers from undergraduate economics and business classes. A total of 96 subjects participated in the experiment, 48 in the standard ultimatum game and 48 in the convex ultimatum game. Each session of the experiment lasted about one hour, and subjects earned on average $17.96 (*s.d.* $4.50), plus a $5.00 show-up fee.

Each session of the experiment required 24 subjects. Subjects were assigned one role, proposer or responder, which they kept throughout the experiment. They played 20 rounds of the bargaining game. Each round was played with a randomly chosen anonymous partner. Subjects never knew which person they were playing with, and were guaranteed never to play the same partner more than twice. They were paid in cash at the end of the study.

In each game, the proposer and responder bargain over how to divide 10 quarters (the US 25-cent piece). The proposer offered a division of each quarter, allocating from 0 to 25 cents to the responder and the rest to himself. The responder, upon seeing the division offered by the proposer, then decides how many of the quarters to divide. In the standard ultimatum game, responders could choose either 0 or 10 quarters. In the convex ultimatum game, the choice was any

to note that the indifference curves plotted in Figures 1 and 2 would be consistent with preferences proffered by Rabin (1993). This is discussed in more detail in Andreoni, Castillo and Petrie (2003).

[7] Andreoni and Miller (2002) found about 25 percent of subjects disliked even favorable inequality. A similar fraction was found in Andreoni, Castillo and Petrie (2003). Rabin (1993), Fehr and Schmidt (1999) and Bolton and Ockenfels (2000) cite evidence of similar effects in many other experiments.

number of quarters from 0 to 10. With this one exception, the two games were identical. All of the parameters of the experiment were known to all subjects.

Implementation of the experiment was standard. Subjects were assigned subject numbers, and names were never recorded. All interactions took place on a computer network. Subjects read the instructions fully and were also taken through several examples of how payoffs were calculated. Their earnings were placed in a closed envelope and presented to them at the end of the study. Full instructions for the games are available from the authors.[8]

Notice that each iteration of the game is a bargain over \$2.50, that is, 10 quarter-dollars. Thus, over the 20 iterations, subjects bargained over the division of \$50.

We ran two sessions of each the regular and convex ultimatum games. Hence, there are 24 proposers and 24 responders in our data from each of the games.

4. Instrument Check

We first check how our repeated standard ultimatum game compares to previous research. The average response function, or the probability of accepting any given offer, is similar to the results from Roth, Prasnikar, Okunofujiwara and Zamir (1991) and Slonim and Roth (1998), who both use a 10-round standard ultimatum game. Acceptance rates in Roth et al. start out at 30 percent for offers of 30 percent and jump to 70 percent for offers of 40 percent. In Slonim and Roth, acceptance rates are 55.5 percent for offers in the range of 30–34.5 percent and 76.3 percent for offers in the range of 40–44.5 percent. In our game, offers around 30 percent are accepted 24 percent of the time, slightly lower than Roth et al. and Slonim and Roth. Offers of 40 percent (offer of 10 cents) are accepted 78 percent of the time in our study, slightly higher than Roth et al. and Slonim and Roth.

The distribution of offers in our standard game is also similar to Roth et al. (1991). The modal offer in Roth et al. is 50 percent, whereas in our game the modal offer is 44 percent.

[8] Go to http://www.ssc.wisc.edu/~andreoni/ (Andreoni), http://www.gsu.edu/~ecomec/ (Castillo) or http://www.gsu.edu/~ecorap/ (Petrie).

5. Results for Responders

Figure 3 illustrates the average response to each offer for both the convex and standard ultimatum games. This figure is consistent with the theory presented above. At low offers the ultimatum game leads subjects to reject more than they would like, and at high offers it leads them to accept more than they would like. Indeed, this suggests that the best-reply function (Figure 2) for the convex game is cutting that of the standard game from below.

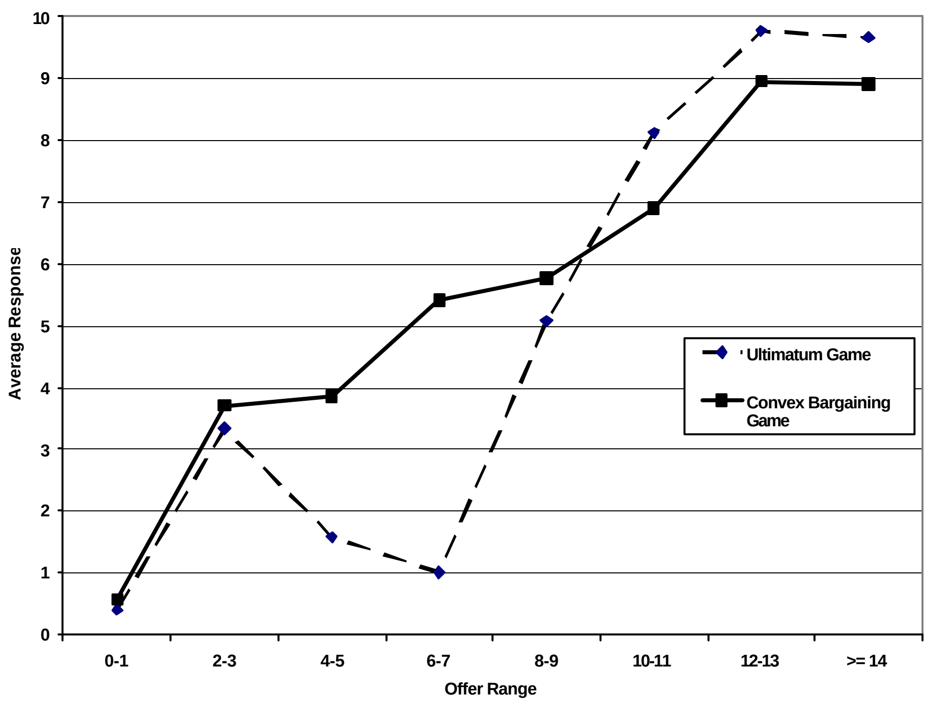

Figure 3: Average Response

Table 1 gives more detail. The Table shows that, as in our second prediction, high offers are made more frequently in the convex game. The ability to partially reject offers shifts bargaining power to responders, thus generating higher offers. Notice, however, that we also observe more density at lower offers in the convex game, especially at the 2 to 7 range. This was not anticipated by the theory section above, but is consistent with our findings in the single-shot version of this game (2003). We hypothesized then that lower offers were less risky in the convex game because so many of the responses were shrunk rather than fully rejected, which made these offers more attractive to proposers. We explore this hypothesis again in the section on proposers.

Table 1

Frequencies of Offers and Responses in Ultimatum and Convex Bargaining Games

	Ultimatum Game				Convex Bargaining Game				
Amount Offered	Freq. of Offer	Frequency of Response 0	10	Average Response	Freq. of Offer	Frequency of Response 0	1 − 9	10	Average Response
0 − 1	25	24	1	0.4	16	15	1	0	0.6
2 − 3	6	4	2	3.3	17	9	2	6	3.7
4 − 5	19	16	3	1.6	37	15	11	11	3.9
6 − 7	10	9	1	1.0	32	9	8	15	5.4
8 − 9	67	33	34	5.1	35	9	10	16	5.8
10 − 11	192	36	156	8.1	147	13	59	75	6.9
12 − 13	131	3	128	9.8	160	3	29	128	9.0
14 − 25	30	1	29	9.7	36	0	9	27	8.9
Total	480	126	354	7.4	480	73	129	278	7.0

A more careful look at Table 1 also reveals some other unexpected observations. First is the curious up-tick in the average response to offers of 2-3 in the ultimatum game. However, since we see only six observations of offers in this range, the effect could simply be due to small numbers. Second, for low offers it seems likely that the conditional probability of choosing 10, that is fully accept, should not be greater in the convex game than in the standard game. However, Table 1 shows that for offers of 0 to 5, for instance, the proportion replying with a 10 is 0.12 in the ultimatum game, but 0.24 in the convex game—those in the convex game are twice as likely to respond with 10 than those in the standard game. It is unclear now whether this result poses a significant deviation from the hypothesis that the same preferences are expressed by responders in both games, or whether it is likely due to random variation that is to be expected when there are only 24 responders in each condition. We will probe this issue more deeply later in the paper.

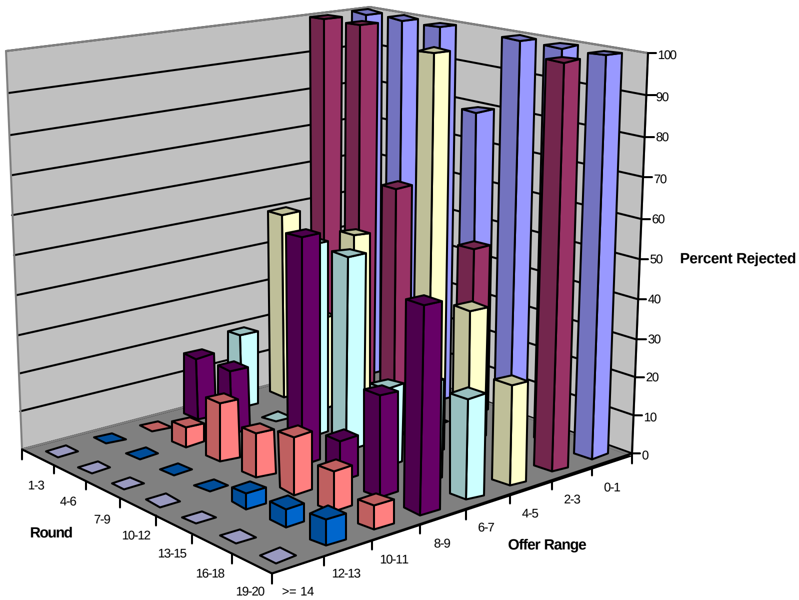

Figure 4: Distribution of Rejections in the Convex Bargaining Game

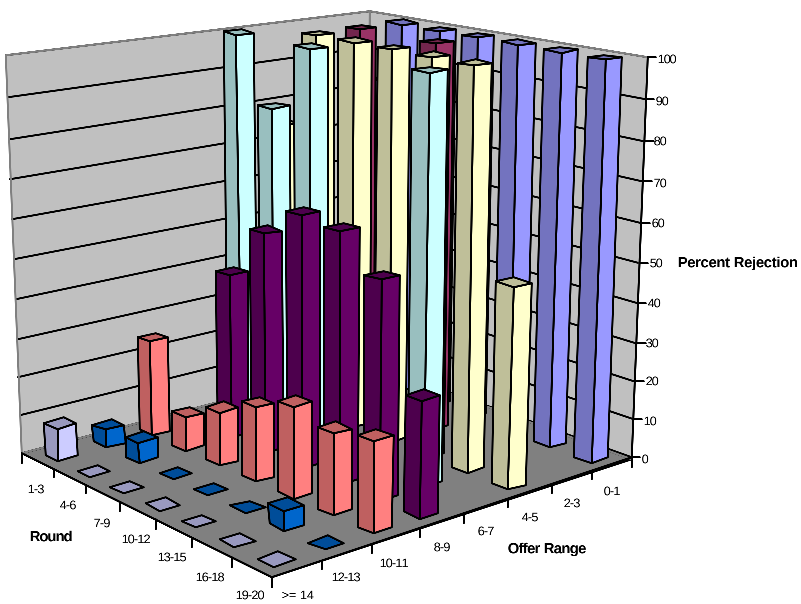

Figure 5: Distribution of Rejections in the Ultimatum Game

5.1. How do responses change over time?

Figures 4 and 5 show how the fraction of offers rejected changes over time for the convex and standard games, respectively. Looking at Figure 4, for the convex game, we see a great degree of variance from round to round, perhaps reflecting the individual heterogeneity in tastes for shrinking the pie. Figure 5, that for the ultimatum game, seems by comparison to be much more orderly. Nonetheless, it is difficult to see any clear progression of rejections over time. This is reminiscent of the findings of Slonim and Roth (1998).

5.2. Are preferences convex and regular?

Next we turn to the individual level data to see if there is some consistency across choices of responders. We see three general hypotheses for the types of preferences that could be underlying choices. First, subjects could have monotonic preferences, that is, preferences that are strictly increasing in own and other's payoffs. These preferences include selfish preferences but also some inequality-averse preferences, as long as indifference curves do not bend back. Monotonicity leads players to accept all offers. Second is non-monotonic utility, but linear indifference curves. For example $U(\pi_s, \pi_o) = \pi_s - \alpha\pi_o$ or $U(\pi_s, \pi_o) = (\pi_s - \pi_o)/\pi_s$ would both have linear indifference curves, the first being parallel lines and the second lines that fan out.[9] These subjects reject offers up to a point and then accept all offers above that point. These players will look the same regardless of whether they play a ultimatum or convex bargaining game. Third is preferences that are strictly convex but not monotonic, as shown in Figures 1 and 2, so subjects may reduce the pie to intermediate levels.

Restricting preferences of the form $u(\pi_s, \pi_o)$ to be strictly convex but not monotonic, however, will not be enough to provide a falsifiable hypothesis about choice. That is, most any set of choices across offers could be generated by convex but nonmonotonic preferences. Hence, we add one more bit of structure to preferences that we call regularity.[10] This assumption is a cousin to normal goods or monotonicity in other settings. To define regular, we first assume that responders have a single most preferred offer.[11] Then we say preferences are

[9] See Levine (1998), Fehr and Schmidt (1999), and Charness and Rabin (2002) for examples of the first type of fairness, and Bolton and Ockenfels (2000) for models with the second type of assumption.

[10] We introduced this concept in our earlier paper (2003).

[11] More generally, we can assume a covex set of most preferred offers.

regular if a person does not shrink an offer more the closer it is to the most-preferred offer. This restriction prevents people from, for instance, fully accepting both low and high offers, but rejecting intermediate offers. While we cannot rule out such preferences as "irrational," we will view them as sufficiently implausible as to cast doubt on our approach. Note that both monotonic and linear preferences meet the definition of convex and regular, as do the strictly convex preferences illustrated in Figures 1 and 2 above.

How do we answer the question, are preferences convex and regular? We first take a nonparametric approach. Adopt the maintained assumption that preferences meet the restrictions of convex and regular, but may vary because of decision errors or learning by subjects. If a subject's choices do not precisely fit the definition of regular preferences, then we ask how much would we have to adjust choices to satisfy regularity. For each subject we then calculate the minimum absolute distance from regular preferences. Since we are measuring distance in the choice space, we scaled the differences to be measured dollar movements. For instance, moving a choice from dividing 7 quarters to dividing 3 would be an absolute distance of 4 quarters, or $1.00. We add these absolute deviations across all budgets. We call this the absolute distance measure. We also report a reweighting of this measure based on deviations in payoffs. That is, if moving a choice from 7 to 3 occurred when offered $0.10 of each quarter, this is counted as costing $0.40, whereas if it happened when offered $0.20 of each quarter it would cost $0.80. This measure gives the best sense of the "cost" of the deviation for a subject who truly has the regular preferences we calculate. We can use these distance measures to get a sense of how well a model of regular preferences fits the data.[12] We then further ask whether these regular preferences are monotonic, linear or strictly convex. The results of this exercise are reported in Table 2.

Look first at the results for the ultimatum game, on the left side of the table. Here, only 3 of the 24 subjects exactly met the definition of monotonic or linear for all 20 decisions (that is, distance is zero). Of the remaining subjects, one was closest to monotonic and the rest were closest to linear. Of course, in the standard ultimatum game it is impossible to tell whether preferences are strictly convex or simply linear, since in both models there should be a critical offer at which responders switch from reject to accept. While there is no natural criterion for deciding when the absolute deviation from regular preferences is high enough

[12] Note that the actual money-equivalent utility cost is likely to be between these two values. Weighting by own-payoff deviations misses the lost utility from altruism or retribution felt by subjects.

to reject regularity, the final three subjects listed appeared to stand out. For this reason alone we are willing to call these subjects not regular.

Table 2
Non-Parametric Categorization of Subjects' Preferences
By Minimum Absolute Distance to Regular Preferences

Ultimatum Game				Convex Bargaining Game			
		Distance Metric				Distance Metric	
Subj	Type	Absolute	Payoffs	Subj.	Type	Absolute	Payoffs
22	Monotonic	\$0.00	\$0.00	5	Monotonic	\$0.00	\$0.00
9	Linear	\$0.00	\$0.00	12	Monotonic	\$0.00	\$0.00
17	Linear	\$0.00	\$0.00	15	Monotonic	\$0.00	\$0.00
20	Near Monotonic	\$2.50	\$1.50	1	Linear	\$0.00	\$0.00
2	Near Linear	\$2.50	\$0.10	10	Linear	\$0.00	\$0.00
3	Near Linear	\$2.50	\$0.50	22	Linear	\$0.00	\$0.00
6	Near Linear	\$2.50	\$0.80	7	Linear	\$0.00	\$0.00
7	Near Linear	\$2.50	\$1.00	24	Strictly Convex	\$0.00	\$0.00
10	Near Linear	\$2.50	\$0.80	6	Near Linear	\$1.25	\$0.35
11	Near Linear	\$2.50	\$0.80	13	Near Linear*	\$3.00	\$1.40
12	Near Linear	\$2.50	\$1.20	21	Near Linear	\$3.25	\$1.28
13	Near Linear	\$2.50	\$0.90	2	Near Strictly Convex	\$0.75	\$0.26
15	Near Linear	\$2.50	\$0.90	17	Near Strictly Convex	\$1.50	\$0.59
18	Near Linear	\$2.50	\$1.20	16	Near Strictly Convex	\$2.00	\$0.95
23	Near Linear	\$2.50	\$0.90	19	Near Strictly Convex	\$2.25	\$0.85
24	Near Linear	\$2.50	\$1.10	23	Near Strictly Convex	\$2.75	\$1.19
5	Near Linear	\$5.00	\$2.20	9	Near Strictly Convex	\$3.25	\$1.28
8	Near Linear	\$5.00	\$2.00	3	Near Strictly Convex	\$3.75	\$1.60
21	Near Linear	\$5.00	\$2.00	14	Near Strictly Convex	\$5.00	\$2.49
24	Near Linear	\$5.00	\$1.10	20	Near Strictly Convex	\$5.00	\$2.15
19	Near Linear	\$7.50	\$3.00	11	Near Strictly Convex	\$8.25	\$2.06
1	Not Regular	\$10.00	\$4.40	4	Not Regular	\$12.25	\$5.54
16	Not Regular	\$10.00	\$3.90	8	Not Regular	\$12.50	\$4.20
14	Not Regular	\$12.50	\$4.20	18	Not Regular	\$13.00	\$4.91

*This subject is equi-distant to linear and strictly convex preferences.

Now turn to the convex game, the right panel of Table 2. By contrast there are eight subjects who exactly fit a utility function over all 20 rounds, three monotonic, four linear and one strictly convex. Contrary to expectations, there were more monotonic and linear subjects in the convex game than in the standard game, even though this gave more opportunity for convex preferences to be expressed. We may, however, be overstating the number of monotonic subjects. The three listed here never received offers below 3 (and for one, 5) so we cannot accurately predict their responses to selfish offers, a point we will return to later. Of the remaining subjects, only three had regular preferences that were closest to linear, while 10 subjects were best described as strictly convex. Overall 46 percent of subjects have preferences that are best measured as strictly convex and regular. As with the standard game, three subjects in the convex game also stood out as not regular.[13]

Figure 6 illustrates four examples of our subject classifications for the convex game. Similar illustrations on all of the subjects are available from the authors' web-sites.[14] The dots are actual choices, while the dotted lines indicate the nearest best reply function under regular preferences. Figure 6*a* shows subject 6 whom we classified as near linear. Here, with the exception of one offer, which was shrunk to 5, choices were exactly linear. Figures 6*b* and 6*c* show two different examples of strictly convex preferences. Subject 2 showed a remarkably smooth pattern of shrinking choices, even many that offered more than 50 percent of the pie. Subject 9 is even more extreme in this behavior. The regular best-reply function that fits this data most closely actually reaches full acceptance at offers of about 50 percent, then shrinks offers above this. Such two-sided inequality aversion, while uncommon, is not without precedent. Our single-shot treatment of this game also found evidence that about 20 percent of subjects had such aversion to even favorable inequality. In this experiment, however, that number is much smaller. Only 8 percent of our subjects showed clear aversion to favorable inequality, although this difference could be explained by the small number of favorable offers made in our study.[15] Figure 6*d* shows a subject whose choices, which appear nearly random, were classified as not regular.

[13] To get an idea of the order of magnitude of these deviations from regular for the six subjects classified as not regular, we can compare their distance to regular to their earnings in the study. For the ultimatum game the earnings (not including the $5.00 show-up fee) were $14.01 (subject 1), $14.80 (subject 14), and $6.40 (subject 14). Likewise, for the convex game earnings were $12.02 (subject 4), $9.01 (subject 8), and $10.29 (subject 18).

[14] See footnote 5 above.

[15] Recall that in the single-shot version of the game, subjects recorded their reactions to all possible offers, not just the offers received. Hence a full spectrum of favorable offers were considered by all subjects. That was not the case here.

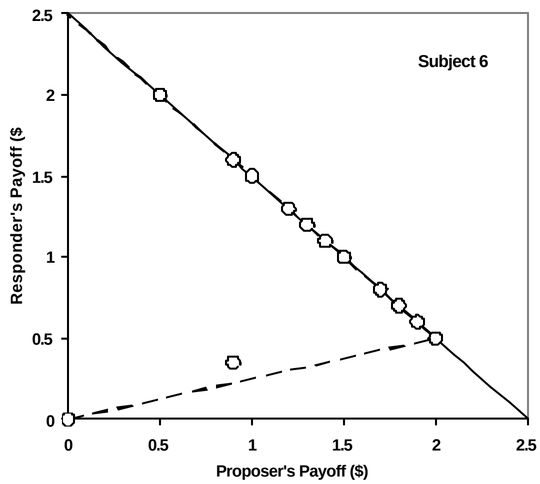

a. Near Linear

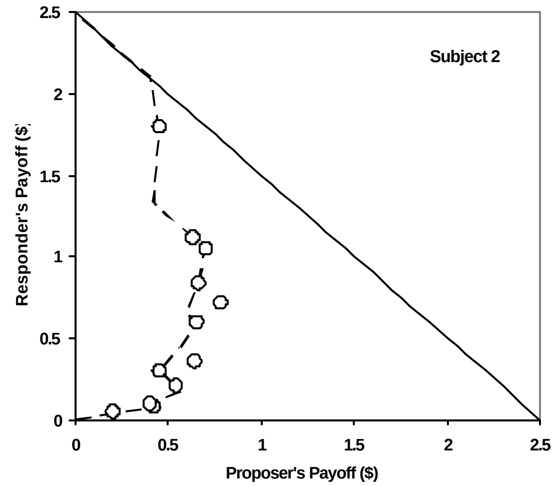

b. Near Strictly Convex

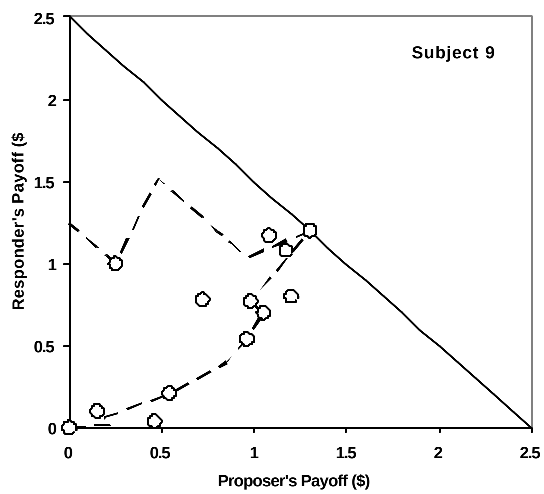

c. Near Strictly Convex

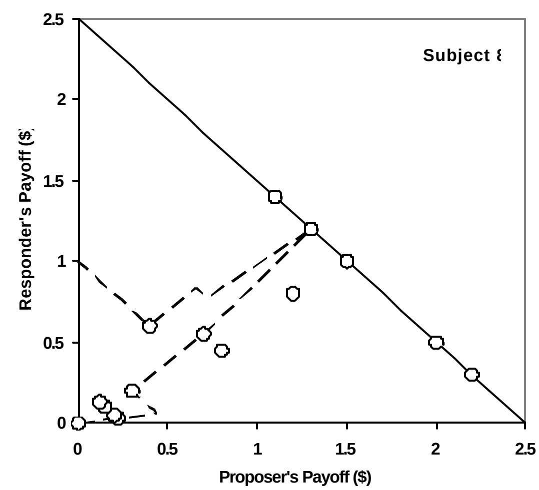

d. Not Regular

Figure 6. Examples of Preference Classifications in the Convex Bargaining Game

We combine the estimates of regular preference of all of our subjects to form a non-parametric estimation of the preference of our subjects as a whole, as represented in best reply functions. This is presented in Figure 7. As can be seen, the average best reply functions match the predictions of the theoretical model, with the curve from the convex game cutting that of the regular game from below. Furthermore, the money maximizing offer in the regular game is 44 percent while in the convex game is 48 percent. Again, this verifies the prediction that the convex game shifts bargaining power to responders.

One can note that the best reply function for the convex game does not meet the 45-degree line. This is due entirely to 4 subjects who, like subject 9 in figure 6, reduce generous offers.

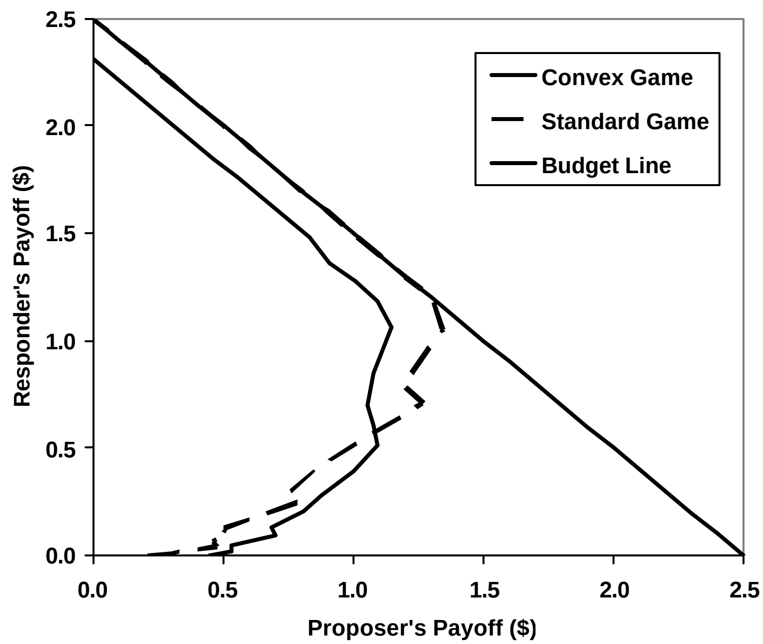

Figure 7: Non-Parametric Response Function

We conclude from this section that regularity is a reasonable organizing criterion to place on the data. In doing so we verify that preferences that are regular can classify almost 90 percent of our subjects. By considering the convex game, we are able to further verify that, of the subjects with regular preferences, half have either monotonic or linear preferences. The remaining half have strictly convex preferences that are not monotonic.

Since we can learn so much about preferences in the convex game, it is worth asking whether the preferences expressed in this game can explain the choices of the subjects in the ultimatum game. While nonparametric results are encouraging, they do not allow us to compare behavior across experiments. It is impossible to know, for instance, whether a responder in the convex game who shrunk a proposal to 5 would accept or reject the same proposal in the standard ultimatum game. To answer this, we will need to estimate utility functions. Thus, we next turn to parametric analysis of the responders in the convex game.

5.3. Parametric Analysis: Estimating Utility Functions

In this section we estimate utility functions for each of our subjects in the convex ultimatum game. For some of the subjects who fit a classification exactly, such a calculation requires no econometrics. For instance, subjects who are monotonic can be described with any monotonic utility function. For subjects whose preferences are exactly the linear specification, it is sufficient to characterize their "switching point." For the other subjects we use maximum likelihood to estimate a random utility model for their choices with a multinomial logit error structure.

In order to allow backward-bending indifference curves we estimated the following quadratic utility function for each subject:

$$U_s(\pi_s, \pi_o) = -(\pi_s - \theta_s)^2 - \beta_1(\pi_o - \theta_o)^2 - \beta_2(\pi_s - \theta_s)(\pi_o - \theta_o). \tag{1}$$

This utility function will produce indifference curve that are ovals. We, of course, do not take seriously the interpretation that (θ_s, θ_o) is the "bliss point" of a subject and that, far away from where choices are made, more for both players yields less utility. We are only trying to approximate the frontier that holds locally in the space of the observed subjects' choices by fitting a simple and easily understood function.

We estimate 14 utility functions like (1), including all the Near Strictly Convex and Not Regular subjects, plus subject 13 (who is also classified as Strictly Convex).[16] Appendix 1 explains how the estimates were conducted. Appendix 2 lists the parameter estimates for each subject, as well as the critical switching points estimated for those with Linear or Near Linear preferences. The overall quality of the fit of the estimation is shown in Figure 8 where we plot the average of the predicted and actual choices for the 24 subjects in the convex game. Note that the data itself shows a great deal of variance between the offers of 3 to 5. Nonetheless, the estimated response seems to track the actual choices quite well. Notice that the prediction fits most poorly at offers from 0–2. This may be due to

[16] We do not estimate the only other Strictly Convex subject (Subject 24) because this subject accepts all offers but an offer of one cent, to which he/she responds with dividing 9 of 10 quarters.

our decision to assume subjects classified as Monotonic would accept these offers, despite having no observations for them in this range. Assuming instead that they would all reject such offers would lower the predicted response by 1.25, thus improving the fit.

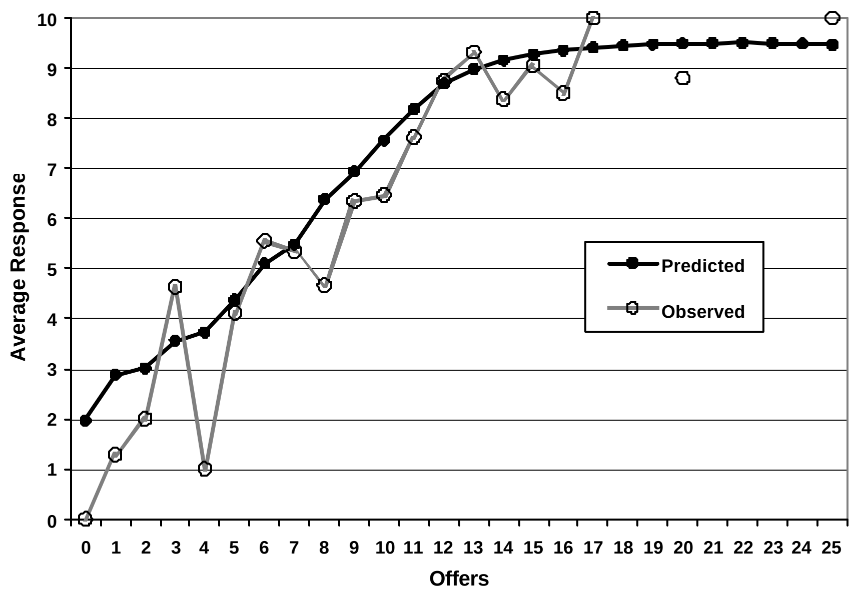

Figure 8: Parametric Prediction and Observed Average Response in Convex Game

Next we apply the estimated utilities of all 24 subjects in the convex game and ask whether this family of utility functions can explain the choices of individuals in the standard ultimatum game. Thus, for each of our estimated utility functions, we determine the probability that the subject would accept or reject any offer presented in the standard ultimatum game, then average this expected response across all 24 subjects to get a predicted expected response in the standard game. Figure 9 plots this prediction against the actual choices in the standard ultimatum game. With the notable exception of three points that have few observations, the fit is nearly perfect, which statistical tests confirm.[17] The Chi-Square test for the

[17] This suggests that the observation in Table 1 regarding low offers (that is, the higher frequency of replies of 10 in the convex game than the standard game) is indeed likely due to random variation in small samples. Again, our prediction for the 0–3 range would improve if we assume the three Monotonic subjects would have rejected these offers.

differences in distributions shows that the two are statistically indistinguishable ($\chi^2[23] = 18.337$, $p = 0.739$).

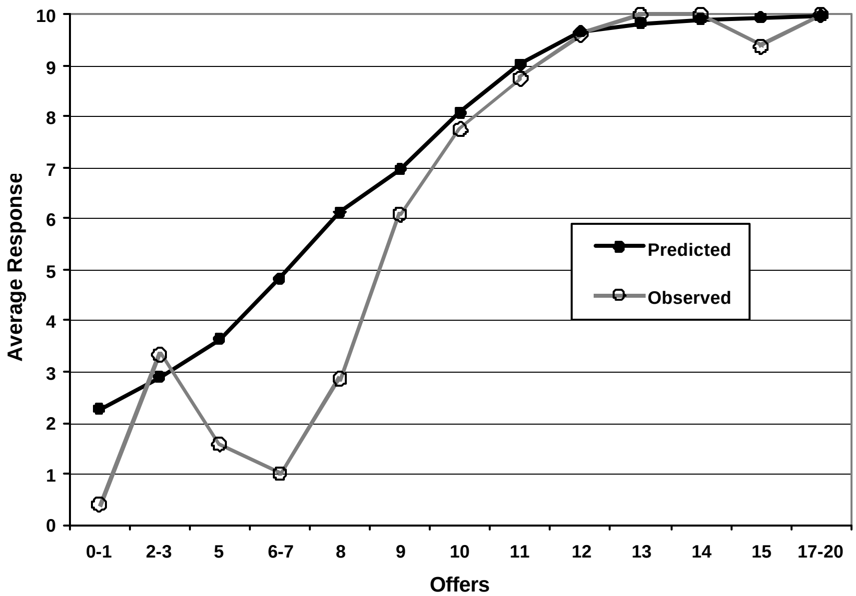

Figure 9: Prediction from Convex Game to Standard Game

Finally, we present one last way to look at the parametric estimates. Figure 10 shows how we would predict our population of estimated utility functions from the convex game would play in either the convex or standard ultimatum game. These reveal that, indeed, the best reply function for the convex crosses that of the standard from below, and that the money maximizing offer in the convex game is 52 percent, which is higher that for the standard game of 48 percent.[18]

[18] Recall that a 50 percent division was not available in our study.

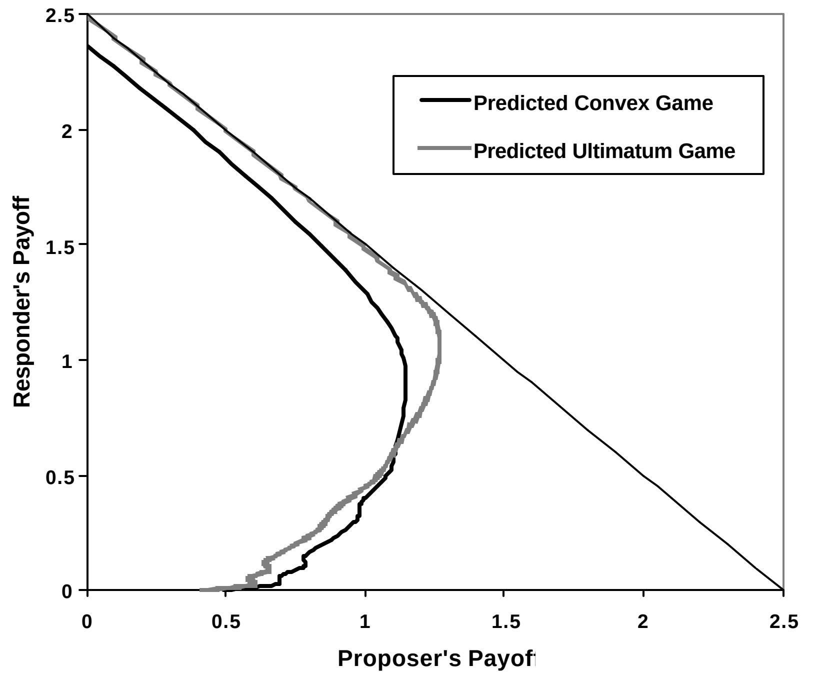

Figure 10: Parametric Predictions for the Average Best-Reply Functions for Convex and Ultimatum Game

6. Results for Proposers

The experiment reported here was designed to learn more about the utility functions of responders. One natural question is whether preferences of proposers are similar to those of responders. Unfortunately neither the convex game nor standard game can give us a pure picture of the preferences of proposers—fear of rejection, expectations on responder types, and fundamental preferences for distributions are all hopelessly conflated in their choices. Moreover, responder preferences also reflect a reaction to the intentions of proposers.[19] Nonetheless, our data does provide some new insights on proposers, which we discuss below.

[19] A better benchmark for calibrating proposer preferences would be dictator game experiments, as in Andreoni and Miller (2002), since, as shown in Andreoni, Brown and Vesterlund (2002), preferences over allocations for responders depend on the intentions and opportunities of the proposers. Also see Manski (2002) for a detailed discussion of the identification problems associated with proposer-response games.

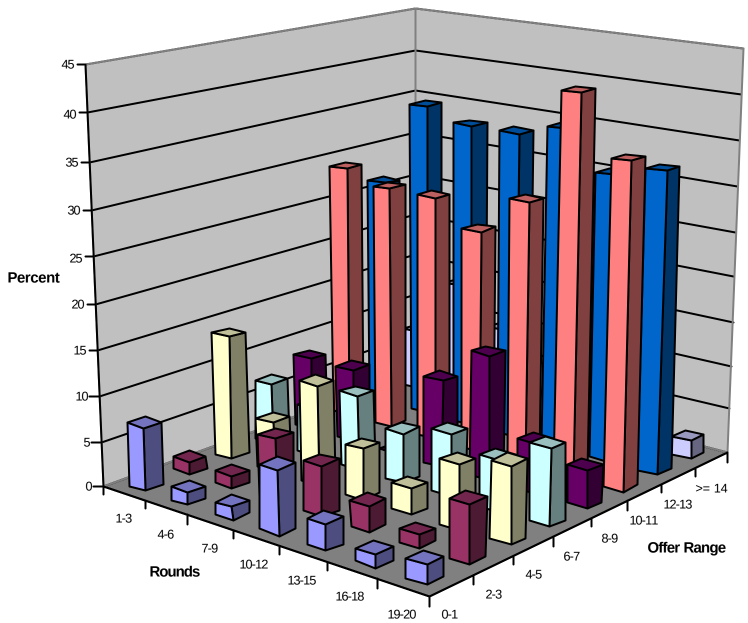

Figure 11: Distribution of Offers in the Convex Bargaining Game

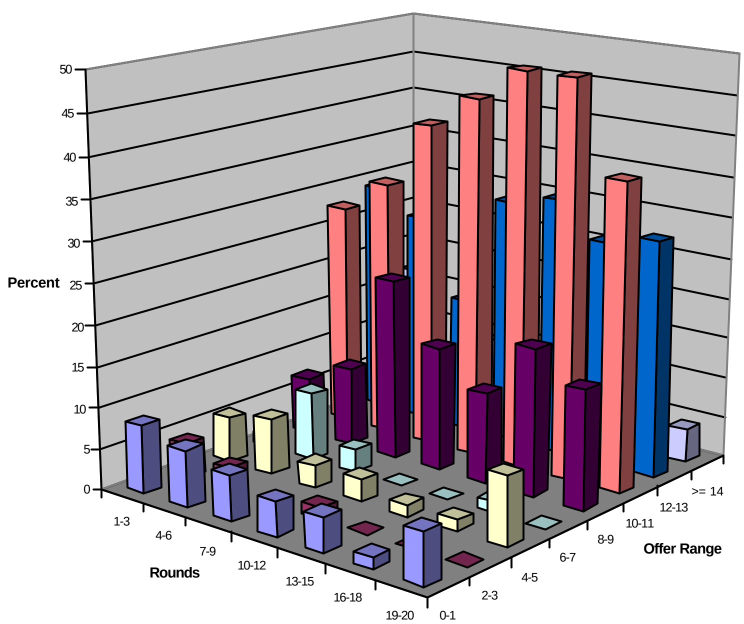

Figure 12: Distribution of Offers in the Ultimatum Game

6.1. Do Proposals Change Over Time?

Figures 11 and 12 show the distribution of offers for the convex and standard ultimatum games respectively. As with the responder data, there do not appear to be any clear and consistent trends in this data over time.

This was confirmed by estimating a model of learning on both the standard and convex games. Table 3 reports estimates of a Roth-Erev (1995) learning model.[20] This is a simple learning model with three parameters. Parameter s indicates the "initial strength" or the aggregate initial propensity to try strategies (a larger s indicates slower learning), e is a parameter meant to capture experimentation, and ρ captures "forgetting" or decay in learning. We find that the parameter estimates for the convex and standard game are nearly identical, and the difference across the models is not statistically significant.[21] This indicates that differences in learning across the games are not likely to have affected outcomes. This adds further confidence to our use of data from the convex game to predict the outcome in standard games.

Table 3
Parameter Estimates (s.e.'s) for a Roth-Erev Reinforcement Learning Model

	s	e	ρ	*Log-likelihood*
Standard Game	1.18	0.37	0.15	-1057.75
	(0.24)	(0.04)	(0.02)	
Convex Game	1.52	0.33	0.14	-1171.43
	(0.30)	(0.04)	(0.01)	
Erev and Roth (1998)	9	0.20	0.10	

How do we interpret the learning parameters? We also report in Table 3 the estimates obtained by Erev and Roth (1998). The parameter s is sensitive to the units of measurement of the payoffs, so are not directly comparable, while e and ρ are less influenced by scale. Notice that Erev and Roth's estimates of e and ρ are

[20] We chose the Roth-Erev model for its simplicity. The results from this model did not suggest any benefit from exploring more complicated models, such as the EWA model of Camerer and Ho (1999), or fictitious play models of Fudenberg and Levine (1998).

[21] A likelihood ratio test of the null hypothesis that the parameters of the convex game are equal to the parameters of the ultimatum game cannot be rejected, $\chi^2[3] = 3.70$, $p = 0.30$.

remarkably similar to the ones we estimated.[22] The similarity of the parameters indicate that the differences across the games are due to responders. As Cooper, Feltovich, Roth, and Zwick (2003) note, proposers are not in the driver's seat in the ultimatum game, and this is reflected in the learning models. This highlights the importance of understanding responders' behavior first.

6.2. Proposals and Predicted Earnings

Figure 13 illustrates both the distribution of offers and the predicted earnings, conditional upon those offers, for both convex and standard ultimatum games. The bars show the frequency of offers, and the lines show the predicted earnings conditional on the offer. To predict earnings we used the non-parametric estimate of preferences from section 5.2. We also calculated predicted earnings using the parametric model and got quite similar results.

First, we see that for the convex game the payoffs reach a maximum over a broad range of offers from 8 to 13, and in the standard game the maximum is reached for offers in the 10 to 13 range. By far the most popular offers for both games is either 10-11 or 12-13. Either of these choices is consistent with money maximization.

Another difference in offers can be seen by exploring lower offers. As noted earlier, there are clearly more offers made at the 2-7 range for the convex game than in the standard game. Eighteen percent of all offers are made in this range for the convex game, but only seven percent are for the standard game. What could explain this difference?

One hypothesis is simply that the expected earnings for these offers is much higher for the convex game–in some cases almost twice as high. Hence, subjects would not have been so foolish to experiment with these low offers in the convex game. A second hypotheses can be found in our earlier (2003) paper. If the expected return from a low offer were the same in the two games, the offer would still be less attractive in the standard game because it comes with more risk. Similarly in this data, we cannot rule out that greater risk aversion is also making subjects more wary of low offers in the standard ultimatum game.[23]

[22] While Erev and Roth estimated parameters using a grid search and a minimum distance criteria, we estimated parameters using maximum likelihood and taking the responders' behavior as given.

[23] One difference between the parametric prediction of earnings and the non-parametric prediction is that the two predictions are much closer in the parametric version. This would add greater support to the risk aversion explanation for this difference. However, as a general propo-

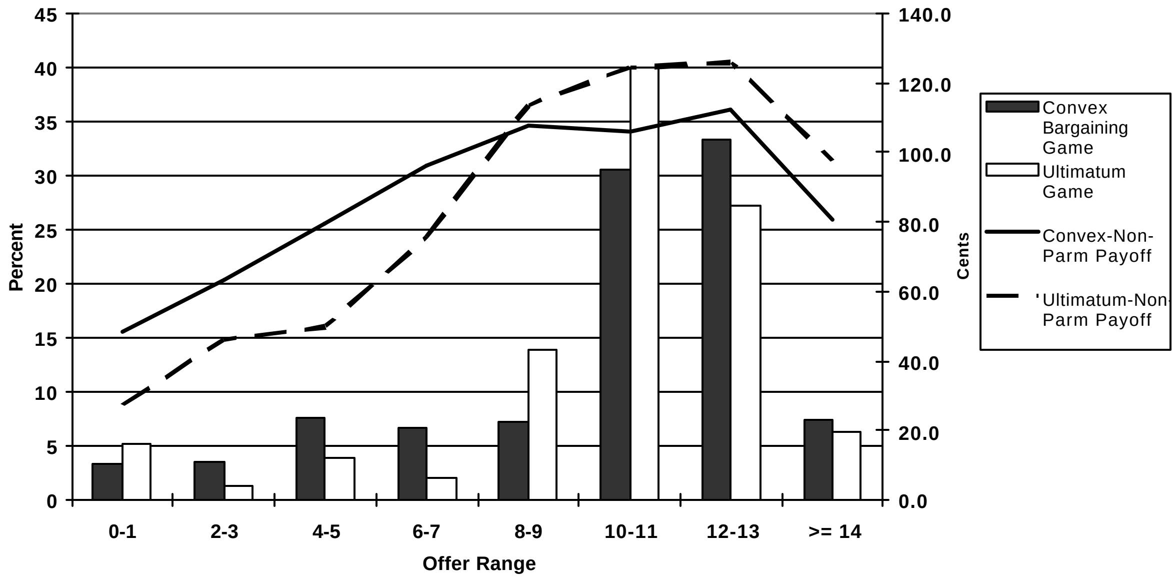

Figure 13: Frequency of Offers (bars, left axis), and Predicted Earnings (lines, right axis)

7. Conclusions

What do responders' preferences look like? This question, which is key to understanding both bargaining and fairness, is nearly impossible to answer using data from an ultimatum game. Here we employed the convex ultimatum game, where responders can "shrink" the proposal as well as accept or reject it. This game collects the information on responders' preferences that allows us to measure preferences and even estimate utility functions for responders. If these estimated preferences can predict the actual choices in the standard ultimatum game, then we have a compelling picture of preferences of responders.

We found that about 90 percent of subjects exhibited well-behaved preferences that were convex and regular over the space of own- and other-payoff. Of these, about 15 percent had monotonic preferences that accepted all offers, and 35 percent had linear preferences that rejected offers up to a point and accepted all above that. The other 50 percent of subjects had strictly convex but not monotonic preferences. These are preferences for which indifference curves bend

sition, nonparametric estimates should be favored, and so we presented them here instead.

back at the extremes. We confirm this with both nonparametric estimates of best reply functions and parametric estimation of utility functions.

When applying the estimated preferences from the convex game to predict the play in the standard ultimatum game we get a strikingly close fit—the preferences measured in the convex game overlap nearly perfectly with the actual choices made by subjects in the standard game. We take this out-of-sample prediction as strong support for our approach and for the set of estimated preferences.

What does this research tell us about ultimatum bargaining? Most importantly, it shows that rejections can be "rationalized" with preferences that are convex, nonmonotonic, and that have the added quality of regularity. Regularity, a concept akin to normality in demand theory, is remarkably successful in organizing responders' behavior. Our results also point to the importance of individual heterogeneity. While most subjects' choices can be captured by regular preferences, there was great variety in individual utility functions. Accounting for this heterogeneity is key to a successful account of bargaining behavior.

What do our results indicate for models of fairness? This data provides the most specific information on preferences to date, and sets a clear goal for models of fairness—modellers need no longer guess about the underlying shape of responders' preferences. Already it establishes that the simplest models of fairness, those that imply linear indifference curves, will capture only a small fraction of the data. Models with more subtle assumptions on preferences, such as those of Rabin (1993) and Li (2004), may be far more productive.

What do our analyses indicate for future research? An important aspect of our data not yet explored is the rich detail about the heterogeneity of preferences. These findings should perhaps put renewed emphasis on theories and experiments on bargaining under incomplete information. In addition, they suggest that there is potential in new research on trading institutions that perform well in the presence of heterogeneous populations of fair-minded bargainers.

Appendix 1: Calculating Expected Responses for the Convex Bargaining Game

Estimation: Following the approach proposed in section 5.3, we use a random utility model to represent responders' behavior. Denote by $\pi_s(x,y)$ responder's payoff when a proposer passes x and a responder chooses y, and by $\pi_o(x,y)$ the corresponding proposer's payoff. We assume that responder's utility can be represented by the following function:

$$u(x,y) = -(\pi_s(x,y)-\theta_s)^2 - \beta_1(\pi_o(x,y)-\theta_o)^2 - \beta_2(\pi_s(x,y)-\theta_s)(\pi_o(x,y)-\theta_o)$$

The probability that a responder chooses $y=k$ when a proposer passes x equals:

$$\Pr(y=k|x) = \Pr(u(x,k)+\epsilon_{xk} \geq u(x,y)+\epsilon_{x\widetilde{y}}, \text{all } \widetilde{y} \neq k)$$

Under the assumption that all ϵ_{xy} are independent and that $\Pr(\epsilon_{xy} \leq \epsilon) = \exp(-e^{-\epsilon})$, we obtain the multinomial logit representation,

$$\Pr(y=k|x) = \frac{e^{u(x,k)}}{\sum_j e^{u(x,j)}}$$

We estimated the parameter vector $(\beta_1, \beta_2, \theta_s, \theta_o)$ using maximum-likelihood methods. To insure the global concavity of the utility function, we impose the restriction that $4\beta_1 \geq \beta_2^2$.

Prediction: In order to predict responders' behavior across games, we make use of the logit's property that the odds ratio between options must remain constant as more options are added. In particular, suppose that responders are restricted to choose between complete acceptance ($y=10$) or complete rejection ($y=0$). We can then calculate the probability of acceptance by:

$$\Pr(y=10|x) = \frac{e^{[\widehat{u}(x,10)-\widehat{u}(x,0)]}}{1+e^{[\widehat{u}(x,10)-\widehat{u}(x,0)]}}$$

where $\widehat{u}(x,y)$ denotes the estimated utility associated to choices x and y.

Appendix 2: Parameter Estimates of Utility Functions for Convex Game

Table A1 presents the parametric estimates for all 24 utility functions estimated in the convex game, including maximum likelihood estimates for the 14 utility functions of the form described in Appendix 1. Standard errors are calculated using the cross-product of the gradient. The table also lists the essential preference parameter for those with monotonic, linear and near linear preferences. The critical value listed is that offer, in percentage terms, at which the responder will choose to fully accept offers and below which they will fully reject them. The critical value is derived from the non-parametric estimates.

Table A1

Parametric Estimates of Utility

Monotonic, Linear, and Near Linear Subjects:

Subj	Accept all offers of at least:
5	0%
12	0%
15	0%
1	4%
10	32%
22	24%
7	20%
24*	4%
6	20%
21	12%

Near Strictly Convex and Not Regular Subjects:

Subj	β_1	s.e.	β_2	s.e.	θ_o	se	θ_s	s.e	$\overline{L}$
13	0.41	0.04	-1.29	0.06	5.09	1.09	-6.71	0.70	-1.54
2	0.13	0.05	-0.73	0.13	2.44	0.61	-1.85	0.22	-2.03
17	0.15	0.05	-0.78	0.12	0.77	0.59	-5.95	0.23	-2.08
16	0.45	0.12	-1.34	0.18	3.20	1.27	-4.73	0.85	-1.81
19	0.66	0.04	-1.63	0.05	5.78	1.77	-7.01	1.45	-1.76
23	0.18	0.03	-0.84	0.07	3.95	1.47	-5.01	0.62	-1.82
9	0.60	0.24	-1.55	0.30	0.04	0.65	-2.47	0.50	-2.12
3	0.18	0.05	-0.85	0.13	2.37	0.48	-0.32	0.20	-2.19
14	0.13	0.05	-0.71	0.13	0.44	0.79	-6.48	0.28	-2.08
20	0.35	0.07	-1.18	0.11	0.64	0.65	-5.44	0.38	-2.00
11	0.41	0.09	-1.28	0.15	1.53	0.79	-4.45	0.50	-2.01
4	0.28	6.39	-1.09	11.71	2.28	740.99	-2.00	1315.88	-2.02
8	0.54	0.14	-1.48	0.20	0.52	0.55	-4.35	0.40	-1.95
18	0.22	0.04	-0.94	0.08	0.38	1.07	-8.43	0.50	-1.91

* Subject 24 is technically classified as Strictly Convex, but only shrank one offer, hence a parametric function could not be estimated. When offered 1 cent, his lowest offer, he chose 9 of 10 quarters to divide. We then approximate his utility with a linear indiffernce curve overlapping the budget at offers of 1.

References

Ahlert, Marlies, Arweb Cruger, and Werner Guth. "How Paulus Becomes Saulus - An Experimental Study of Equal Punishment Games." Working Paper, Humboldt University of Berlin, 2001.

Andreoni, James, Paul M. Brown and Lise Vesterlund, "What Makes an Allocation Fair? Some Experimental Evidence." *Games and Economic Behavior*, 40, July 2002, 1-24.

Andreoni, James, Marco Castillo, and Ragan Petrie. "What Do Bargainers' Preferences Look Like? Experiments with a Convex Ultimatum Game." *American Economic Review,* June 2003, 93(3), pp.672–685.

Andreoni, James and John H. Miller. "Giving According to GARP: An Experimental Test of the Consistency of Preference for Altruism." *Econometrica,* March 2002, 70(2), pp. 737–753.

Bolton, Gary E. and Axel Ockenfels. "ERC: A Theory of Equity, Reciprocity and Competition." *American Economic Review*, March 2000, 90(1), pp. 166–193.

Camerer, Colin. "Social Preferences in Dictator, Ultimatum, and Trust Games." Chapter 2 in *Behavioral Game Theory*, Princeton University Press: Princeton, NJ, 2003.

Camerer, Colin and Teck Ho. "Experience-weighted Attraction Learning in Normal-from Games." *Econometrica*, 1999, 67,827–74.

Charness, Gary and Matthew Rabin, "Understanding Social Preferences with Simple Tests." *Quarterly Journal of Economics,* August 2002, 117(3), pp. 817-869.

Cooper, David, Nick Feltovich, Alvin Roth, and Rami Zwick. "Relative versus Absolute Speed of Adjustment in Strategic Environments: Responder Behavior in Ultimatum Games," *Experimental Economics*, October 2003, 6(2), 181-207.

Eckel, Catherine C. and Philip Grossman. "Chivalry and Solidarity in Ultimatum Games." *Economic Inquiry,* April 2001, 39(2), pp. 171–88.

Erev, Ido and Alvin E. Roth, "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria" *American Economic Review*, September 1998, 88(4), pp. 848-81.

Fehr, Ernst and Klaus M. Schmidt. "A Theory of Fairness, Competition and Cooperation." *Quarterly Journal of Economics*. August 1999, 114(3), pp. 817–868.

Forsythe, Robert, Joel Horowitz, N.S. Savin and Martin Sefton. "Fairness in Simple Bargaining Games." *Games and Economic Behavior*, May 1994, 6(3), pp. 347–69.

Fudenberg, Drew and David K. Levine, *The Theory of Learning in Games*, Cambridge: MIT Press, 1998.

Gneezy, Uri, Ernan Haruvy, and Alvin E. Roth, "Bargaining under a Deadline: Evidence from the Reverse Ultimatum Game." *Games and Economic Behavior*, 2003, 45(2), pp. 347-68 .

Guth, Werner, R. Schmittberger, and B. Schwartz. "An Experimental Analysis of Ultimatum Bargaining" *Journal of Games and Economic Behavior*, December 1982, 3(4), pp. 367–88.

Harbaugh, William, Kate Krause and Steve Liday. "Children's Bargaining Behavior," Working paper, University of Oregon, 2002.

Li, Jing, "A Model of Ethical Behavior." working paper, University of Wisconsin, 2004.

Levine, David K. "Modeling Altruism and Spitefulness in Experiments," *Review of Economic Dynamics*, July 1998, 1, pp. 593-622.

Kennan, John, and Robert Wilson. "Bargaining with Private Information" *Journal of Economic Literature*, Mar 1993, 31(1), pp. 45-105.

Manski, Charles. "Identification of Decision Rules in Experiments on Simple Games of Proposal and Response," *European Economic Review*, 2002, 46(4-5), pp. 880-891.

Ochs, Jack and Alvin E. Roth. "An Experimental Study of Sequential Bargaining." *American Economic Review*, June 1989, 79(3), pp. 355–84.

Rabin, Matthew. "Bargaining Structure, Fairness, and Efficiency." Working Paper, University of California, 1997.

Rabin, Matthew. "Incorporating Fairness into Game Theory and Economics." *American Economic Review*, December 1993, 83(5), pp. 1281–1302.

Roth, Alvin E. "Bargaining Experiments," in J.H. Kagel and A.E. Roth, eds, *The Handbook of Experimental Economics*, Princeton, NJ: Princeton University Press, 1995.

Roth, Alvin E., and Ido Erev "Learning in Extensive-form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term." *Games and Economic Behavior,* 8, 1995, 164–212.

Roth, Alvin E., V. Prasnikar, M. Okunofujiwara and S. Zamir. "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study." *American Economic Review*, December 1991, 81(5), pp. 1068–95.

Slonim, Robert and Alvin E. Roth, "Learning in High Stakes Ultimatum Games: An Experiment in the Slovak Republic." *Econometrica*, May 1998, 66(3), pp. 569–96.

Suleiman, Ramzi. "Expectations and Fairness in a Modified Ultimatum Game." *Journal of Economic Psychology*, November 1996, 17(5), pp. 531–554.