



Learning to bargain

William T. Harbaugh^{a,b,*}, Kate Krause^c, Lise Vesterlund^d

^a *Department of Economics, University of Oregon, 1285 University of Oregon, Eugene, OR 97403-1285, United States*

^b *National Bureau of Economic Research, Cambridge, MA 02138, United States*

^c *Department of Economics, University of New Mexico, Albuquerque, NM 87131, United States*

^d *Department of Economics, University of Pittsburgh, Pittsburgh, PA 15260, United States*

Received 21 July 2005; received in revised form 21 March 2006; accepted 30 March 2006

Available online 30 June 2006

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. Average offers by even the youngest children were close to the amount that maximized their expected monetary payoff. Furthermore offers and responses were similar to what others have reported for adults. The variance of proposals was higher among younger children. 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 respond to rejections with larger than payoff-maximizing increases in proposals. We found mixed support for social learning: while proposals increased after other children made larger proposals, they did not increase after proposals by others were rejected.

© 2006 Elsevier B.V. All rights reserved.

JEL classification: C70; D10; D63

PsycINFO classification: 2840; 3020; 3040

Keywords: Children; Ultimatum game; Bargaining; Learning

* Corresponding author. Address: Department of Economics, University of Oregon, 1285 University of Oregon, Eugene, OR 97403-1285, United States. Tel.: +1 541 346 1244; fax: +1 541 346 1243.

E-mail address: bill@harbaugh.org (W.T. Harbaugh).

1. Introduction

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 actually happens. Many papers have looked at ultimatum-bargaining behavior by adults since the game was first designed by [Guth, Schmittberger, and Schwarze \(1982\)](#), (see [Roth, 1995](#) for a review). Among adults, proposers will typically offer about 40% of the money, and many responders will reject offers of less than 30% – even though they decrease their earnings by doing so. These results come from experiments with real cash payments, and they generally hold even when the amounts of money to be split are equivalent to several days, or even months, worth of earnings, as in [Slonim and Roth \(1998\)](#) and [Cameron \(1999\)](#). People's proposals tend to be remarkably close to what is optimal – that is, they are close to the amount that maximizes the proposer's expected return, given responders' actual probabilities of rejecting different sized proposals. Interestingly this result holds despite large inter- and intra-cultural differences in bargaining behavior. That is, proposers tend to make small proposals in environments where responders are more likely to accept small proposals as discussed in [Roth, Prasnikar, Okuno-Fujiwara, and Zamir \(1991\)](#), [Henrich et al. \(2001\)](#), [Harbaugh, Krause, and Liday \(2003\)](#).

These cultural differences suggest that to some extent bargaining behavior is learned rather than innate. In this paper we study such 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.

To study how children learn to bargain we had 256 children from age 8 to 18 participate in a bargaining game. We use two treatments to examine how learning takes place. In each treatment the game is repeated five times. A child has the same role for each repetition, but is matched with a different anonymous partner each time. 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 distributions of proposals and responses that are made by others in the experiment. This allows us to determine how observational learning affects participants' strategies over the course of the experiment.

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 such adjustments to be evidence of learning, and we focus on how behavior changes over the five rounds of each experiment.

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. Similarly, using face-to-face interaction, [Fonzi, Schneider, Tani, and Tomada \(1997\)](#) look at the correlation between sharing and friendship in eight-year-olds. Both of these papers 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, as reviewed in [Eisenberg and Fabes \(1998\)](#), with [Benenson, Markovits, Roy, and Denko \(2003\)](#) as a recent example with learning. This literature examines 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. By repeating the game players are provided 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 10 repetitions of an ultimatum game the mean and within-round variance of proposals decrease, but rejection rates do not change. [Duffy and Felto- vich \(1999\)](#) 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 because of rejections, proposals in the ultimatum game will with repetition increase away from the prediction of zero, and this increase will be more pronounced when players are able to observe outcomes for other players. Consistent with their prediction they find that proposals increase with repetition when others' outcomes are observed, however they do not increase in the usual limited-information treatment. They also show that, when it is provided, adults use the information about rejections by other responders.

Developmental psychologists have long been concerned with questions regarding strategic behavior and learning in social settings. This literature is large, and we will only discuss the most relevant parts. 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 might expect 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. However, bargaining is a relatively complex social interaction, and it may instead be argued that the ability to interact with others in a strategic bargaining setting improves long after children are able to master the relatively simple theory of mind tasks. For example, [Gurucharri and Selman \(1982\)](#) argue that it is not until around age 11 that children can balance their own viewpoint with that of another person, an ability that is precisely what is called for in the ultimatum game.

Psychologists have long emphasized the importance of the sort of social learning that occurs through observation and imitation of others as in [Bandura \(1986\)](#). In combination, the theory of mind and principles of social learning suggest that our participants can be expected to enter the experiments with some ability to consider and learn from both their own and other participants' strategies. We would also expect that there might be a considerable increase in these abilities as we look at older children.

In addition to this general work, there is also research in developmental psychology that focuses specifically on how children learn to engage in economic interactions, as reviewed in [John \(1999\)](#) and [Webley \(2004\)](#). Presumably because of the public policy implications, much of this literature focuses on the question of how children respond to advertising and peer pressure. The advertising work is wide ranging, and perhaps most relevant are studies that examine children's ability to acquire and use information rationally. The general conclusion is that children develop this skill quite quickly. For example, [Gregan-Paxton and John \(1995\)](#) find experimental evidence that suggest that children as young as 6–7 years old incorporate cost and benefit information in their product search decisions. In general, the results that are reported in the developmental psychology literature for the sorts of knowledge and skills that are needed in our bargaining task suggest that we should expect changes in bargaining behavior as we move from eight-year-old to older subjects.

As mentioned earlier we will use a limited and a full-information treatment to study learning by children. 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 learning or 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 proposals others are making and what responses these proposals are getting. Thus participants are in part receiving information on the proposals and responses that are made by potential future partners.

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. Our primary hypothesis follows directly from principles of observational learning. 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 level that maximizes the expected monetary payoff.

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 may be better at assimilating the available information.

2. Method

We followed the standard methodology for economic experiments, which forbids any deception by the experimenter and requires that all decisions have real consequences. [Herwig 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. The children in ninth grade and above were told that 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 do not 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 those in ninth grade and above. 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 8–18, with an average age of 12.8, of whom 57% 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

Table 1
Numbers of participants and average age, by grades, across treatments

Grade	Average age	Limited information	Full information	Total
3	8.8	16	16	32
4	10.0	16	16	32
5	10.9	16	16	32
6	11.9	32	32	64
9	15.3	0	16	16
10	16.0	32	16	48
11	16.7	0	16	16
12	17.8	16	0	16
Total	12.8	128	128	256

their return. Our recruiting procedure means that the participants in a given group are homogenous in age, and know each other quite well. However the children did not know which of their classmates they were matched with. Although we do not have children in every grade in each treatment, the children were approximately balanced across treatments by grades, as shown in Table 1. This table also shows the average ages of the children in the experiment, by grade. In the US, children are typically six years old when they begin first grade.

Each repetition of the experiment was for 10 tokens. First, the proposers were given a form on which they could mark their proposed allocation of these tokens. Each responder was then given a form telling them the proposal that had been made to them, and were asked if they wanted to accept this proposal or reject it. In the event of a rejection 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 anyone else in the experiment know the identity of their partner. We identified them using code numbers and we transcribed their decision forms, to prevent the possibility that a partner might recognize their “mark.” These procedures were explained in advance, and the children seemed convinced by them. Anonymity is generally used 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. This said, we can of course not prevent the children from talking about their own decisions – truthfully or not – after the experiment.

We refer to our two treatments as the “limited information” and the “full-information” treatment. 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.

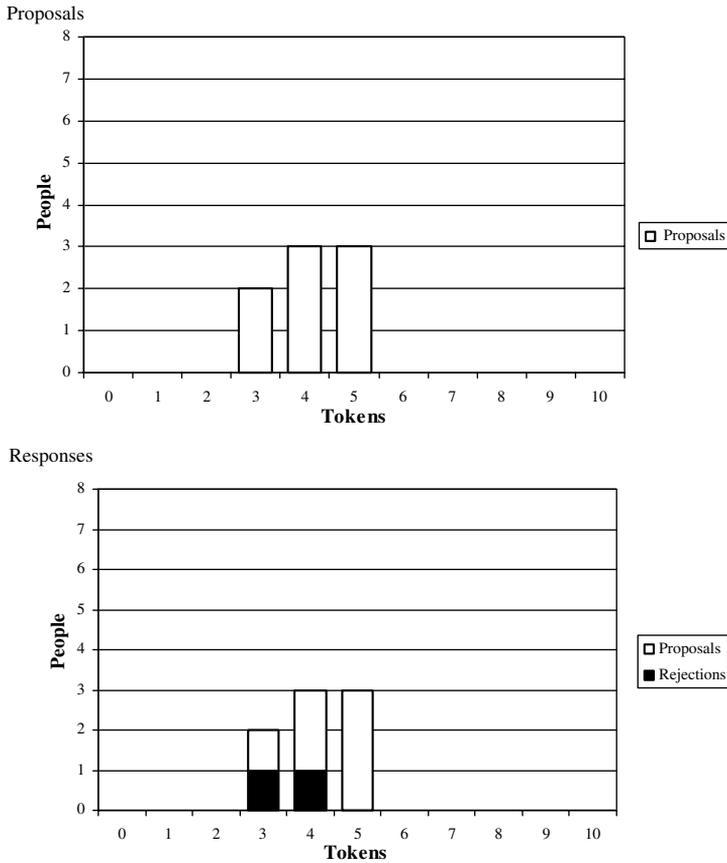


Fig. 1. Sample full-information displays.

We display bar graphs on an overhead projector to provide information in the full-information treatment. Samples of the displays are shown in Fig. 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 proud to tell 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.

3. 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 only slightly 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, and then 68% of the zero token proposals. Surprisingly, 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 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 a proposal that maximizes the expected monetary payoff. However, we do not see this decrease in the aggregate data. We will address this issue 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

Table 2A
Proposals by treatment and round

Round	Limited information			Full information		
	<i>N</i>	Mean	(SD)	<i>N</i>	Mean	(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
Proposals by treatment and grade

Grade	Limited information			Full information		
	<i>N</i>	Mean	(SD)	<i>N</i>	Mean	(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
Proposals by age and round

Round	Young			Old		
	<i>N</i>	Mean	(SD)	<i>N</i>	Mean	(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)

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.

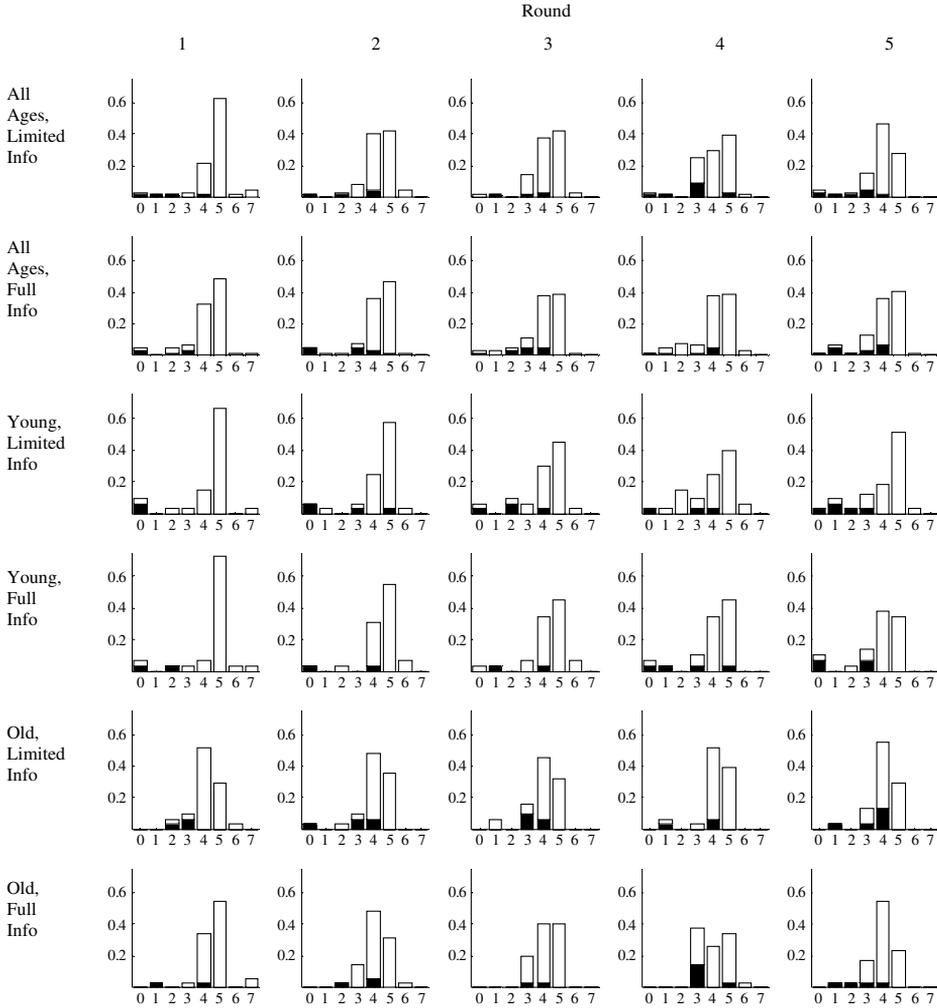
Fig. 2 shows the distribution of proposals, and responses conditional on proposals, across treatments, rounds, and age. The first two rows compare the treatments. The first round proposals are fairly similar, with five as the modal proposal and a smaller spike at four. A general trend away from proposals 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 with 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.

3.1. Exploratory regressions

Table 3 shows ordinary least squares regressions for the proposals, and probit regressions for the rejection decisions shown in Fig. 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, it is neither large nor significant in the limited-information treatment. To test whether this decline differed systematically by the age of the participant, we also estimated regressions that include a term for the interaction between round and grade. The estimated coefficient was not significant.

Not surprisingly the amount offered to a responder 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



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.

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

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 being significant for 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,

Table 3
Exploratory regressions

	Proposals		Rejections	
	Limited information	Full information	Limited information	Full information
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)
<i>R</i> -squared (pseudo for probit)	0.010	0.068	0.33	0.29
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%.

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.

Next we turn to the question of whether the children are learning to make proposals that maximize their expected monetary payoff, 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 is 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

Table 4
Predicted monetary payoffs to proposers conditional on proposal

Proposal amount	Predicted monetary 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

tokens, on the other hand, lead to noticeably lower own payoffs, thus it is not surprising that their frequency decreases over time.

3.2. Specific learning models

We next examine learning by looking at variation in behavior. Over the course of the experiment 90% 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. The average range between the largest and smallest proposal that an individual offers is 2.5 for the younger children, versus 1.5 for the older children. This difference is significant with a p -value of 0.001 or less with any of the usual parametric or non-parametric tests.

We might expect less experimentation in the full-information treatment, since proposers typically see responses to a range of proposals every round, and the marginal value of their own experimentation is therefore likely to be lower. The average ranges between the largest and smallest proposal that individuals offer are 1.9 in the full versus 2.0 in the limited-information treatment. This difference is not statistically significant with the usual tests. We note that the result that even young children are searching for payoff-maximizing strategies is consistent with the results in Gregan-Paxton and John, reported earlier. 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.

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.

Table 5
Average earnings by round

Round	Proposers		Responders	
	Limited information	Full information	Limited information	Full information
<i>Older children</i>				
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
<i>Younger children</i>				
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 in goods for the younger children, and \$0.25 in cash for older children.

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. Thus the demonstration effect is not supported by the data. Since all these possible explanations are questionable, 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. The results are 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 larger proposals in the following round, and this effect is large and significant. The coefficient on the interaction of the size of the lagged proposal (i.e., the proposal in the previous round) 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 their personal bargaining experience.

Table 6
Learning regressions

	1 Limited information	2 Full information	3 Full information	4 Full information
Own rejection lagged	2.335*** (0.643)	3.005*** (0.626)	2.795*** (0.672)	3.573*** (0.892)
Interaction between own proposal, lagged and 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)
Interaction between grade and own rejection, lagged				-0.153** (0.065)
Interaction between grade and rejection rate for others, lagged				-0.171 (0.284)
Interaction between grade and 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 (adjusted)	0.17	0.17	0.22	0.22
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%.

Regression 3 considers social learning and social referencing, by showing the effect public information of the full-information treatment has on behavior. 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 adjusted 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. Although the rejection by other responders has a positive effect on proposals, the effect is not significant.

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. Interestingly we find that the proposals of older children are 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 other responders do not have a significant effect on proposals, 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 three tokens by 1.7 tokens after a rejection, while eighteen-year-olds increase theirs by only 0.3 tokens. This difference can perhaps be explained by older participants having more prior bargaining experience. Furthermore, the older participants may have been better able to understand the implications of being matched with a different partner in each round.

4. Discussion

In ultimatum-bargaining games, children make proposals that are very close to the level that maximizes the expected monetary payoff from the beginning. Their proposals are very much in line with the almost 50–50 splits that adults offer. This result is consistent with the literature demonstrating children's social understanding of relatively complicated economic considerations. While our study does not allow us to determine whether such equal division is innate or learned behavior, we would argue that the obvious selfishness of very young children makes it hard to argue that such behavior is innate. We also find that children respond appropriately to experience, during the experiment. In response to rejections of past proposals individuals increase future proposals, in contrast they decrease future proposals in response to past acceptances. These adjustments are very large in young children, and much smaller in older ones. This age effect seems entirely sensible, as younger children can be expected to have less experience and therefore less confidence in their prior beliefs about the likelihood that a given proposal will be rejected.

We find some evidence of social referencing. When shown information about what other proposers are doing, children tend to adjust their proposals in the direction of the proposal amounts that they see others making. Interestingly, the interaction effects in our regressions show that this tendency increases with age. This suggests that when making decisions with a social component, older children pay more attention to the example that others set. This argument for social referencing is strengthened by the fact that the proposal changes that are made do not lead to higher monetary payoffs, and it is consistent with other studies of children from developmental psychology. For example, [Steinberg and Silverberg \(1986\)](#), report survey data showing that as children enter adolescence they become more responsive to peer influences.

One interpretation of the result that changes in proposals do not increase earnings is that 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 their monetary payoffs from the experiment.

Our study shows that while the response to a child's own action causes them to change their strategies in a payoff-maximizing direction, feedback on other peoples' actions and the associated response has a very different effect. Rather than responding to whether or not proposals were rejected or accepted, children tend to respond to the proposals that others are making. During the limited-information treatment, proposers are only provided with information on their own monetary payoff, and have no information about other proposals and responses to those proposals. In this treatment, proposers use this information in a way that is consistent with reinforcement learning. During the course of the full-information treatments the participants also see information about what other proposers are doing, and whether their proposals are accepted or rejected. We find reinforcement learning as before, from the child's own experience. However children do not respond to information on whether the proposals of others were accepted or rejected, instead they adjust their proposals to mirror those of others. This suggests that the social learning that we are observing is driven by social referencing rather than reinforcement.

It could also be argued that, in the full-information treatment, players might assume that the decisions of the other proposers are rational and self-interested reactions to the rejections those players have experienced. Matching their proposals is therefore a way of incorporating the experience of other players, not just a desire for conformity. 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 personal bargaining experience, and conclude that children will refine their own strategies as they gain more direct bargaining experience. We also find evidence that children learn from the behavior of others, but only in a particular context. Proposers, particular older ones, tend to change their proposals in the direction of the proposals they observe others making. On the other hand, while proposers modified their proposals in a sensible way, according to the responses they received to prior proposals, they did not adapt to information about the responses that other proposers received. We argue that this result supports social referencing theory, but provides only weak support for social learning theory in bargaining situations.

Last, we note that the results of this study come from a controlled experiment with little of the sorts of social context that is present in the kind of bargaining that goes on between peers and within families. Our focus has been on showing that useful results on children's behavior can be obtained with this procedure. We think that future work using this basic experimental paradigm, but with more variation in social context, has the potential for producing interesting results on questions involving such bargaining. As one example, it would be reasonable to ask if children would behave differently if they were bargaining with other children with greater social distance, such as children in a different classroom, or even a different school.

Acknowledgements

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 Meindinger, Martin Dufwenberg, Andrew Shipley, and several anonymous referees for their comments on the paper.

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 (5th ed). In W. Damon (Ed.). *Handbook of child psychology* (Vol. 3, pp. 701–778). New York: John 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.
- Gregan-Paxton, J., & John, D. R. (1995). Are young children adaptive decision makers? A study of age differences in information search behavior. *Journal of Consumer Research*, 21(4), 567–580.
- Gurucharri, C., & Selman, R. L. (1982). The development of interpersonal understanding during childhood, preadolescence, and adolescence: A longitudinal follow-up study. *Child Development*, 53(4), 924–927.
- 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.
- John, D. R. (1999). Consumer socialization of children: A retrospective look at twenty-five years of research. *Journal of Consumer Research*, 26(3), 183–213.
- 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.
- Steinberg, L., & Silverberg, S. B. (1986). The vicissitudes of autonomy in early adolescence. *Child Development*, 57(4), 841–851.
- Webley, P. (2004). Children's understanding of economics. In M. B. E. Buchanan-Barrow (Ed.). *Children's understanding of economics*. London: Psychology Press.
- 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.