

Running head: LEARNING TO BARGAIN

Learning to Bargain

William T. Harbaugh
University of Oregon Economics Department
and National Bureau of Economic Research

Kate Krause
University of New Mexico Economics Department

Lise Vesterlund
University of Pittsburgh Economics Department

Author Note:

This work is based upon work supported by the National Science Foundation under Grant No. 0112157. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

We thank Steve Liday and the Coquille Oregon School District and its teachers and students for their help in collecting and generating the data, and Gerald Seidel, Claude Meidinger, and Martin Dufwenberg for their comments on the paper.

Author's address for correspondence:

William T. Harbaugh
University of Oregon Economics Department
1285 University of Oregon, Eugene, OR, 97403-1285
bill@harbaugh.org

Learning to Bargain

Abstract

This paper studies how children learn to bargain. We performed simple anonymous bargaining experiments with real payoffs with 256 children from age 8 to 18. On average, offers by even the youngest children were close to optimal, given the responses. Both offers and responses were similar to the results that others have reported for adults. Younger children showed more variance in the size of proposals. Children showed clear evidence of reinforcement learning, responding to a rejection by increasing subsequent proposals. This pattern was strongest for the youngest children, who tended to over-react to rejections. We found mixed support for social cognitive theory: while proposals increased after other children made larger proposals, they did not increase after proposals by others were rejected.

Learning to Bargain

A bargain is a mutually agreeable division of benefits. Bargaining is a common and important activity – it is a central feature of most economic transactions, and also of many non-economic interactions. Among adults, variations in bargaining skill can lead to differences in such economic outcomes as salaries, car purchase prices, and divorce settlements. In children, bargaining helps determine the outcomes of a host of day-to-day interactions between friends and within the family, involving things such as who controls a group activity, who does what work, and who gets what share of the family's resources.

We use the ultimatum game to investigate how children learn to bargain in a simple well structured setting. An ultimatum game involves two people. The “proposer” is given a sum of money, and makes an offer on how to split it with a “responder.” If the responder accepts the proposed split, the money is shared accordingly. If the responder rejects the offer, no one gets anything and the game ends. In this game a selfish responder should accept any offer greater than zero. Knowing this, a selfish proposer is expected to make the smallest possible offer to that responder, and to walk off with the bulk of the money.

Of course, this is not what really happens. Many papers have looked at ultimatum-bargaining behavior by adults since the game was first designed by Guth, Schmittberger, and Schwarze (1982). (A review is in Roth, 1995.) Among adults, proposers will typically offer about 40% of the money, and many responders will reject offers of less than 30% – even though they themselves lose by doing so. These results hold even when the amounts of money to be split are several days, or even months, worth of earnings. (Slonim & Roth, 1998 and Cameron, 1999. All these results come from experiments with real cash payments.) People's proposals tend to be remarkably well calibrated – that is, they are close to the amount that maximizes the expected

return, given people's actual probabilities of rejecting different sized proposals. This calibration holds up even though there are large inter- and intra-cultural differences in bargaining behavior. (Roth, Prasnikar, Okuno-Fujiwara, & Zamir, 1991; Henrich et al., 2001; Harbaugh, Krause, & Liday, 2003.) When responders are more likely to accept small proposals, proposers tend to make small proposals.

These cultural differences suggest that bargaining behavior is learned rather than innate. In this paper we study this learning in children, in part because we think childhood is the most likely place to observe this sort of learning, and in part because bargaining by children is intrinsically important. Effective bargaining is particularly important for people with limited resources, and few people have as little direct control over economic resources as do children.

The existing literature on bargaining by children is mostly concerned with styles of negotiation in face to face interactions. For example, Ram and Ross (2001) analyze sibling behavior in a toy sharing task and compare it to reports of other dimensions of the sibling's relationship. Fonzi, Schneider, Tani, and Tomada (1997) look at the correlation between sharing and friendship in eight-year-olds, also in a face to face interaction. These papers both find that children who are more cooperative in other settings are also more cooperative negotiators. There is also a very large literature on children's prosocial behavior. (This is reviewed in Eisenberg and Fabes, 1998, and for a recent example which also looks at learning see Benenson, Markovits, Roy, and Denko, 2003.) This literature includes altruism and sharing, two behaviors which can be important in bargaining situations. The general finding is that older children are more generous.

To our knowledge there are only two papers that directly examine strategic considerations by children in a bargaining setting, where one person can reject another person's

proposal. Murnighan and Saxon (1998) report results from hypothetical ultimatum games by age and gender, and Harbaugh et al. (2003) compare ultimatum bargaining among children of different ages with sharing behavior in a “dictator game,” where the responder does not have the option of rejecting a proposal. Neither of these papers addresses learning. We think the ultimatum game is particularly interesting for studying learning about bargaining because the optimal decisions by proposers depend heavily on how responders will react. The repetitions of the game provide the players with information about responses to different proposals, and with strong feedback about the consequences of mistakes.

There are several papers that study learning by adults in the ultimatum game. Slonim and Roth (1998) show that over ten repetitions of an ultimatum game the mean and within-round variance of proposals decrease, but rejection rates do not change. Duffy and Feltovich (1999) apply a modified reinforcement learning model to repeated ultimatum games, with and without observation of one other player’s proposals and responses. They devise a model which predicts that, with repetition, proposals in the ultimatum game will increase away from the prediction of zero, (because of rejections) and that this increase will be more pronounced when players are able to observe outcomes for other players. In their experiments they find that proposals do increase with repetition when others’ outcomes are observed, but do not increase in the usual limited-information treatment. They also show that, when it is provided, adults do use the information about rejections by the other partner.

Developmental psychologists have long been concerned with questions that are related to strategic behavior and to learning in social settings. While few of the results of this literature are directly transferable to bargaining, some findings are very relevant. The “theory of mind” literature addresses the development of the ability to understand other people’s mental states. For

example, a meta-analysis by Wellman, Cross, and Watson (2001) report that even three-year-olds understand that other people can believe things that are false. Other work, such as Perner and Wimmer (1983) shows that by around age five children can make predictions about people's actions from information about their desires. A good model of what another person knows and how other people make decisions is clearly important in bargaining situations such as the ultimatum game. Since the theory of mind research establishes that these models are developed quite early, we expect to find that even the youngest participants in our experiments have the ability to think carefully enough about other people and their potential actions to act strategically. "Social cognitive theory" is another related area of research, which considers ways by which learning can take place through observation and imitation of others, rather than just through direct experience (See Bandura, 1986.) In combination, the theory of mind and social cognitive theory suggest that our participants can be expected to enter the experiments with the ability to consider and learn from both their own and other participants' strategies.

We use two treatments of the ultimatum game described above to investigate learning about bargaining. In each treatment we repeat the game five times, and each time the children have the same role but a different anonymous partner. The treatments differ in the amount of information the children receive about others' decisions. In the limited-information treatment, children are only shown what their partners do. In the full-information treatment both the proposers and the responders are shown the full distribution of the proposals and responses being made by the others in the experiment. In this way we can clearly determine how the observational learning of social cognitive theory affects participants' strategies over the rounds of the experiment. Social cognitive theory implies that children, particularly the older ones, will

consider the information about proposals and rejections that we present in the full-information treatment.

In the limited-information treatment, a proposer learns only one thing in each round: whether their proposal was acceptable to a responder with whom they will not interact again. This limited information allows for reinforcement learning, but not the kind of observational learning that is important in social cognitive theory. Similarly, in the limited-information treatment responders only learn what offer they received, and not how that offer compares to the other proposals. In the full-information treatment, on the other hand, people learn what sorts of proposals others are making and what sorts of responses are being made to these proposals. Proposals and responses by potential future partners are part of this information.

In addition to providing information about other participants' strategies, the full-information treatment gives participants an opportunity to reveal information that they cannot reveal in the limited-information game. For example, a proposer may make a small proposal to indicate to the other proposers that small proposals are acceptable. A responder might reject a proposal in order to convince proposers, some of whom they will be matched with later, that they need to increase their proposals if they want their proposal to be accepted. None of these motives make sense in the limited-information treatment with different partners for every round.

In this paper we investigate factors that influence ultimatum proposals and responses during the course of the experiment, and the relationship between these factors and age. Specifically, we are interested in how children of different ages adjust their behavior in response to the success or failure of different strategies, and whether they incorporate information about the decisions and payoffs of others. We consider these adjustments to be evidence of learning, and so our focus is on how behavior changes over the five rounds of each experiment.

Our primary hypothesis follows directly from social cognitive theory. We expect proposals in the full-information treatment to converge more quickly than in the limited-information treatment for two reasons. First, if proposals are driven in part by a desire to conform to a norm about fair or reasonable behavior in this game, then proposals should converge toward the mean. Second, the treatment provides more information about what proposal amounts are acceptable to the responders. Even self-interested proposers should incorporate this information and adjust their proposals toward the payoff maximizing level.

In addition, we hypothesize that proposals should start off closer together among the older participants, and that conditional on this, convergence should be quicker. Older children presumably have more consistent beliefs about social norms and more bargaining experience, and more ability to assimilate the available information.

Method

We followed the standard methodology for economic experiments, which forbids any deception by the experimenter and requires that all decisions have real consequences. (Hertwig and Ortmann, 2001 provide an interesting evaluation of the differences between the experimental methods used in psychology and economics, and their consequences.) We told the younger children that they would be bargaining for tokens which they could exchange for toys and school supplies. We told the older children they would get to keep the cash they would be bargaining for. These procedures, and the display of the money or the toys and supplies, got the children's attention. We did not use cash for the younger children because many report that their parents don't allow them to go to the store and spend money on things they want, and therefore it is not as salient a reward as the tokens and the opportunity to spend them immediately. Average payoffs were about \$3 in goods for the younger kids, and about \$7 in cash for the older ones.

According to our surveys, this is several days of discretionary expenditure for children of these ages, and in that sense these might be considered “high-stakes” experiments.

The experiments were done in Coquille, Oregon, a logging and fishing town with an area population of about 10,000. Our sample consists of 256 children aged 8 to 18, with an average age of 6.8, of whom 57 percent were males. We recruited participants by contacting public school teachers and getting permission to run the experiments in their classrooms. While students could choose not to participate, none did. Because of this, and the fact that the vast majority of children of these ages are enrolled in public schools, our samples are very representative of the general population of Coquille. On the other hand, the population of Coquille is not nationally representative – it is a mostly white, working class community.

Children participated in groups of 16, and in one of two treatments. Each treatment had five repetitions of the ultimatum game, with each repetition played against a different randomly chosen and anonymous partner. Each person was randomly assigned the role of proposer or responder, and kept that role throughout the experiment. The experiments were conducted in regular classrooms with children who were enrolled together in class. If a class had more than 16 students, randomly selected extras were sent to the library during the experiment and paid a consolation fee (slightly higher than average earnings) on their return. Our recruiting procedure means that the participants in a given group are homogenous in age, and know each other quite well. Although we don’t have children in every grade in each treatment, the children were approximately balanced across treatments by grades, as shown in Table 1.

Each repetition of the experiment was for 10 tokens. First, the proposers were given a form on which they could mark how their proposed allocation of these tokens. Each responder was then given a form telling them the proposal that had been made to them, and asking them if they

wanted to accept this proposal or reject it, so that neither player would get anything. We then gave each participant their tokens, and started the next round. We told the children that they would not know which person they were matched with, nor would the other know. We identified them using code numbers and we transcribed their decision forms, to prevent the possibility that a partner might recognize their “mark”. All this was explained in advance, and the children seemed convinced by this procedure. Anonymity is the general rule in economic experiments where one person’s decisions can affect the payouts of another. In this experiment it prevents complications from pre-existing relationships, safeguards the children from the possibility of post-experiment retaliation, and ensures that our results are not affected by the anticipation of such retaliation.

Table 1

Numbers of Participants by Grades, Across Treatments

Grade	Limited Information	Full Information	Total
3	16	16	32
4	16	16	32
5	16	16	32
6	32	32	64
9	0	16	16
10	32	16	48
11	0	16	16
12	16	0	16
Total	128	128	256

The “limited-information” treatment is the standard ultimatum-bargaining game. Players know their own history and their partners’ decisions, but nothing else. In the “full-information” treatment everyone is also shown the entire distribution of proposals, and of responses conditional on proposals, in each round. This is done in the following way. After we collect the proposals, the proposers and responders are shown the full distribution of proposals, and responders are shown the proposal made to them. Responders then make their decisions. After we collect those decisions, everyone is shown the full distribution of responses by the corresponding proposal amount. Proposers and responders are then given their earnings for that round (in tokens or quarters), and the next round is begun.

We use bar graphs displayed on an overhead projector to provide information in the full-information treatment. Samples of the displays are shown in Figure 1. All of the children, even the third-graders, had seen bar graphs before as part of their math classes, and many of them were quite proud about telling us this when we asked. We carefully explained the graphs again anyway.

After completing the five-round bargaining experiment, we measured risk aversion by asking children to decide how much they wanted to stake on a better than fair gamble. This gamble consisted of a spinner wheel of the type used in board games, with one third of the wheel colored red and two thirds green. Children could stake up to five of their tokens on this gamble. If the spinner stopped in the red area we kept their stake, if it stopped in the green we gave them an extra token for each one they had staked.

Results

Overall, the data are surprisingly similar to what others have reported for adults. Forty-three percent of the proposals are at the mode of five tokens, and a bit more than 20% of the

proposals are for three or fewer tokens. Responders reject 8% of the four token proposals, and this rate increases monotonically, to 57% of the one token proposals. Interestingly, of the 19 zero token proposals that are made, responders accept 6. So, while the rejection rate for zero token proposals is higher than that for one token proposals, it is still considerably less than 100%.

Table 2A shows descriptive statistics for proposals by round and by grade, both by treatment type. Mean proposals are 0.1 tokens higher in the full-information treatment, but this difference is not significant with a t-test or with a rank-sum or chi-squared test. (Because these tests do not account for the correlation between proposals across rounds, they are biased toward finding a significant difference.) Next we note that the mean proposal falls with repetition in both treatments. It starts higher and falls by more in the full-information treatment. The standard deviation of proposals is included to give a measure of the consistency of proposals across individuals. We might expect this to decrease with repetition, as proposers learn the responses to different proposals and converge to an optimal proposal. In fact, we do not see this decrease in the aggregate data. We will address this issue below, in the subsection that uses regressions to control for individual experience in testing specific learning models.

Table 2B does not show an obvious pattern to the mean proposals across grades. However, as measured by the standard deviation, proposals by the oldest proposers are considerably more consistent than those of the youngest proposers. For Table 2C, we split the entire sample at age 11 years and 9 months, to give two approximately equal sized age groups, and compare how behavior across rounds differs by age. For proposals, we find that the trend for smaller proposals with repetition is similar across ages. The lack of an obvious trend in the standard deviation of proposals means there is no clear convergence of proposals with repetition for either age group.

Table 2A

Summary Behavioral Data, by Treatment and Round

Round	Limited Information			Full Information		
	N	M	(SD)	N	M	(SD)
1	128	4.22	(1.29)	128	4.56	(1.25)
2	128	4.16	(1.25)	128	4.31	(0.99)
3	128	4.00	(1.25)	128	4.23	(1.02)
4	128	4.03	(1.26)	128	4.00	(1.17)
5	128	4.03	(1.22)	128	3.83	(1.20)
Overall	640	4.09	(1.25)	640	4.19	(1.15)

Table 2B

Summary Behavioral Data, by Treatment and Grade

Grade	Limited Information			Full Information		
	N	M	(SD)	N	M	(SD)
3	80	4.13	(1.83)	80	4.30	(1.52)
4	80	3.72	(1.72)	80	3.80	(1.94)
5	80	4.25	(1.03)	80	4.30	(0.69)
6	160	4.03	(1.32)	160	4.45	(0.69)
9	0			80	4.38	(1.03)
10	160	4.11	(0.75)	80	4.28	(0.91)
11	0			80	3.55	(0.81)
12	80	4.32	(0.69)	0		
Overall	640	4.09	(1.25)	640	4.19	(1.15)

Table 2C

Summary Behavioral Data, by Age and Round

Round	Young			Old		
	N	M	(SD)	N	M	(SD)
1	126	4.40	(1.54)	130	4.38	(0.97)
2	126	4.35	(1.29)	130	4.12	(0.94)
3	126	4.15	(1.35)	130	4.09	(0.91)
4	126	3.94	(1.44)	130	4.09	(0.96)
5	126	3.87	(1.50)	130	3.98	(0.87)
Overall	630	4.14	(1.43)	650	4.13	(0.93)

Figure 2 shows the distribution of proposals, and responses conditional on proposals, across treatments, rounds, and age. The first two rows compare the treatments. The first round proposals are fairly similar, with five as the modal proposal and a smaller spike at four. A general trend away from proposals of five tokens and towards smaller proposals can be seen. The trend is more pronounced in the limited-information treatment. There the modal proposal has fallen to four tokens by the last round.

Rows three and four compare results across treatments for the younger children. There is very little difference between proposals in the two treatments until the last round, when those participating in the limited-information treatment are much more likely to increase their proposals to five tokens. Rows five and six, for the older children, show that their proposals in the limited-information treatment are initially smaller than in the full information. Rejections appear to be a bit more frequent in the limited-information treatment, and in the last two rounds proposals in that treatment are higher than in the full-information version.

Exploratory Regressions

Table 3 shows ordinary least squares regressions for the proposals, and probit regressions for the rejection decisions shown in figure 2. (We note that while the measures of statistical significance from these regressions are adjusted to account for the correlation across decisions by the same person, they are still biased towards significance because they do not account for the possibility of correlations within sessions of the experiment.) These regressions generally confirm the impressions from the descriptive statistics. For proposals, grade alone has no effect while the decline in proposals over the rounds is relatively large and statistically significant in the full-information treatment. To test whether this decline differed systematically by the age of

the participant, we also estimated regressions with a round X grade interaction term. The estimated coefficient was not significant.

The partner's proposal amount has a large and highly significant effect in the rejection decision regressions. Older children reject more frequently, but the effect is not significant. We find weak evidence for the possibility of a demonstration effect, where small proposals are rejected in early rounds in an effort to induce larger proposals in later rounds. While the general increase in rejections over rounds (controlling for the proposal) would seem to suggest against this, if we include dummy variables for the rounds to allow for a nonlinear effect, we find the rejections peak in round four, and then drop. (Available from authors.) However, like the linear effect, this pattern is not statistically significant, though it is larger and closer to significance among the older children. Given the relative complexity of the demonstration effect strategy, it seems plausible that older participants, if anyone, would be more likely to use it.

The variable "gamble" is the number of tokens, from zero to five, that the person staked on the gamble. It can be interpreted as a measure of risk attitudes, with low values implying risk aversion. (Boys are on average slightly less risk averse by this measure, but the difference is not significant. Older children are significantly less risk averse.) In theory, more risk averse people should be expected to make larger proposals, since larger proposals act as a form of insurance against variance in payoffs. However, we don't find this among children - this variable's only statistically significant relationship with behavior is at the 0.10 p-value, for limited-information rejections. In that regression children who make larger gambles are more likely to reject. We do not find any systematic difference in proposals or responses between boys and girls.

Table 3

Exploratory Regressions

	Proposals		Rejections	
	Limited Info	Full Info	Limited Info	Full Info
Grade	0.0320 (0.0378)	-0.0329 (0.0322)	0.0728 (0.0703)	0.0924 (0.1213)
Round	-0.0500 (0.0399)	-0.1781*** (0.0392)	0.1127 (0.0999)	0.0774 (0.1237)
Male	0.0980 (0.2286)	-0.2666 (0.1729)	-0.0712 (0.3729)	0.4648 (0.5672)
Gamble	-0.0236 (0.0722)	-0.0112 (0.0537)	0.2474* (0.1343)	0.1183 (0.1903)
Partner's Proposal			-0.8853*** (0.1532)	-1.3296*** (0.4097)
R-squared	0.010	0.068		
Observations	320	320	320	320
Participants	64	64	64	64

Notes:

Panel regressions, ordinary least squares for proposals, probit for rejections.

Marginal effects shown (at the means for probit), with standard errors in parentheses.

* means significant at 10%, ** at 5%, *** at 1%.

Next we turn to the question of whether children are learning to make optimal proposals over the rounds of the experiment. We estimate a probit model in which we regress rejections on proposals, by treatment. We then use the resulting equations to predict earnings from different sized proposals for each treatment. The results are shown in Table 4. Given the estimated rejection behavior, proposals of either three or four tokens maximize the expected own payoff to the proposers, but proposals of four tokens or more are less risky, so it's not surprising to see that actual proposals are usually four or higher. (However, recall that our measure of risk aversion does not show a statistically significant relationship with proposals.) Proposals of five tokens, on the other hand, lead to noticeably lower own payoffs, so it makes sense that their frequency decreases over time.

Table 4

Predicted Payoffs to Proposers

Proposal Amount	Predicted Payoff to Proposer	
	Limited Information	Full Information
0	1.9	1.9
1	3.4	3.4
2	4.8	4.9
3	5.5	5.7
4	5.5	5.6
5	4.9	4.9
6	4.0	4.0

Specific Learning Models

We next examine learning by looking at variation in behavior. Over the course of the experiment ninety percent of the children vary their proposals, and about half change them by two or more tokens. Splitting the sample at the median age of 11 years and 9 months, we find that the younger children tend to make larger changes. Their average range is 2.5, versus 1.5 for the older children, and this difference is significant with a p-value of 0.001 or less with any of the usual parametric or nonparametric tests.

We might expect less experimentation in the full-information treatment, since proposers typically see responses to a range of proposals every round, and so the marginal value of their own experimentation is presumably lower. Average proposal ranges are 1.9 in the full versus 2.0 in the limited-information treatment, but this difference is not statistically significant with the usual tests.

However, while we find evidence of significant amounts of experimentation over rounds, Table 5 shows that this actually seems to lead to smaller rather than larger earnings among the proposers in later rounds of the experiment. Additionally, a simple regression of the range of proposals and grade on proposer's final earnings gives a significantly negative coefficient for the range of proposals. One interpretation of these results is that the children are experimenting, but five rounds is not enough time to expect much learning. Learning models often distinguish between a period of exploration, when people accumulate information, and exploitation, when they take advantage of what they have learned. While it would be optimal to stop experimenting relatively quickly in a game with only a few periods, children may not realize this, or they may be innately curious and persist in experimentation even though it is not optimal in terms of cash payoffs.

Table 5

Earnings by Round

Older Children

Round	Proposers		Responders	
	Limited Information	Full Information	Limited Information	Full Information
1	5.09	4.86	4.06	4.48
2	4.79	4.97	4.06	4.39
3	5.00	5.21	3.85	4.13
4	5.39	4.97	3.88	3.71
5	4.70	5.10	3.73	3.52

Younger Children

Round	Proposers		Responders	
	Limited Information	Full Information	Limited Information	Full Information
1	5.16	4.97	4.13	4.42
2	4.77	5.26	3.55	3.82
3	4.97	5.43	3.35	3.97
4	5.16	4.83	3.68	3.39
5	4.65	5.23	3.48	3.64

Note: amounts are in tokens, worth \$0.10 for younger children, \$0.25 for older.

Next we consider learning among the responders. One measure of variation among responders is the consistency of their responses to a given proposal. For example, 73 responders receive at least two proposals of four tokens, and 11 of these responders give different responses. Of the 15 responders who received more than one proposal of three tokens, two made differing responses. In short, different responses to the same proposals are rare, but not unheard of.

The different responses can be explained either as errors, as the outcome of a purposefully randomized strategy, or by assuming that (some) experimentation is occurring among the responders. For example, they may be unsure how they will feel after a rejection, and some may try rejecting or not to find out. Another possibility, particularly in the full-information treatment, is that the responders reject early offers in order to induce proposers to make larger proposals. However, we have already found that, conditional on proposals, there is no significant decline in rejection in the later rounds of the experiment. In fact if anything the rate increases. This demonstration effect is not supported by the data. Since all these possible explanations are problematic, we are encouraged by the fact that there is relatively little experimentation among the responders, both in absolute terms and relative to the degree of experimentation among the proposers.

To further explore learning by proposers we use regressions on the change in the proposals, with results shown in Table 6. As would be expected, in the limited-information treatment (regression 1) we find strong evidence for reinforcement learning. When a person experiences a rejection they make a larger proposals next round, and this effect is large and significant. The coefficient on the interaction of the size of the lagged proposal and rejection of that proposal is negative, showing that the increases in proposals are reduced when the rejected proposal was already large. Again, this is as expected, though it is only marginally significant in

this regression. Regression 2 shows the same basic pattern for the full-information treatment, with the interaction effect now being significant. This result is consistent with what others have found with adults, and we think the fact that it is also found in children shows evidence of fairly sophisticated responses to direct experience.

Regression 3 considers social cognitive learning by showing the effect of the public information that is available in the full-information treatment. A likelihood-ratio test favors this model over the model without this information, with a p-value of 0.001, and we can see that the R-squared increases substantially. On this basis we conclude that the public information does have an effect on proposals. Specifically, larger proposals by others lead to larger own proposals next round. Rejections by the partners of other players, on the other hand, do not have a significant effect on own proposals, though the coefficient is positive.

Regression 4 adds interactions of each kind of lagged information with grade. This allows us to consider whether children of different ages use information differently. We find that older children's proposals are actually less sensitive to rejections of their own previous proposals, and more sensitive to previous proposals by others (though the statistical significance of this last effect is marginal.) Again, rejections by partners of other proposers are not significant, and there is no significant age interaction. (Gender and the risk aversion measure are not significant in these regressions.)

The estimated age differences in the responsiveness to rejections that are implied by these interaction effects are quite large. Our regression estimates imply that on average eight-year-olds increase a proposal of 3 tokens by 1.7 tokens after a rejection, while eighteen-year-olds increase theirs by only 0.3 tokens. This sort of difference seems plausible, given that the older participants have more prior experience to apply to the situation. Furthermore, the older participants may

have been better able to understand the implications of being paired with a different partner in each round.

Table 6
Learning Regressions

	1	2	3	4
	Limited Info	Full Info	Full Info	Full Info
Own Rejection, Lagged	2.335*** (0.643)	3.005*** (0.626)	2.795*** (0.672)	3.573*** (0.892)
Own Proposal, Lagged X Own Rejection, Lagged	-0.328* (0.194)	-0.508*** (0.164)	-0.483** (0.187)	-0.404** (0.169)
Rejection Rate for Others, Lagged			0.444 (0.632)	1.556 (2.040)
Average Proposal by Others, Lagged			0.488*** (0.133)	0.477*** (0.126)
Grade X Own Rejection, Lagged				-0.153** (0.065)
Grade X Rejection Rate for Others, Lagged				-0.171 (0.284)
Grade X Mean Proposals by Others, Lagged				0.0122 (0.0077)
Constant	-0.220*** (0.0524)	-0.330*** (0.0604)	-2.124*** (0.520)	-2.539*** (0.584)
R-squared	0.17	0.17	0.23	0.24
Observations	256	256	256	256
Subjects	64	64	64	64

Notes: Dependent variable is the change in proposal. Standard errors in parentheses.
* is significant at 10%, ** at 5%, *** at 1%.

Conclusion

In ultimatum-bargaining games, children make what appear to be very close to optimal proposals from the beginning. This result is consistent with the theory of mind results that show that even very young children have the sort of understanding that seems necessary for playing these sorts of simple social games. We also find that children respond appropriately to direct experience. If their bargaining proposal is rejected, they tend to respond by increasing their subsequent proposal, and if it is accepted they tend to decrease their subsequent proposal. These changes are very large in young children, and much smaller in older ones. This age effect seems entirely sensible. Younger children can be expected to have less experience and therefore less confidence in their prior beliefs.

We also find some evidence of social or observational learning. When shown information about what other proposers are doing, children tend to make proposals that match that information. They do not, however, appear to use information about whether other players' proposals are rejected, and the changes they make in their decisions do not lead to higher earnings.

We think these results can most easily be interpreted in terms of the distinction between two basic models of learning – reinforcement and social cognitive theory. Suppose that the children playing these games care about their own payoff and also have a desire to conform to what everyone else is doing. Reinforcement learning occurs during the course of the experiment as they observe their payoffs in terms of these two measures. The reinforcement regarding their taste for their own payoff comes from seeing their payoffs from various proposals, and the reinforcement regarding their taste for conformity comes from comparing their proposal with the

proposals of others, when that information is available. The relevant distinction between the two protocols is that only the full-information one provides information about the proposals of others.

However, it is also possible that the information about others' proposals is used for motives other than a simple desire to do what others are doing. For example, it is possible that players assume that the decisions of the other proposers are a rational and self-interested reaction to the rejections those players have experienced. Matching their proposals is therefore a way of incorporating the experience of other players. However, this explanation is difficult to reconcile with the finding that children do not respond to information about the rejection rates that others face.

We began by asking how children learn the skills necessary to become effective bargainers. We find strong evidence that they respond to direct experience and conclude that children will refine their own strategies as they gain more direct bargaining experience. We also find that proposers tend to learn from the proposals they observe others making. By extension, a child who observes an adult bargaining can learn to use the same sorts of strategies in similar situations. However, a child might not learn to anticipate others' responses by observing bargaining among adults. In our experiments, the proposers did not adapt to the responses that other proposers received. Perhaps this information involved too many cognitive steps. Our participants learned through direct experience and by watching other participants who were in the same situation. However, our participants did not learn by watching the responses of others to children who were in their same situation. While the theory of mind literature indicates that children age eight and above are capable of thinking about what the other person might do, the children in our experiments did not use all of the available information to refine their bargaining strategies.

References

- Bandura, A. (1986). *Social foundations of thought and action*. Englewood Cliffs, NJ: Prentice-Hall.
- Benenson, J. F., Markovits, H., Roy, R., & Denko, P. (2003). *Behavioral rules underlying learning to share: Effects of development and context*. Unpublished manuscript.
- Cameron, L. A. (1999). Raising the Stakes in the Ultimatum Game: Experimental Evidence from Indonesia. *Economic Inquiry*, 37(1), 47-59.
- Duffy, J., & Feltovich, N. (1999). Does observation of others affect learning in strategic environments? An experimental study. *International Journal of Game Theory*, 28(1), 131-152.
- Eisenberg, N., & Fabes, R. A. (1998). Prosocial Development. In W. Damon (Ed.), *Handbook of child psychology* (5th ed., Vol. 3, pp. 701-778). New York: J. Wiley.
- Fonzi, A., Schneider, B. H., Tani, F., & Tomada, G. (1997). Predicting children's friendship status from their dyadic interaction in structured situations of potential conflict. *Child Development*, 68(3), 496-506.
- Guth, W., Schmittberger, R., & Schwarze, B. (1982). An Experimental Analysis of Ultimatum Bargaining. *Journal of Economic Behavior and Organization*, 3(4), 367-388.
- Harbaugh, W. T., Krause, K., & Liday, S. (2003). *Bargaining by Children*. Unpublished manuscript.
- Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H., et al. (2001). In search of Homo economicus: Behavioral experiments in 15 small-scale societies. *American Economic Review*, 91(2), 73-78.

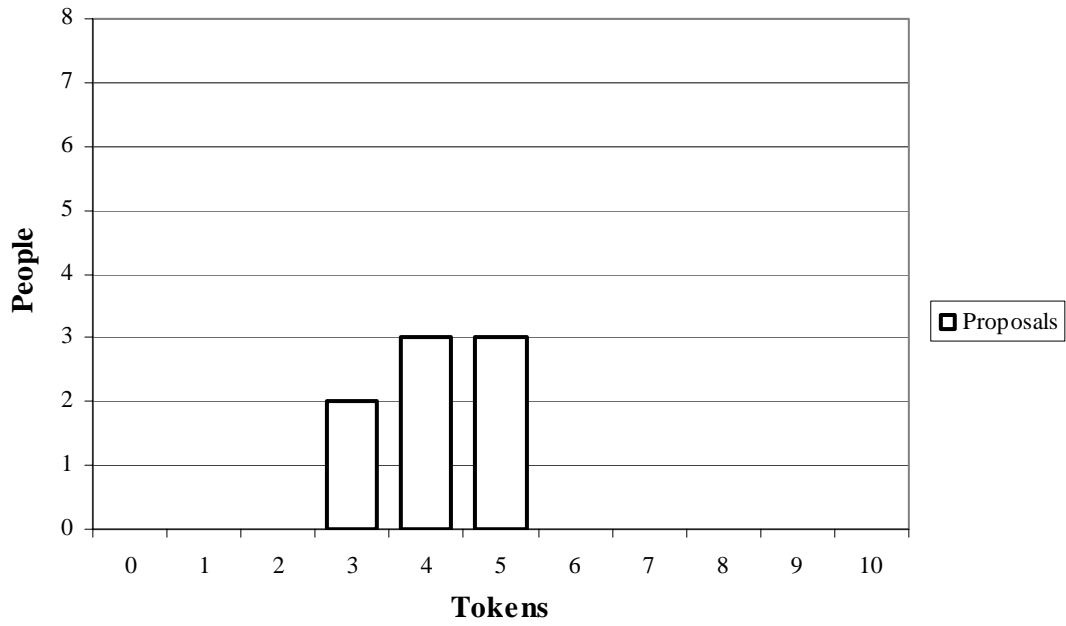
- Hertwig, R., & Ortmann, A. (2001). Experimental Practices In Economics: A Methodological Challenge For Psychologists? *Behavioral and Brain Sciences*, 24(3), 383-403.
- Murnighan, J. K., & Saxon, M. S. (1998). Ultimatum Bargaining by Children and Adults. *Journal of Economic Psychology*, 19(4), 415-445.
- Perner, J., & Wimmer, H. (1983). Beliefs about beliefs: representation and constraining function of wrong beliefs in children's understanding of deception. *Cognition*, 13, 103-108.
- Ram, A., & Ross, H. S. (2001). Problem-solving, contention, and struggle: How siblings resolve a conflict of interests. *Child Development*, 72(6), 1710-1722.
- Roth, A. E. (1995). Bargaining Experiments. In J. H. Kagel & A. E. Roth (Eds.), *The Handbook of Experimental Economics*. Princeton: Princeton University Press.
- Roth, A. E., Prasnikar, V., Okuno-Fujiwara, M., & Zamir, S. (1991). Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study. *American Economic Review*, 81(5), 1068-1095.
- Slonim, R., & Roth, A. E. (1998). Learning in High Stakes Ultimatum Games: An Experiment in the Slovak Republic. *Econometrica*, 66(3), 569-596.
- Wellman, H. M., Cross, D., & Watson, J. (2001). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development*, 72(3), 655-684.

Figure Captions


Figure 1. Full information displays.

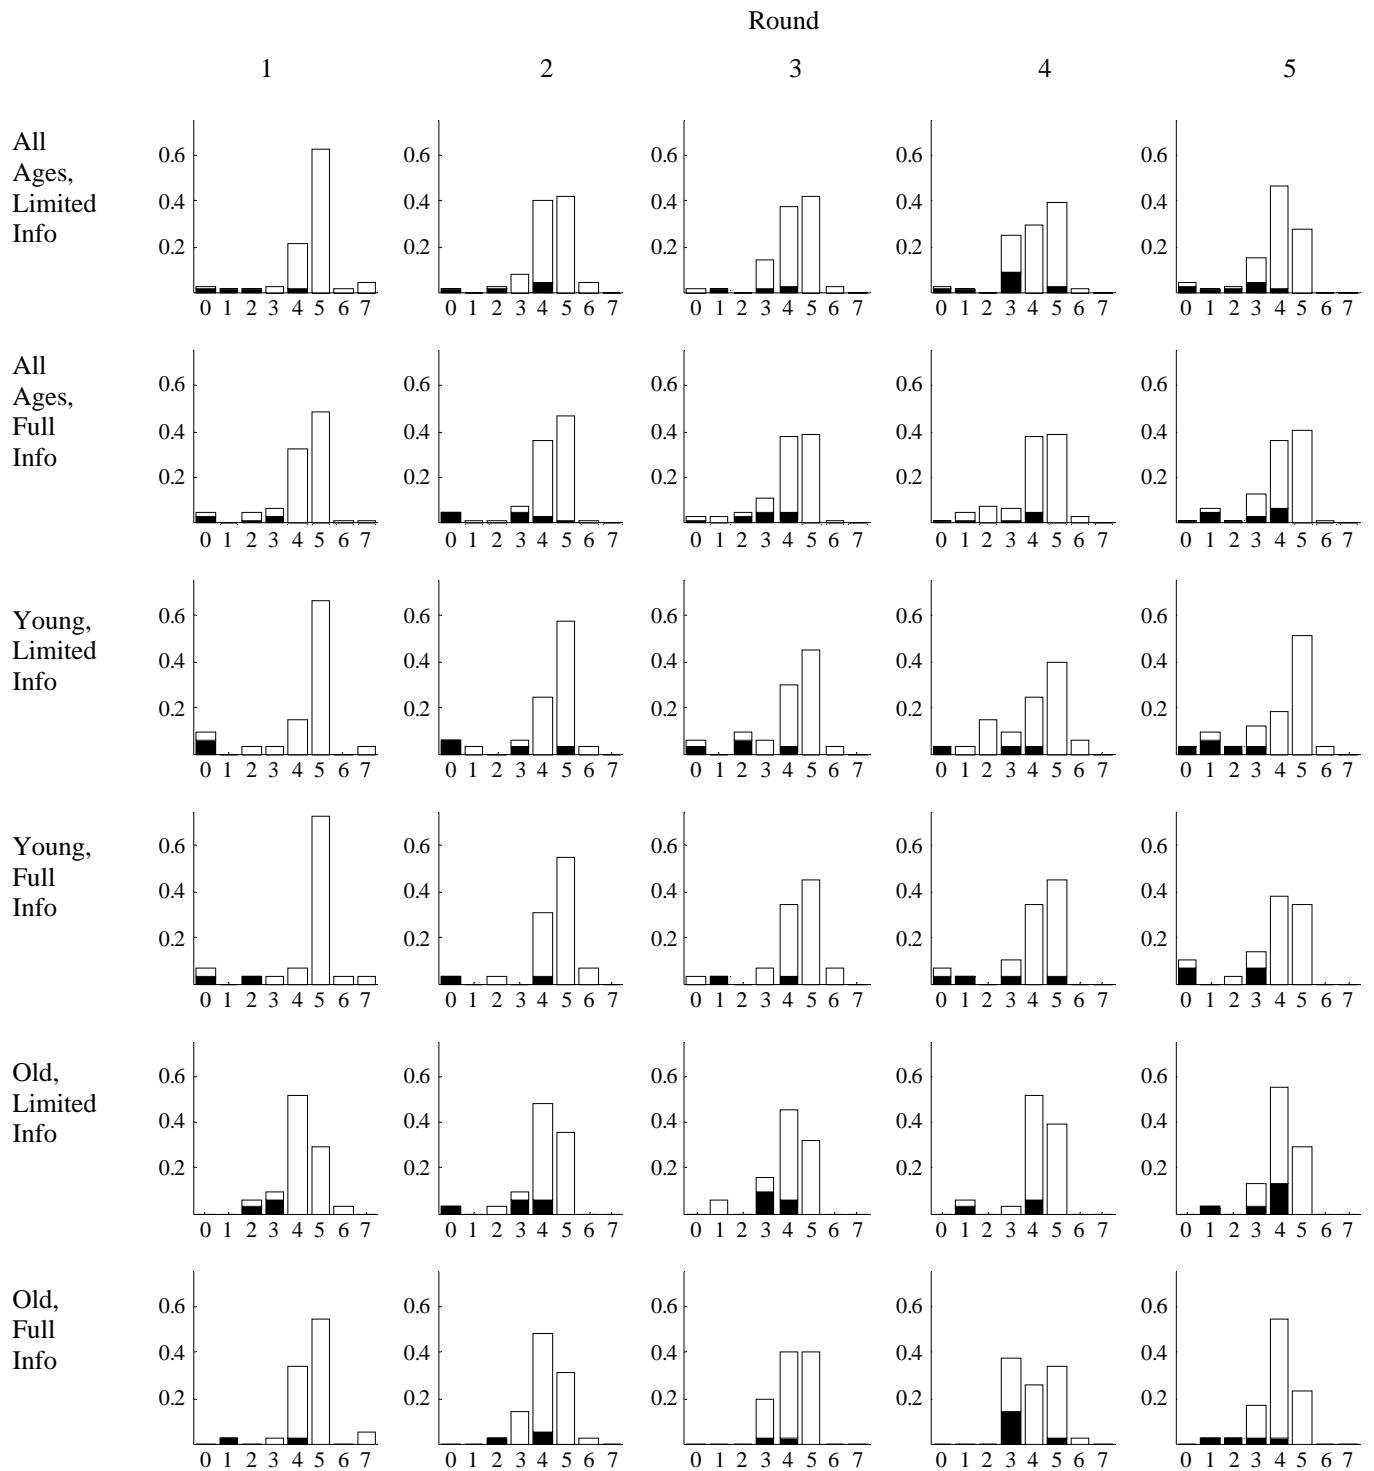
Figure 2. Distribution of proposals and rejections, by treatment and round.

Proposals



Responses





Notes: The horizontal axes give the proposal size, vertical axes give the proportion of observations. The total bar height gives the proportion of the sample making proposals of that size, the height of the solid portion shows the proportion of the sample who received and rejected a proposal of that size.