

What Do Bargainers' Preferences Look Like? Experiments with a Convex Ultimatum Game

By JAMES ANDREONI, MARCO CASTILLO, AND RAGAN PETRIE*

The ultimatum game, by its all-or-nothing nature, makes it difficult to discern what kind of preferences may be generating choices. We explore a game that convexifies the decisions, allowing us a better look at the indifference curves of bargainers while maintaining the subgame-perfect equilibrium. We conclude that bargainers' preferences are convex and regular but not always monotonic. Money-maximization is the sole concern for about half of the subjects, while the other half reveal a preference for fairness. We also found, unexpectedly, the importance of risk aversion among money-maximizing proposers, which in turn generates significant bargaining power for fair-minded responders. (JEL C78, D64, C92)

The ultimatum game experiment was designed to test theories of alternating-offers bargaining. The consistent failure of ultimatum bargaining to reach the subgame-perfect prediction has raised many fundamental questions about bargaining models and about individual preferences.¹ What, for instance, is the role of fairness in bargaining, and how is this expressed in utility functions of bargainers?

Although the ultimatum game is a two-person interaction, the most puzzling behavior is that of responders. Why do so many responders reveal a preference for allocations that give both players less consumption? A common answer is that they must care about the relative allocations of the players. This concern has been described

variously as fairness, spite, or reciprocity. The challenge to economists is to understand these choices in terms of optimizing individuals with transitive binary preference orderings. Discussing "motivations for fairness" or "reciprocity" in an economic framework reduces to a description of a choice set, a hypotheses on the elements of preferences, and tests of whether actions conform to the model of choice. We take this approach here by asking, What do bargainers' preferences look like?

If we have defined the choice set carefully, we would hope to see preferences that are continuous and convex, and most likely strictly convex. Unfortunately, given the existing data on ultimatum games, it is impossible to tell whether bargainers' preferences meet this expectation. The reason is that the choice set of responders is itself not convex. By allowing subjects to only accept or reject an offer, we are losing precious information that would allow us to learn more about responders' preferences.

In this experiment we present subjects with a convex ultimatum game. That is, after a proposer makes an offer, say 80 percent of the pie for the proposer and 20 percent for the responder, the responder has the option of shrinking the size of the pie. Shrinking the pie to zero is the same as rejection, and not shrinking it at all is the same as acceptance. In the convex game, however, the pie can be shrunk to an intermediate level. Note that the convex game and the standard game have the

* Andreoni: Department of Economics, University of Wisconsin, 1180 Observatory Drive, Madison, WI 53706 (e-mail: andreoni@wisc.edu); Castillo: Environmental Policy Program, Georgia State University, Atlanta, GA 30303 (e-mail: prcmecx@langate.gsu.edu); Petrie: Department of Economics, Georgia State University, Atlanta, GA 30303 (e-mail: rpetrie@gsu.edu). Andreoni thanks the National Science Foundation for financial support. We are grateful to Bill Harbaugh, Lise Vesterlund, and two anonymous referees for helpful comments.

¹ For important early contributions, see Werner Guth et al. (1982), Jack Ochs and Alvin E. Roth (1989), Roth et al. (1991), and Robert Forsythe et al. (1994). For a review of the literature on bargaining and ultimatum games, see Roth (1995) and Colin F. Camerer (2003). For important recent contributions, see Robert Slonim and 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.

same subgame-perfect Nash equilibria: responders should never shrink an offer, so the most selfish offer is made. When we fail to reach subgame perfection due to responses that only partly shrink the pie, we can learn much more about the shape of the indifference curves of responders.²

Our experiment also has the feature of asking subjects to make decisions in both roles of proposer and responder before their true role is determined. With this information we can look at the strategic choices of proposers and connect expressions of fairness in both roles. Merging our findings on proposers and responders, we can begin to assemble building blocks for a successful theory of behavior in ultimatum games.

We find that, at the aggregate level, the expected response to any proposal is largely the same across both games. However, in the convex game, there are significant numbers of subjects who shrink the pie to an intermediate level. Looking at the individual data, these partial rejections are consistent with preferences for own and other payoffs that are continuous, nonlinear, strictly convex, and nonmonotonic. Interestingly, subjects demonstrated a great deal of consistency both across games and across roles. In both games about half of the subjects indicate perfectly selfish preferences, and half care about fairness. Across roles, those who show a willingness to reject unfair offers as responders are also more generous proposers. Finally, our data reveal a surprising finding that risk aversion among responders in the ultimatum game confers a great deal of bargaining power to proposers. The convex game, which is less risky for proposers, robs bargaining power from fair-minded responders and shifts it to selfish proposers.

Section I formally presents the convex ultimatum game. Section II describes the experiment and Sections III through VI provide our

results. Section VII is a discussion and conclusion.

I. The Convex Ultimatum Game

Consider the following bargaining game with two players. The proposer first specifies the proportion of the money that will go to the responder. The responder then determines how much money to divide, from zero to M dollars. Formally, let $a \in [0, 1]$ be the proportion of money that the proposer allocates to the responder and let $m \in [0, M]$ indicate the amount of money that the responder chooses to divide. Then the payoff functions for the proposer, π_p , and the responder, π_r , are

$$\pi_p = (1 - a) \times m,$$

$$\pi_r = a \times m.$$

Note that by restricting $m \in \{0, M\}$, the game reverts to the standard ultimatum game. That is, the responder can only choose to divide the whole pie or none of it. In the convex ultimatum game, by contrast, the responder has the option to “reject” only part of the offer.³

In addition to being more general than the standard game, there are reasons to be interested in the convex game in its own right. For instance, bargaining in the real world is seldom so stark as the ultimatum game. Responders typically have more responses than to accept or reject. They can, for instance, take longer to reply to a proposal, workers can slow down their work effort, police officers can get the “blue flu,” airline pilots can refuse to work overtime, and teaching assistants can “lose” their grade books. All of these tactics are ways of shrinking the pie when someone is unhappy with an offer.

The convex ultimatum game is illustrated in Figure 1. The downward-sloping line represents all possible divisions of M . The proposer chooses a point along this line. Draw a line segment from the chosen point to the origin. The responder in the convex game then chooses a point along this line segment. In the standard

² There have been other convex games introduced in the literature. Ramzi Suleiman (1996) presents one similar 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 amount 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.

³ This game was also presented by Matthew Rabin (1997) as the “Squishy Game.”

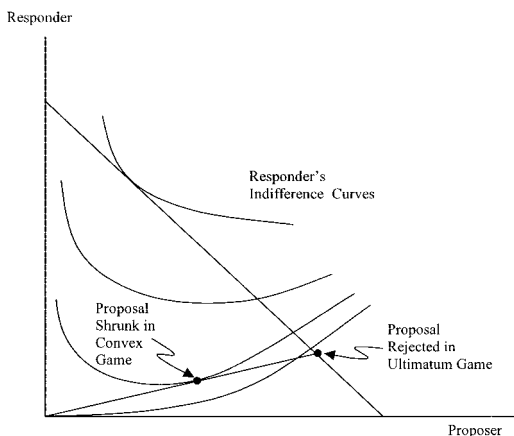


FIGURE 1. STANDARD AND CONVEX ULTIMATUM GAMES

ultimatum game, the responder can only choose between the endpoints of this line segment, that is, complete rejection or complete acceptance.

What predictions do we have for the convex game versus the standard game? First, consider the assumption that both players are selfish, risk-neutral money-maximizers. Then the unique subgame-perfect Nash equilibrium in the convex game is identical to that in the standard ultimatum game—since a responder will never shrink any offer, and the proposer knows this, then the proposer will make the most selfish offer possible. Note that the same subgame-perfect outcome would hold under slightly weaker conditions in which the proposer is a money-maximizer, and the responder has preferences, say $u(\pi_r, \pi_p)$, that are monotonic. Again, the responder would never shrink the pie and, knowing this, the proposer makes a selfish offer.

Second, consider responder preferences that are nonmonotonic in both players' payoffs, but in which the indifference curves are restricted to be linear, such as those suggested by Ernst Fehr and Klaus M. Schmidt (1999). Such indifference curves would be upward-sloping lines that are parallel, or which fan out. Then we predict the standard and convex ultimatum games would yield the same results. The linearity of preferences implies that any pie should be shrunk all the way to zero or not at all, making the convexity of the choice set uninteresting. This also holds for a model such as that suggested by Gary E Bolton and Axel Ockenfels

(2000). If a responder finds a particular relative payoff acceptable, she would always choose the point on the line where her payoff was the highest. The pie would never be shrunk to an intermediate value.

Third, suppose we maintain the assumption of nonmonotonic preferences for responders but allow indifference curves to be nonlinear, such as those shown in Figure 1. This figure shows preferences in which more for the proposer, but no less for the responder, could make the responder worse off.⁴ We now can expect a difference between the standard and convex ultimatum games. As we can see in Figure 1, a proposal that would be completely rejected in the standard game might be shrunk in the convex game.⁵ Likewise, other proposals that may be completely accepted in the standard game could also be shrunk in the convex game. Such preferences could, for instance, be consistent with Rabin's (1993) approach to fairness.

If responders behave as depicted in Figure 1, how should proposers behave? Consider proposers who are money-maximizers and are aware of the preferences of responders. In this case, the proposer will offer the sharing rule that maximizes own monetary payoffs given the responder's reaction curve. This is shown in Figure 2. As can be seen, the convex game strengthens the bargaining power of responders. That is, the ability to partially reject an offer will entice proposers to make more equal offers.

This analysis, of course, assumes that the proposer knows the preferences of the responder. If, instead, there are heterogeneous tastes of responders, this will lead to uncertainty about the response to a particular offer. Since the standard game only allows the extreme responses, then there are cases in which a given offer will be more risky in the standard game

⁴ The opposite case could also occur. That is, more for the responder but no less for the proposer could make the responder worse off. This would be the case if the responder disliked inequality that favored him more than he liked the extra income. Our experiment will allow us to identify subjects with this type of preference as well.

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.

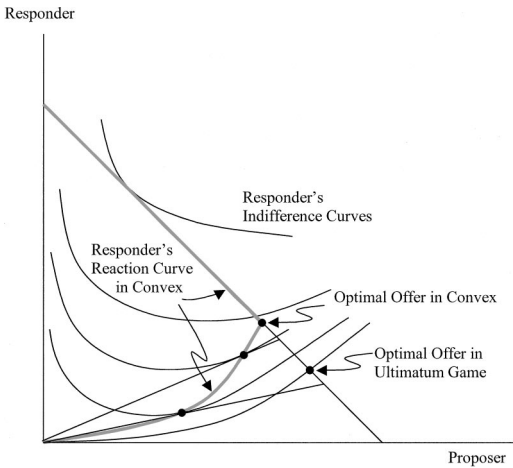


FIGURE 2. PREDICTED CHOICES

than in the convex game. Hence, by reducing the risk faced by proposers, the convex game may actually make selfish offers more appealing to proposers, and thus rob responders of some bargaining power. This effect, as we will see, may be quite important.

Finally, note that for proposers who care about altruism or equity, the structure of the game may not bind their actions. That is, the responder's ability to punish partially would not affect a proposer who, because of fairness, prefers to make a 50-50 offer, assuming of course it will be fully accepted.

II. The Experiment

We explore both the standard and convex ultimatum games. The instructions and decisions sheets in both games refer to the proposer as the Divider, since he divides each dollar, and the responder as the Designator since she designates the number of dollars to divide. However, we will retain the names of proposer and responder to describe the experiments here. Complete instructions for the convex game are available from the authors.⁶

The objective of the game is to divide 12 dollars. The proposer indicates how each dollar is to be divided. The responder then chooses how many dollars to divide. In the standard

ultimatum game, the responder has a choice of 0 or 12 dollars. In the convex game the responder can choose any number from 0 to 12.

Figure 3 shows the decision sheets used in the convex game. The proposer (the top panel) circles a letter from **a** to **k** indicating the division of each dollar. For each possible division chosen by proposers, the responder indicates how many dollars to divide (the bottom panel). The standard game is presented identically to the convex game, except in the decision sheet of responders the 13 numbers in column **B** are replaced with "0 or 12." In each session the subjects were asked to make decisions as both the proposer and the responder. They were told that their ultimate role would be determined randomly after all decisions were made. This method allows us to examine how a single person's attitude toward fairness is expressed in both roles. Notice that we also map the preferences of responders for each possible proposal. This gives us a richer characterization of the responders' preferences.⁷

To conduct the experiment, we recruited 116 students from intermediate-level economic and business classes at the University of Wisconsin-Madison. There were two sessions. In each session, 58 participants were assembled in one large room and then randomly assigned to one of two adjacent rooms, one for the standard game and one for the convex game. In each session, therefore, each treatment had 29 participants. One participant in each session was chosen at random to be a monitor. The monitor made no decisions but verified to the other participants that the correct procedures were followed. Hence, we collected data on 28 subjects per treatment per session, for a total of 112 subjects.

Once the participants were assembled and the monitor was chosen, the instructions were then read out loud to the participants, who read along with their own copy. A quiz was administered to make sure the participants understood the experiment, and answers were explained on the

⁷ This method has been used many times in ultimatum games. Jordi Brandts and Gary Charness (2000) explicitly tested for differences in this "strategy elicitation method" and the "sequential elicitation method" in the Prisoner's Dilemma and the Chicken Game. They found no significant difference in behavior between treatments.

⁶ Go to <http://www.ssc.wisc.edu/~andreoni/>.

A. Divider chooses a rule by circling one letter in this column	B. Possible Dividing Rules <i>Of each Dollar to divide, the rule is:</i>
a	Divider gets 99¢ and Designator gets 1¢
b	Divider gets 90¢ and Designator gets 10¢
c	Divider gets 80¢ and Designator gets 20¢
d	Divider gets 70¢ and Designator gets 30¢
e	Divider gets 60¢ and Designator gets 40¢
f	Divider gets 50¢ and Designator gets 50¢
g	Divider gets 40¢ and Designator gets 60¢
h	Divider gets 30¢ and Designator gets 70¢
i	Divider gets 20¢ and Designator gets 80¢
j	Divider gets 10¢ and Designator gets 90¢
k	Divider gets 1¢ and Designator gets 99¢

	A. If the Divider chooses this Dividing Rule....	B. ...then I choose to divide this many dollars (circle one for each Dividing Rule):
a	Divider gets 99¢ and Designator gets 1¢	0 1 2 3 4 5 6 7 8 9 10 11 12
b	Divider gets 90¢ and Designator gets 10¢	0 1 2 3 4 5 6 7 8 9 10 11 12
c	Divider gets 80¢ and Designator gets 20¢	0 1 2 3 4 5 6 7 8 9 10 11 12
d	Divider gets 70¢ and Designator gets 30¢	0 1 2 3 4 5 6 7 8 9 10 11 12
e	Divider gets 60¢ and Designator gets 40¢	0 1 2 3 4 5 6 7 8 9 10 11 12
f	Divider gets 50¢ and Designator gets 50¢	0 1 2 3 4 5 6 7 8 9 10 11 12
g	Divider gets 40¢ and Designator gets 60¢	0 1 2 3 4 5 6 7 8 9 10 11 12
h	Divider gets 30¢ and Designator gets 70¢	0 1 2 3 4 5 6 7 8 9 10 11 12
i	Divider gets 20¢ and Designator gets 80¢	0 1 2 3 4 5 6 7 8 9 10 11 12
j	Divider gets 10¢ and Designator gets 90¢	0 1 2 3 4 5 6 7 8 9 10 11 12
k	Divider gets 1¢ and Designator gets 99¢	0 1 2 3 4 5 6 7 8 9 10 11 12

FIGURE 3. DECISION FORMS FOR PROPOSER (TOP) AND RESPONDER (BOTTOM) FOR THE CONVEX ULTIMATUM GAME

chalkboard. Subjects then filled out the experimental decision forms.

Participants placed completed decision forms in plain envelopes. These were collected, shuffled in front of the participants and randomly separated into two even piles, one for proposers

and one for responders. The forms were then taken to a nearby room to calculate payments. These payments, along with a \$6 show-up fee, were placed in a private envelope with only the participant's identification number on the outside. Another experimenter, not involved in the

TABLE 1—AVERAGE NUMBER OF DOLLARS TO DIVIDE CHOSEN BY RESPONDER FOR EACH DIVIDING RULE

	Dividing Rule: Cents per dollar that proposer allocates to responder										
	1	10	20	30	40	50	60	70	80	90	99
Standard	6.86	8.14	9.43	10.93	11.57	11.79	11.79	11.79	12.00	11.79	11.36
Convex	6.95	7.95	8.84	9.75	10.45	11.57	11.11	11.16	11.09	10.82	10.91
Difference	-0.09	0.19	0.59	1.18	1.12	0.22	0.68	0.63	0.91	0.97	0.45

calculation of payments, handed out the envelopes to the participants, who were then escorted from the room. While payments were being calculated, all participants filled out a postexperiment questionnaire intended to check their understanding of the game. The experiment lasted less than an hour, and participants earned an average of \$11.58 (standard deviation 3.44). The following four sections describe our results.

III. Instrument Check

First, did our procedures generate results for the standard ultimatum game that are similar to those found by others? On the proposer side, more than half of all offers were equal splits, and only 12 percent were at the subgame-perfect equilibrium of (99,1). The average offer was 38 percent. The overall rejection rate in our experiment was 11 percent, while 43 percent of subgame-perfect offers were rejected. While our proposers are somewhat more selfish on average, our results are well within the range of those from previous experiments (see Roth, 1995, and Camerer, 2003, for reviews).

Next, were our methods understood by the subjects? We attempted to measure this with a postexperiment questionnaire. We asked two questions intended to test whether subjects could correctly calculate payoffs.⁸ We found a very high degree of understanding. Of the 112 subjects, three subjects in the convex game did not understand the payoffs, and all subjects in

⁸ For instance, one of the questions was this: Suppose there was a pairing in which the Divider chose Dividing Rule c: "Divider gets 80¢ and Designator gets 20¢." a) If the Designator has selected \$7 on line c, what would be the payoff for the Divider and the Designator? b) How many dollars should the Designator choose to divide to make the most money possible for herself?

TABLE 2—CLASSIFICATION OF RESPONDER BEHAVIOR

Category	Convex UG	Standard UG
Monotonic	26 (46 percent)	31 (55 percent)
Linear	7 (13 percent)	22 (39 percent)
Strictly Convex	23 (41 percent)	3 (5 percent)
Total	56	56

the standard game understood the payoffs. Nonetheless, two subjects in the standard game described their strategy as choosing randomly, and indeed their responses bounced asystematically across offers. All of the analysis to follow includes these five noisy subjects, although similar results hold if they are excluded.

IV. Responder Preferences

Table 1 shows that the average response, conditioned by proposal, is consistently less than \$12, except for offers of 80 cents in the standard game. In general, responses appear to suggest nonmonotonic preferences. Responders are slightly more accepting in the standard game relative to the convex game. However, this tends to be significant primarily for generous offers, those greater than 50 cents.

Closer examination of the convex game reveals significant amounts of intermediate values chosen. Overall, 26 percent of all choices were something other than 0 or 12 dollars. This ranged from a peak of 36 percent for offers of 20 to 40 cents per dollar, to a minimum of 14 percent for offers of 99 per dollar. This suggests that there may be a significant number of subjects with strictly convex but nonmonotonic preferences. To explore this, we look to the "reaction curves" of responders. In Table 2 we place responders into three categories based on patterns in their reaction curves. Responders in the Monotonic category choose to accept all offers. This

TABLE 3—CLASSIFICATION OF NONMONOTONIC AND CONVEX RESPONDERS BY MINIMUM DISTANCE OF CHOICE TO LINEAR AND REGULAR PREFERENCES, AVERAGED PER CHOICE

	Convex ultimatum game, 23 subjects (41 percent of total)				Standard UG, 3 subjects (5 percent of total)	
	Regular one-sided	Regular two-sided	Weakly two-sided	Not regular preferences	Regular two-sided	Not regular preferences
Number of subjects	9	4	7	3	1	2
Percent of total	16	7	13	5	2	3
Minimum distance to linear:						
Straight average	\$1.34	\$2.36	\$2.94	\$2.73	\$2.18	\$3.82
Choice-weighted average	\$1.72	\$2.32	\$1.89	\$3.20	\$0.00	\$4.07
Minimum distance to regular:						
Straight average	0	0	\$0.34	\$1.39	0	\$2.73
Choice-weighted average	0	0	\$0.15	\$1.49	0	\$4.07

group has monotonic preferences which could be selfish, linear, or strictly convex. Responders in the Linear category completely accept all offers up to a point and then completely reject any offer less than that. Their preferences are consistent with linear indifference curves and nonmonotonic preferences, and with the models of Fehr and Schmidt (1999) and Bolton and Ockenfels (2000). Responders in the Strictly Convex category partially reject offers. They appear to have nonmonotonic preferences and nonlinear indifference curves. Note that if preferences are convex and strictly increasing in one's own payoff, players in the standard games would never reject an offer if they had accepted a lower offer. Of the three subjects that violate this, one of the responders showed an aversion to unequal offers even if they were favorable to the responder, while two reported to choose their strategy randomly, as discussed earlier.

Examining Table 2, the first thing to note is that, within each game, there is a diversity in preferences across subjects. There appears to be several types of responders with various combinations of monotonicity and convexity in preferences. Looking across games, we see that the number of subjects classified as Monotonic responders make up about half of the sample in both conditions. This is remarkable for two reasons. First, it suggests that the degree of purely selfish responses is independent of the condition. Second, it shows that about half of the subjects have preferences that display concerns for something besides their own payoffs—a re-

sult similar to findings elsewhere.⁹ Finally, Table 3 also indicates that the standard game masks the strict convexity of indifference curves for most of the nonselfish responders. Of the subjects that are not money-maximizers in the convex game, at least 75 percent of them show some sort of nonlinearity in their indifference curves. This could not be detected in the standard ultimatum game.

Next we ask what more we can say about the subjects in the Strictly Convex category in Table 2. Do these choices seem systematic, as if they came from a well-behaved preference ordering, or are there errors? To examine this, it is helpful to consider some reasonable conditions on preferences in these games. Were this an ordinary consumer choice problem, we could, for instance, explore certain regularity conditions, such as an assumption of normal goods. Is there a corollary to normal goods in this choice space? Note that responders with convex preferences will, by definition, have a "favorite" dividing rule among the 11 rules offered. That is, there is one rule that they most prefer proposers to choose. Call this dividing rule R^* . Then it would seem reasonable to expect that as other dividing rules get closer to R^* the responder should not *decrease* the amount of dollars to divide. Stated differently, as proposers' offers get closer to the responder's favorite offer, the responder will not shrink the pie more.

⁹ See Andreoni (1995) and Andreoni and John H. Miller (2002), for instance.

We will refer to preferences with this quality as *regular preferences*. Such a quality of preferences is consistent with the fairness model of Rabin (1993), for example, which, in our language, rules out nonregular preferences.¹⁰

Plotting a reaction curve for responders with regular preferences means that the reaction curve will have a “single peak.” That is, as the upward-sloping budgets for responders sweep counterclockwise, the choice along the budget will move up monotonically and, perhaps, down monotonically, but it will never go down and back up.¹¹ Choices that violate this will not be regular preferences (although we cannot rule out that they come from a consistent, albeit peculiar, preference ordering).

Applying the criteria of regular preferences revealed several interesting patterns in the choices of the 23 Strictly Convex subjects. Figure 4 illustrates the patterns we observed. First, nine subjects had regular preferences that could be characterized as having one-sided inequality aversion. An example is shown in panel a. These subjects shrunk offers that were unfavorable, up to a point, and then fully accepted all offers after that. Four subjects had preferences showing two-sided inequality aversion like that shown in panel b. These subjects shrunk both the most selfish and the most selfless offers.

This leaves 10 of the 23 subjects who did not strictly meet the definition of regular preferences. So, next we ask whether errors by subjects are important. Subjects could, for instance, be close to linear, monotonic, or regular preferences. To check this, we calculated the minimum distance, measured in dollars, that subjects’ choices would need to be moved in order to meet the definition of either linear (including monotonic) or regular convex preferences. We averaged this total deviation in two ways. First is a straight average across the 11 choices. Second, in recognition that subjects

may know that some dividing rules were less likely to be chosen than others (and hence did not take the choice as seriously), we weight the deviations by the probability that the dividing rule was actually chosen. Thus, one can think of these averaged deviations as the expected loss from failing to optimize according to linear or regular preferences.

Table 3 shows the results for the minimum distance to both linear and regular preferences for all the Strictly Convex subjects. First, this reveals that those with regular preferences discussed above are, on average, quite far from having linear preferences. They can expect to lose \$1.72 to \$2.32 on average from their deviations. Second, of the ten who did not strictly meet the definition of regular, seven are extremely close. These seven can expect to lose only \$0.15 from their deviation from regular, but \$1.89 for their deviation from linear. We can call these seven weakly regular. They all, it turns out, showed a desire for two-sided inequality aversion. An example of a subject with weakly regular preferences is shown in panel c of Figure 4. This subject’s choice-weighted deviation from linear is \$1.95 and from regular is \$0.09. Finally, there were three subjects in the convex game whose preferences were clearly not regular. An example of such a subject is given in panel d of Figure 4.

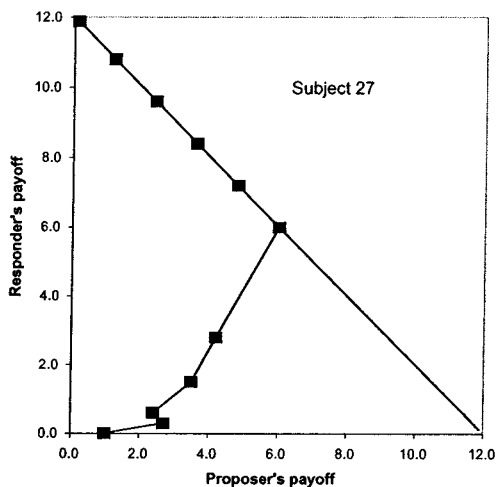
We can also examine whether preferences are regular in the standard ultimatum game. The final two columns in Table 3 show that of the three Strictly Convex players identified in Table 2, one has preferences that are regular and, when weighted by choice probabilities, is consistent with linear preferences. Two other subjects, however, made choices that are not from regular preferences.

Overall, the evidence in this section presents a picture of consistent choices among responders with strictly convex, but not always monotonic, preferences.¹²

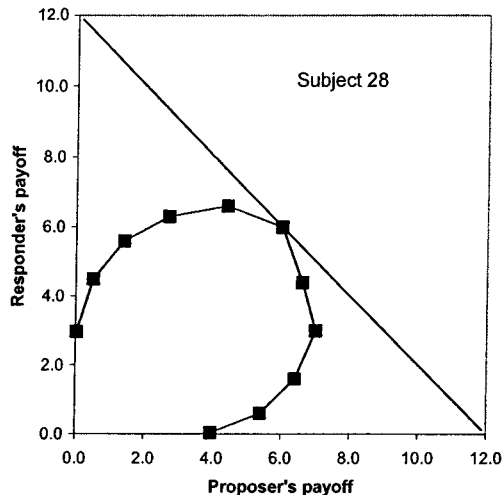
¹⁰ In fact, a strict application of Rabin’s (1993) model would predict that as offers get more generous, responders’ choices would rise monotonically to dividing \$12 by at least the 50-50 offer, and begin shrinking the pie again (if ever) after the 50-50 offer. This characterizes all of our one- and two-sided inequality-averse subjects.

¹¹ We, of course, mean a weak version of monotonicity. A responder’s choice may “plateau” at a given level of shrinking the pie, either on or off their “peak,” and still be consistent with our notion of regular.

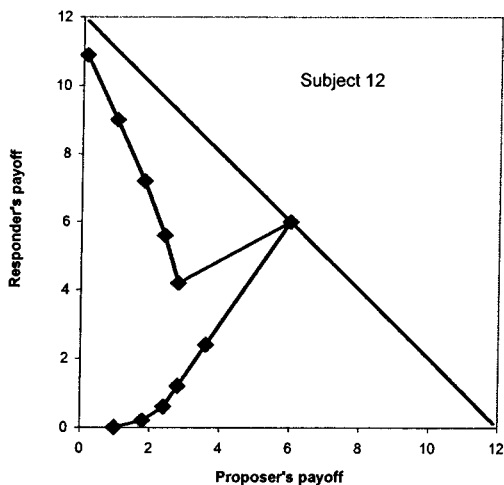
¹² Note the similarities to the findings of Andreoni and Miller (2002) in nonstrategic settings. Andreoni and Miller also found about half of subjects were selfish, while half cared for altruism, that 6.2 percent of subjects had linear preferences, compared to 13 percent above, and that 4 percent had “irrational” preferences, compared to 5 percent above. However, the move to a strategic setting generated some interesting differences. On upward-sloping budgets




a. One-Sided Inequality Averse



b. Two-Sided Inequality Averse



c. Weakly Regular Two-Sided



d. Not Regular Preferences

FIGURE 4. EXAMPLES OF SUBJECT'S CHOICES

V. Proposer Preferences

Figure 5 shows the distribution of proposals across the two games. This indicates that pro-

posers make far more 50-50 offers in the standard game, and they make far more selfish offers of (99,1) in the convex game. Indeed, the standard game results in significantly more fair offers.¹³ How does this result match up with the

they found that 23.5 percent of subjects shrank the pie on disadvantageous inequality, while 8.8 percent shrank the pie with advantageous inequality, compared to 41 percent and 20 percent, respectively, above.

¹³ A simple test comparing means indicates a significantly higher mean proposal for the standard game, with $t = 2.4$.




FIGURE 5. DISTRIBUTION OF OFFERS

various predictions? First, this finding is clearly at odds with both the subgame-perfect Nash equilibrium prediction, and with the fairness models of Fehr and Schmidt (1999) and Bolton and Ockenfels (2000) which predict no difference in proposals across the games.

How does the result compare to models which assume strictly convex but nonmonotonic indifference curves among responders? Under the assumption of perfect information, this model predicts that the convex game should give more bargaining power to responders, and hence result in more fair offers. This is the opposite of what we found. What could explain this?

We see two plausible explanations. First, it is possible that, in fact, these proposals are an optimal reply to the distribution of responses. Hence, we can compare the money-maximizing response of proposers across the two games. Second, since the experimental setting is more appropriately described as proposers having imperfect information about responder preferences, we can investigate whether the differences in risk across the two games had an effect. We answer these questions in turn.

A. Are Proposers Choosing a Best Reply?

Figure 6 shows the empirical version of Figure 2, that is, the average of responders’ reactions in each game. This maps the expected payoffs for the proposer and the average choice of responders for each sharing rule.¹⁴ If the

¹⁴ Note that it is also possible to convert Figure 6 into an “average demand curve” for reductions in payoff. The slope

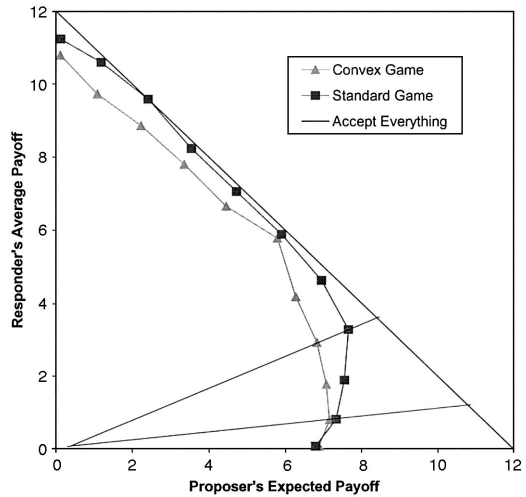


FIGURE 6. AVERAGE RESPONDER REACTIONS

proposers correctly infer this distribution of responses, what should they choose?

From Figure 6, we see that the optimal sharing rule for a money-maximizing risk-neutral proposer is (70,30) in the standard game and (90,10) in the convex game. This suggests that more generous proposals in the standard game and more selfish proposals in the convex game may indeed be an optimal behavioral response. However, both within and across each game, there is no statistically significant difference in expected payoffs for offers between (99,1) and (70,30). To see this, we bootstrapped the difference in expected payoff, and a proposer making any offer between \$0.01 and \$0.30 should not expect a significantly different payoff from any of these offers. This is true within each game. Across games there is no statistically significant difference in expected payoffs for offers between (99,1) and (80,20).¹⁵

of the line through the origin acts as the price for reducing the proposer’s payoff. The horizontal projection of the distance from this line to the “Accept Everything” line is the reduction in payoff. It can be seen that the demand curve is downward sloping until the price is 1, that is, when the proposal is a 50-50 split. After that the demand in both games is flat.

¹⁵ We bootstrapped the distribution of responses in each game and observed that the 95-percent confidence interval

Given this lack of significant difference, it is perhaps not surprising that the money-maximizing choices are not the most common selection. However, it may be surprising that they are chosen so infrequently—only 18 percent of the offers in the standard game are (70,30), and no proposer offers (90,10) in the convex game. In fact, over 50 percent of the offers in the standard game are (50,50). The modal choice in the convex game is (99,1) at 34 percent, followed by (50,50) at 30 percent (see Figure 5). Hence, proposers in the standard game could gain by offering less, and proposers in the convex game should be more cautious. The differences observed seem not to be captured simply by optimal responses by proposers.

B. *Could Imperfect Information on Responder Types Matter?*

The result above, that both within and across games the expected return for offers between (99,1) and (70,30) is not significantly different, suggests an interesting alternative. If the mean return is the same, then maybe the difference in choice is driven by differences in risk. Notice that, given similar means, the distribution of returns in the standard game fits the classic definition of a mean-preserving spread on the returns in the convex game (Michael Rothschild and Joseph E. Stiglitz, 1970). Hence the standard game is more risky. For instance, as can be seen in Table 1, the expected return to the proposer from choosing the (99,1) offer will differ across the standard and convex games by only nine cents (out of about \$6.80), a statistically and economically insignificant amount. The risk however, is vastly different. In the standard game, the chance that the (99,1) offer gets a response of dividing zero dollars is 43 percent in the standard game, but only 25 percent in the convex game. Similarly for an offer of (90,10)—expected returns differ by 18 cents out of \$7.25, yet 32 percent of standard game proposers will see a reply of zero dollars, but

only 16 percent will do so in the convex game. Hence, a plausible hypothesis seems to be that by reducing the risk faced by proposers, the convex game is actually forfeiting even more bargaining power from responders, which results in even more aggressive offers by proposers.¹⁶

Note that this finding complements results of Slonim and Roth (1998) in high-stakes ultimatum games. Slonim and Roth's study, however, differs from ours in that their subjects had repeated exposures to the game. The difference between Slonim and Roth's low- and high-stakes players was evident on the first round of play, and it grew over time. Since their experiments are more suggestive of equilibrium behavior than in our one-shot experiments, it suggests future work with a repeated convex ultimatum game could be valuable.¹⁷

VI. Heterogeneity

Recall that we collected choices of subjects in both the role of proposer and responder. The information on responders is pure—someone will reject all or part of an offer if they feel it is unfair. However, the information on proposers is contaminated—if they make generous offers it may be out of generosity or, as seen in the last section, fear. We can use our data to ask, Can the attitudes expressed as responders help us understand and predict behavior as proposers?

Consider this simple classification. We might presume that subjects who fully accept all offers as responders care primarily with maximizing their money earnings, and care little or none for fairness. By contrast, those who do reject some offers demonstrate a willingness to pay for fairness. Are the proposals of these two groups different?

Figure 7 shows the offers by those subjects who will reject some offers when they play the

of offers overlapped. We cannot reject the hypothesis that either expected payoffs for offers ranging from 99 to 70 cents are equal in either treatments or that expected payoffs for offers ranging from 99 to 80 cents are equal across treatments. This result could, obviously, be due to our sample size, especially given the diversity in preferences.

¹⁶ Note that term "risk" here is used in the Rothschild-Stiglitz sense. However, a model of risk in which subjects have a special fear of losing everything—that is, suffering a complete rejection—would also be consistent with this result. See Rabin (1998) for a discussion of alternative notions of risk preferences.

¹⁷ For instance, the arguments presented in this section rely on the assumption that the proposer's expectations of the responder's choice are somewhat accurate. This could be explored more in a setting with repeated play.

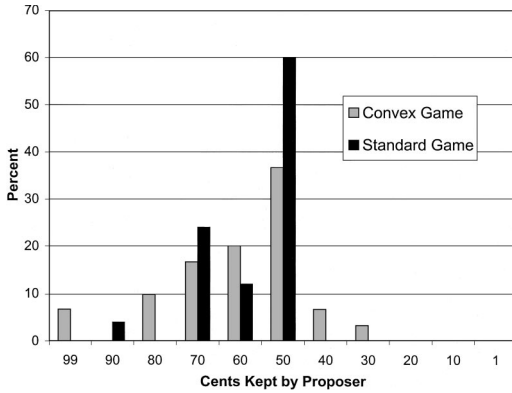


FIGURE 7. OFFERS BY PEOPLE WHO, AS RESPONDERS, DO NOT ACCEPT EVERYTHING

role of responders. These people have demonstrated, as responders, that fairness is important to them. Figure 7 shows the behavior of this group is virtually identical across the two games, which statistical tests confirm.¹⁸ The modal offer is (50,50) in both groups and there are almost no offers below (70,30). Hence, their behavior as proposers is consistent with their behavior as responders, as both illustrate a concern for fairness. In addition, it appears that this concern is not affected by the different opportunities the games offer.

Figure 8 tells a different story. These are people who accept all offers as responders, and thus demonstrate no willingness to pay for fairness. By contrast to Figure 7, this population differs significantly across the two games.¹⁹ Two-thirds offered the least amount of money possible in the convex game, while only one-fifth did so in the standard game. Also, about one-fifth offered an egalitarian sharing rule in the convex game, compared to almost 50 percent in the standard game. Hence, the differences seen across the standard and convex game seem largely due to this group of people. Those who show no concern for fairness are seizing an

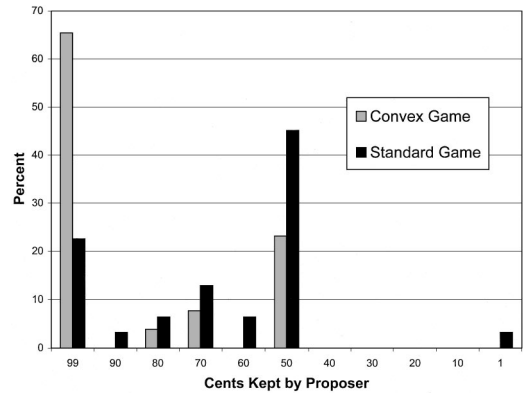


FIGURE 8. OFFERS BY PEOPLE WHO, AS RESPONDERS, ACCEPT EVERYTHING

opportunity to make more money in the convex game, whereas those who care for fairness do not take advantage of the new opportunities. This gives a new and interesting dimension to analysis of experimental games. At the aggregate level it would appear that the convex game diminishes concerns for fairness. But at the individual level this would be incorrect—people who care for fairness are unaffected. Rather, the convex game takes away some bargaining power from those who care for fairness and gives it to those who do not.

VII. Conclusion

Our goal has been to reach a deeper understanding of bargainers' preferences. By convexifying the ultimatum game we were able to uncover these preferences and to confirm that a simple utility function that is continuous, convex, regular, but not monotonic, can rationalize the choices of our subjects. The theoretical model of Rabin (1993), for instance, could capture behavior.

A key finding is that bargainers are very heterogeneous. Overall, about half of the subjects had preferences that could be characterized as monotonic, that is, the responders would not reduce the pie. This behavior is consistent with selfishness among responders. Interesting, this fraction was similar across both the standard and convex games. Second, the convex game revealed that about 7 percent of subjects had linear preferences, either rejecting or accepting all offers, while 41 percent had strictly convex preferences. Among those with convex preferences,

¹⁸ There are 29 subjects from the convex game and 25 from the standard game in Figure 7. Condense the distributions to ranges of the distribution to 99, 90 to 60, and 50 and below, in order to get sufficient numbers in each cell. Then a test of differences in distributions for those in Figure 7 is $\chi^2 = 2.26$, which is insignificant.

¹⁹ There are 27 subjects from the convex game and 31 from the standard game in Figure 8. Repeating the test from the previous footnote for those in Figure 8, we find $\chi^2 = 10.66$, which is significant beyond the 0.01 level.

about half showed a distaste for disadvantageous inequality, while half disliked even advantageous inequality. Finally, 5 percent of all subjects made choices that did not conform to notions of regular preferences.

An unexpected but important result to come from our data is the significance of the risk that responders create for proposers. In both games, the expected return to a given offer is almost exactly the same. The difference is that the possible responses in the ultimatum game are more extreme than in the convex game—the responses in the standard game are approximately a mean-preserving spread of the responses in the convex game. As a result, risk-averse proposers may be more willing to make aggressive offers in the convex game. This means that the convex game may actually have diminished the bargaining power of responders, leading the money-maximizing proposers to make more selfish offers.

While our data revealed heterogeneity across subjects, it also revealed consistency within subjects. By allowing subjects to play the roles of both proposer and responder, we were able to use the concerns for fairness expressed by subjects as responders to help understand their choices as proposers. Half of our subjects chose to reject some offers, thus revealing a value for fairness. As proposers they did not exploit the shifting bargaining power across the games, but made fair offers in both games. The other half—those who accepted all offers—reveal only an interest in making money. These subjects clearly responded to shifts in bargaining power, appearing much more aggressive in the convex game. Looking at the aggregate data, one might be tempted to conclude that the convex game diminished concerns for fairness. By looking at heterogeneity we were able to see that fairness was unaffected by the change in the game—since those who care for fairness were unaffected—but that bargaining power shifted from those who cared about fairness to those who did not.

What do our results suggest for models of fairness? First, a utility-based approach is on sound footing, and preferences of many subjects are indeed convex but nonmonotonic. Second, within a game we must account for the heterogeneity of preferences. For instance, our results suggest that selfishness and altruism are both

relatively stable characteristics of individuals. Indeed, this experiment complements several others that find that about half of all subjects show a clear concern for something other than money-maximization. Third, across games we need to study how people's motivations shift as situations change and how choices are molded by strategic and nonstrategic cues in the environment. Some environments facilitate more altruistic and fair behavior. In our study the simple change of convexifying the choice set of responders caused a dramatic change in many proposers' offers. Fourth, our results suggest that by approaching the data as revealing to us the values and cares of the players, we can, with enough carefully constructed variation, integrate the information into meaningful statements about the underlying preferences.

REFERENCES

- Ahlert, Marlies; Cruger, Arweb and Guth, Werner.** "How Paulus Becomes Saulus—An Experimental Study of Equal Punishment Games." Working paper, Humboldt University of Berlin, 2001.
- Andreoni, James.** "Cooperation in Public Goods Experiments: Kindness or Confusion?" *American Economic Review*, September 1995, 85(4), pp. 891–904.
- Andreoni, James and Miller, John H.** "Giving According to GARP: An Experimental Test of the Consistency of Preference for Altruism." *Econometrica*, March 2002, 70(2), pp. 737–53.
- Bolton, Gary E and Ockenfels, Axel.** "ERC: A Theory of Equity, Reciprocity, and Competition." *American Economic Review*, March 2000, 90(1), pp. 166–93.
- Brandts, Jordi and Charness, Gary.** "Hot vs. Cold: Sequential Responses and Preference Stability in Experimental Games." *Experimental Economics*, March 2000, 2(3), pp. 227–38.
- Camerer, Colin F.** *Behavioral game theory: Experiments in strategic interaction*. Princeton, NJ: Princeton University Press, 2003.
- Eckel, Catherine C. and Grossman, Philip.** "Chivalry and Solidarity in Ultimatum Games." *Economic Inquiry*, April 2001, 39(2), pp. 171–88.
- Fehr, Ernst and Schmidt, Klaus M.** "A Theory of

- Fairness, Competition and Cooperation." *Quarterly Journal of Economics*, August 1999, 114(3), pp. 817–68.
- Forsythe, Robert; Horowitz, Joel; Savin, N. S. and Sefton, Martin.** "Fairness in Simple Bargaining Games." *Games and Economic Behavior*, May 1994, 6(3), pp. 347–69.
- Guth, Werner; Schmittberger, R. and Schwartz, B.** "An Experimental Analysis of Ultimatum Bargaining." *Journal of Games and Economic Behavior*, December 1982, 3(4), pp. 367–88.
- Harbaugh, William; Krause, Kate and Liday, Steve.** "Children's Bargaining Behavior." Working paper, University of Oregon, 2002.
- Ochs, Jack and Roth, Alvin E.** "An Experimental Study of Sequential Bargaining." *American Economic Review*, June 1989, 79(3), pp. 355–84.
- Rabin, Matthew.** "Incorporating Fairness into Game Theory and Economics." *American Economic Review*, December 1993, 83(5), pp. 1281–302.
- _____. "Bargaining Structure, Fairness, and Efficiency." Working paper, University of California, 1997.
- _____. "Psychology and Economics." *Journal of Economic Literature*, March 1998, 36(1), pp. 11–46.
- Roth, Alvin E.** "Bargaining Experiments," in J. H. Kagel and A. E. Roth, eds., *Handbook of experimental economics*. Princeton, NJ: Princeton University Press, 1995, pp. 253–348.
- 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), pp. 1068–95.
- Rothschild, Michael and Stiglitz, Joseph E.** "Increasing Risk I: A Definition." *Journal of Economic Theory*, September 1970, 2(3), pp. 225–43.
- Slonim, Robert and Roth, Alvin E.** "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. 531–54.