1

Cognition and Framing in Sequential Bargaining for Gains and Losses

Colin F. Camerer, Eric J. Johnson, Talia Rymon, and Sankar Sen

Introduction

Noncooperative game-theoretic models of sequential bargaining give an underpinning to cooperative solution concepts derived from axioms, and have proved useful in applications (see Osborne and Rubinstein 1990). But experimental studies of sequential bargaining with discounting have generally found systematic deviations between the offers people make and perfect equilibrium offers derived from backward induction (e.g., Ochs and Roth 1989).

We have extended this experimental literature in two ways. First, we used a novel software system to record the information subjects looked at while they bargained. Measuring patterns of information search helped us draw inferences about how people think, testing as directly as possible whether people use backward induction to compute offers. Second, we compared bargaining over gains that shrink over time (because of discounting) to equivalent bargaining over losses that expand over time.

In the games we studied, two players bargain by making a finite number of alternating offers. A unique subgame-perfect equilibrium can be computed by backward induction. The induction begins in the last period and works forward. Our experiments use a three-round game with a pie of \$5.00 and a 50-percent discount factor (so the pie shrinks to \$2.50 and \$1.25 in the second and third rounds). In the perfect equilibrium the first player offers the second player \$1.25 and keeps \$3.75.[1]

In previous experiments (including ours; see Johnson et al. 1991) subjects actually offer the second player something between the \$1.25 equilibrium and \$2.50, an equal split of the initial pie. Mean offers are around \$2.00. Lower offers, including equilibrium offers of \$1.25, are often rejected.

In our expanding-loss experiments subjects began with $5 for each round. Then the players bargained over division of a loss of $5.00 in the first round, a loss of $7.50 in the second round, and a loss of $8.75 in the third round. If they reached no agreement, each lost their $5.00 and the game was over. By adding the $5 stakes to the potential losses at each stage, it is easy to see that the shrinking-gain and expanding-loss games are equivalent if players maximize net wealth. Using backward induction, we derive the same equilibrium as in the shrinking gain game: player 1 offers player 2 a loss of $3.75 (a net gain of $1.25 when the stake is added) and accepts a loss of $1.25.

Comparing behavior in gain-and-loss games tests whether players bargain over final wealth positions, or are instead sensitive to changes in wealth (as suggested by "framing effects" in studies of individual choices, e.g., Kahneman and Tversky 1979). If players "segregate" their $5 stake from the losses they bargain over, and react differently to losses than to gains, then behavior in the two games might be different, even though the games have the same implications for net wealth.

Our work is part of a broader attempt to identify whether there are systematic deviations between actual behavior in games and behavior predicted by solution concepts. The hope is that the study of such deviations might throw light on the basic reasoning processes players use.[2] The presumption is that bounded rationality might cause players to reason differently than theorists do. (Our experiments also shed light on how players learn, if they do not reason game-theoretically.) The two-person alternating-offer bargaining we study is a fruitful setting for searching for deviations because the backward induction that underlies perfect equilibrium calculations is difficult.

We suspect there are three classes of deviations from game-theoretic reasoning: (1) players do not look forward and backward sufficiently; (2) players do not reason sophisticatedly about the choices of others; and (3) players violate expected utility maximization (much as they do in making individual choices). Our experiments test for deviations of the first and third type.

Explaining Observed Anomalies: Fairness vs. Learning

Observed departures from equilibrium bargaining have two possible sources. *Fairness* theories account for the departures by assuming that play is in equilibrium, but players have a preference for fair rules or a utility for fairness (e.g., Bolton 1991). (Rejections of offers above $1.25 among player

2s suggests a *dis*utility for *un*fairness, which might induce an apparent utility for fairness among player 1s.) *Learning* theories account for the observed departures by suggesting that players do not initially understand how the structure of the game conveys bargaining power. Confused, or naive, players resort to conventions of fairness, but experience can teach them how the game's structure creates bargaining power; then they will offer equilibrium divisions (with no concern for fairness).

There is evidence to support both views. In ultimatum games with one round, or dictator games in which player 1 dictates a division player 2 has to accept, players do *not* always offer small sums to others, as perfect equilibrium based on self-interest predicts; many offers are equal splits (e.g., Thaler 1988; Forsythe et al. 1988; but see Hoffman et al. 1991). These data suggest fairness plays some role.

Evidence consistent with the learning hypothesis comes from three experiments on sequential bargaining in which convergence to equilibrium offers occurred with suitable experience. In Binmore, Shaked, and Sutton 1985 experience taught player 2s to make equilibrium offers in a second bargaining trial, when they were placed in the role of player 1. In Harrison and McCabe 1992 offers converged to the perfect equilibrium in several repetitions of a three-stage game, when subjects also played the second-stage subgame between each repetition of the three-stage game. In Neelin, Sonnenschein, and Spiegel 1988 economics students, who probably had some exposure to game theory, chose the perfect equilibrium immediately in a two-stage game.

Our intention is not to resolve the debate over the two interpretations. Instead, we cast the learning hypothesis as an assertion that players do not reason game-theoretically at first, and we test that assertion directly. Our evidence does not resolve the debate because players who do *not* backward induct might be concerned about fairness, too. But if backward induction is a poor description of players' thinking, then any theory which claims the data are *equilibrium* offers, reflecting both fairness and a cognitive understanding based on backward induction, is falsified.

Framing Effects in Sequential Bargaining

An interesting question is whether subjects bargain in the same way over the division of shrinking gains and the division of expanding losses from a stake. Prospect theory (Kahneman and Tversky 1979) accounts for differences between gains and losses which result in the same net wealth, by assuming that people value outcomes compared to a reference point. Then outcomes can be expressed as positive or negative deviations, gains or losses, depending on the "frame" (or reference point) one adopts. Inducing gains and losses by allowing a reference point is different from the orthodox approach, in which people have utility over net wealth, in two ways: (1) people are assumed to have diminishing sensitivity to increasing changes and (2) they are assumed to be more sensitive to losses than to gains (or "loss-averse"). Diminishing sensitivity implies a value function that is concave for gains (corresponding to diminishing marginal utility) and convex for losses (diminishing marginal *dis*utility), reflecting in shape at the origin. Reflection implies that people will take risks when the possible outcomes are all losses.

Many applications of prospect theory have focused on the framing effects of gains and losses. For example, Bazerman (1983) compared the efficiency of bargains achieved by subjects when outcomes were formulated as gains or as losses. His subjects were more likely to disagree when bargaining over losses, consistent with the principle that subjects are more risk-seeking in the domain of losses. Similarly, if our subjects are more risk-seeking in the domain of losses, we expect to see more rejected offers in the expanding-loss game.

Weg and Zwick (1991) studied sequential bargaining over gains and losses when a *fixed* cost is imposed each period (rather than a percentage cost, due to discounting, as in our setup). They found no difference in bargaining over gains and losses. Furthermore, their results, and experiments by Rapoport, Weg, and Felsenthal (1990), suggest that behavior in fixed-cost games is much better predicted by perfect equilibrium than behavior in discounting games is.

Methodology

Subjects were Wharton undergraduates recruited from business or economics classes, or from general sign-up sheets posted around the campus. Ten students met in our lab at a certain time. The methods in the gain-and-loss sessions were the same, except subjects in loss sessions were given $60 at the start, from which losses were later subtracted. The $60 stake represents $5 in each of eight rounds, plus a $20 flat payment. We gave subjects the entire stake at the beginning of the experiment, and encouraged them to put it in their pockets, so they would be more likely to mentally segregate the stake from subsequent losses.

An experimenter read instructions aloud to make them common knowledge. Afterwards, subjects worked through several examples (balanced to

avoid biasing their responses) and took a quiz to ensure they understood the instructions.

Each group of ten subjects played eight three-round alternating-offer bargaining games, with a different anonymous opponent each time. This general design helps subjects to learn from "stationary replication" of the game, while avoiding the reputation effects that might arise if two subjects were playing each other repeatedly and knew it.

At the end of the sessions, subjects were paid half of the dollar amounts they actually earned from bargaining, in cash. In the loss sessions, subjects physically paid us back some of their initial stake.

Recording Information Search with MouseLab

The novelty of our experiments in this study, and in Johnson et al. 1991, is that we recorded the information search of the players.

Subjects were not told the pie sizes in each of the three stages. Instead, the pie sizes were hidden behind boxes on a computer screen, shown in figure 1.1. The computer screen has six boxes. Behind each box is the amount of the pie in a round (left-hand column boxes) or the role of the subject in a round[3] (right-hand column boxes), for each of the three rounds of the game. To see what is in a box, subjects use a mouse to move a cursor

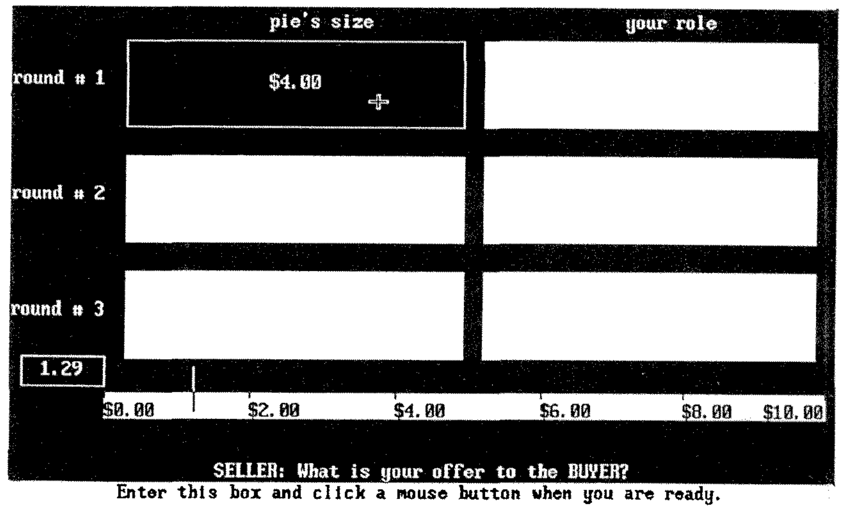

Figure 1.1
The MOUSELAB display of the shrinking-pie game subjects saw

into the box. Once the cursor enters the box, the box automatically opens, revealing the pie size. In figure 1.1, for example, the player opens the first box and sees that the first-round pie is a loss of $4.00.

The computer system we use, called MouseLab, records the location of the cursor every 60 milliseconds, giving a fine-grained measure of the time when each box is entered and exited. With these data we can study the order in which boxes were opened, and how long each box was open.

Measuring patterns of information search provides an indirect way of learning about cognitive processes that are not directly observable. Think of the brain as a factory that produces decisions from raw materials (information). If we cannot observe the production process in the factory, the next best thing is to observe the flow of raw materials into it (information into the brain): the order in which materials arrive, how long they are used in the production process, and so forth. A theory about how materials are assembled (or how information is combined to reach a decision) can be tested indirectly by observing the flow of materials.

Results

We conducted three sessions comparable to those reported in Johnson et al. 1991 with 10 subjects and 8 three-round games (or "trials"), for a total of 120 observations.

We will first discuss and compare the offers players made in the gain-and-loss games, and how often offers were rejected. We then will turn to the information search data revealed by MouseLab.

Offers, Rejections, and Counteroffers

Three important stylized facts[4] have emerged from alternating-offer bargaining experiments (e.g., Ochs and Roth 1989):

(1) Subjects do not choose equilibrium divisions. The average offer lies somewhere between equal split and equilibrium.

(2) In equilibrium no offers should be rejected, but some are. (About half of the equilibrium offers are rejected in gains, and most of them are rejected in losses.)

(3) *Most* counteroffers (about 80% in gains and 52% in losses) are "disadvantageous": they give less to the person making the offer than he or she previously rejected.

Our first concern is whether these stylized facts are replicated in our experiments. Figures 1.2a and 1.2b, for gains and losses respectively, show histograms of offers made in the first round (pooling across sessions and trials) and plots of rejected first round offers versus resulting second-round counteroffers. (To make the gain data comparable to the loss data, all offered losses were transformed into $5 minus the offer).

In the gain domain the average offer is $2.11. The equal-split point ($2.50) and equilibrium prediction ($1.25) are marked on the figures. Most offers are closer to the equal split than to the equilibrium, but only a few are within a dime of the equal split. The shaded portion of each bar indicates the number of rejected offers. Offers were rejected 10.8 percent of the time in gains, a rate comparable to the rejection rates found in prior experiments. Note that equilibrium offers (between $1.20 and $1.40) are rare, and are rejected about half the time.

In the loss domain the average offer was $2.22 (i.e., player 1 offered player 2 a loss of −$2.78, leaving her with a net gain of $2.22). Offers were rejected 22.5 percent of the time, and low offers were very frequently rejected. For example, gain offers of less than $1.80 were accepted about half the time (see figure 1.2a), but loss offers giving player 2 a net payment less than $1.80 were rejected all but once (see figure 1.2b). Offers were also much more dispersed in the expanding-loss game than in the shrinking-gain game. In the loss domain there are several offers *above* the equal split (indicating player 1's offer to accept more than half the initial loss) and *below* the perfect equilibrium.

First-round rejections result in a second round of play. The right panel of figure 1.2a shows first-round offers rejected by player 2 (vertical axis) plotted against counteroffers player 2 proposed for herself in the second round (horizontal axis). A point above the diagonal line indicates a case where player 2 rejected an offer (e.g., $1.80) then proposed a division of the second-round pie which gave her even less ($1.25). Figure 1.2a shows that 85 percent of second-round offers are disadvantageous, falling above the diagonal line. This striking frequency of disadvantageous counteroffers is close to the fraction observed in earlier studies. However, in the loss domain only 52 percent of the counteroffers are disadvantageous.

Dark circles in figures 1.2a and 1.2b indicate second-round offers which were rejected. In the gain domain 23 percent (3/14) of second-round offers are rejected; in the loss domain, 37 percent (10/27) are rejected. (Two of the three subsequent third-round offers were rejected in the gain domain, and eight of ten in the loss domain.)

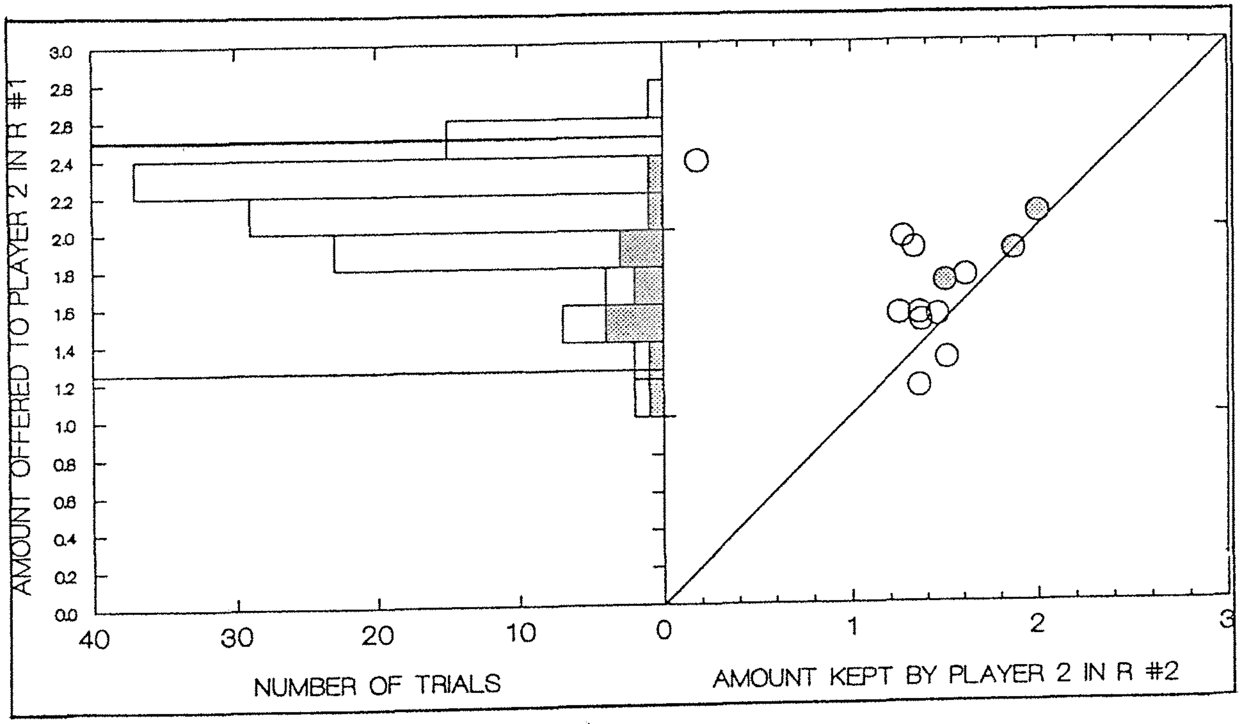

Figure 1.2a
Histogram of first-round offers and scatterplot of offers vs. counteroffers, shrinking-gain game

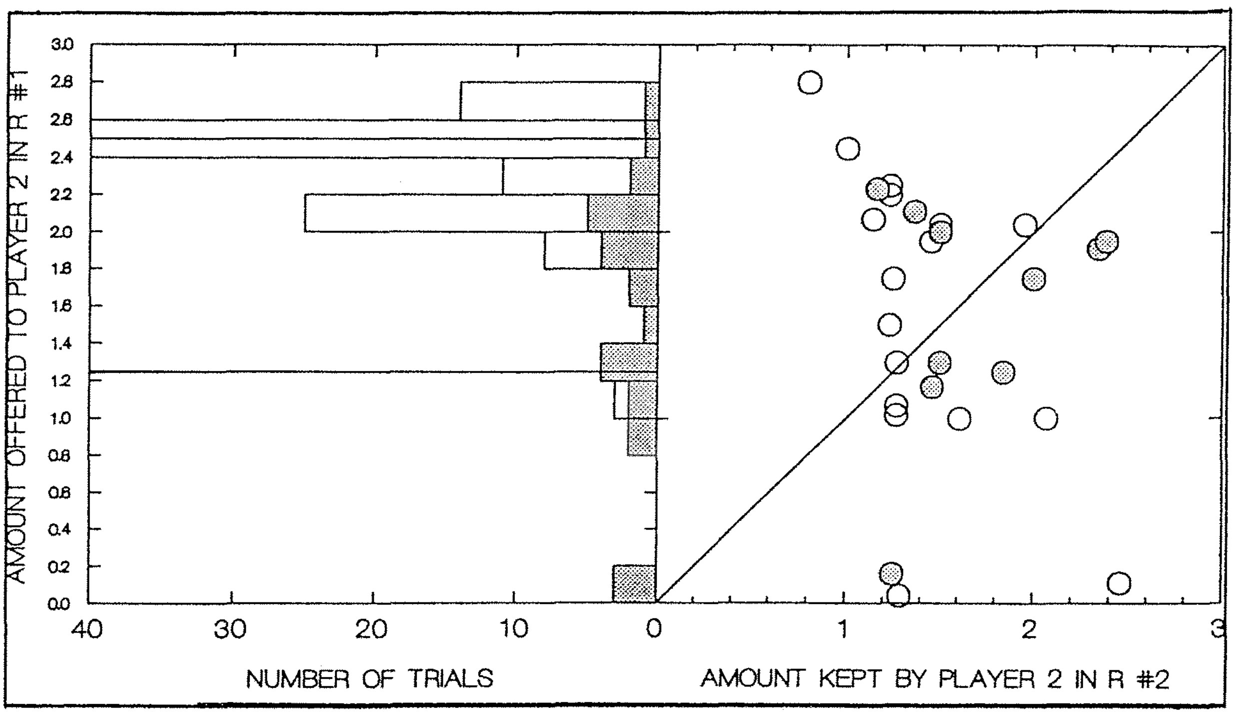

Figure 1.2b
Histogram of first-round offers and scatterplot of offers vs. counteroffers, expanding-loss game

These results roughly replicate the stylized facts discovered in earlier experiments: Mean offers lie between equal split and perfect equilibrium; a substantial minority of offers (10–20%) are rejected; and most counteroffers are disadvantageous. However, in the expanding-loss games offers are more dispersed, more offers are rejected, and fewer counteroffers are disadvantageous.

Information Search: Theory

We examine three measures of information search: (1) *number of acquisitions* per box (number of times each box is opened in a period); (2) *total time examining payoff*, or looking time, per box (amount of clock time each box is open in a period[5]); and (3) *number of transitions* from one specific box to another.

Before presenting data, it is useful to ask: If subjects were using backward induction to compute perfect equilibrium offers, what would their pattern of information search look like?

One possible answer is that equilibrium analysis predicts only *outcomes* (equilibrium offers); it says nothing about the *process* by which outcomes are derived. We think this answer is unproductive and wrong[6] If we test only the theory's outcome predictions, it is clearly rejected. We turn to process predictions as a way to see where the theory fails. If we are not allowed to specify the theory as a process and test it with information search data, we are left with a rejected theory and no obvious way of figuring out why it failed.

Players using backward induction to calculate equilibrium offers might search for information as follows:

(1) First look at the third-round payoff, and figure out the equilibrium offer in that round. (This calculation is simplest.)

(2) Then look at the second-round payoff, and figure out the equilibrium offer in that round. This calculation will take more time, and may require subjects to glance back and forth between the second- to the third-round boxes.

(3) Finally, look at the first-round payoff.[7]

Thus, we take equilibrium analysis to predict: transitions from box to box will be predominantly backward transitions, from the nth-round box to the $(n-1)$th-round box; the longest looking time will be in the second box; looking times in the first- and third-round boxes will be shorter.

The reader may object to our characterization of the information search process that is inherent in equilibrium analysis. We are eager to hear alternative characterizations. We also derived an *empirical* characterization of equilibrium information search by training a group of subjects in backward induction (see Johnson et al. 1991) and rewarding them for computing the perfect equilibrium offer. Those subjects looked mostly at the second and third boxes, and moved backward (from future round boxes to current ones) more often than forward.

There is no obvious prediction from game theory about how bargaining over losses and gains will differ, either in offers or in search for information. However, the principle of loss-aversion (losses are more painful than equivalent-sized gains are pleasurable) suggests that people might sacrifice more time and effort—they will work harder—to avoid losses in bargaining than to reap gains. For example, Maule (1989) found that people who were asked to speak aloud while making a choice spoke more when making choices involving losses rather than gains. Similarly, we suspect our subjects might spend more time processing information revealed by boxes, opening boxes, and making transitions between boxes, when bargaining over losses instead of gains.

Information Search before Player 1's First-Round Offer

We start our analysis of information search patterns by studying information search by player 1 subjects in the first round of each game (pooled across all eight trials of all three experimental sessions). Table 1.1 presents these data for the loss domain (for the gain domain data see Johnson et al. 1991).

In both domains most of the looking time (12.91 seconds out of 20.82 total seconds in gains, and 24.17 seconds out of 32.97 total seconds in losses) is spent looking at the first-round pie size. Half as much time in gain, and less than a fourth in loss, is spent looking at the second round. Only 1

Table 1.1
Information search measures, player 1's first-round offer (all trials), expanding-loss game

Round (loss)	Number of Acquisitions	Total Time Examining Payoff	Transitions 1	Transitions 2	Transitions 3
Round 1 ($5.00)	4.10	24.17	—	1.75	0.64
Round 2 ($7.50)	2.85	5.40	1.20	—	1.20
Round 3 ($8.25)	2.25	3.42	0.64	0.68	—

second in gains and 3 seconds in losses are spent looking at the third-round pie size. In fact, in 10 percent of the trial in gain and 14 percent in loss the players did not even open the third-round box (although they opened the first-round box in *every* trial). Without opening each box, players cannot possibly be calculating equilibrium offers by using backward induction.

Notice that each box is opened about 2–4 times each trial in both domains. These data suggest subjects are opening and reopening boxes frequently, rather than memorizing the numbers in the boxes. The pattern of transitions between boxes is shown in the last three columns of table 1.1. Entries show the average number of transitions from the row box to the column box. (For example, players moved from the round 3 box to the round 2 box an average of .68 times per trial.) Contrary to the backward induction prediction, there are always more forward transitions (above the diagonal) than backward ones (below the diagonal).

The pattern of search evident in table 1.1 does not conform to our characterization of equilibrium search. Most looking time is concentrated on the first-round pie, with decreasing attention paid to second- and third-round pies, and subjects make forward transitions rather than backward ones. Figure 1.3 is an icon graph which displays the information processing measures given in table 1.1 (marked LOSS), and corresponding measures for the shrinking-gain game results (marked GAIN) reported in Johnson et al. 1991. Each box corresponds to the approximate position of the three payoff boxes on the MOUSELAB computer display. The width of each box is proportional to the number of acquisitions of that box (the second column in table 1.1). The height of the shaded area in each box is proportional to the amount of time spent looking at that box (the third column in table 1.1). The shaded areas are standardized so that the box which is open longest (in this case, the first-round payoff in the loss domain) is completely filled. (Horizontal lines show the midpoints of each box to make comparisons easy.)

The arrows represent the number of transitions between boxes (the three right-hand columns in table 1.1). The thickness of each arrow is proportional to its frequency. To simplify the display, we left out arrows marking transitions that occurred less than once a trial, on average.

The icon graphs express visually what table 1.1 shows numerically: in both domains, subjects open the first and second boxes most often, and look at the first box longest. They barely glance at the third box (although in loss they open it more often, for more time, than in gain). They move back and forth between the first and second boxes relatively often, and

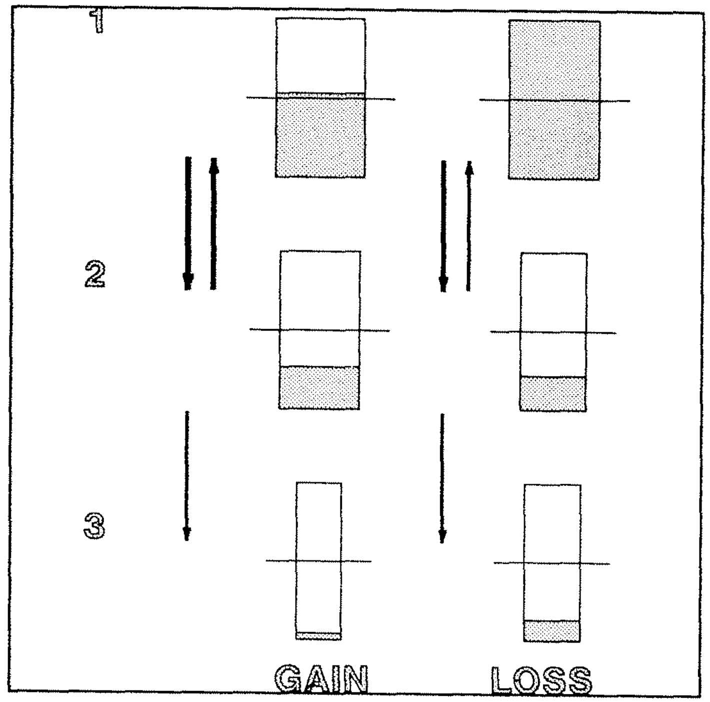

Figure 1.3
Icon graphs of player 1 first-round information search measures

more in gain than in loss. In terms of total looking time and number of box openings, subjects do appear to work harder when bargaining over losses than when bargaining over gains.

The Relation between Information Search and First-Round Offers

Table 1.1 and figure 1.3 show that in both domains, subjects appear to search for information in ways that are inconsistent with backward induction. It is also useful to look for differences in information processing, to see if they are correlated with differences in offers.

Based on first-round offers to player 2s, we divided trials into three groups—"near-perfect" (<\$2.01, $n = 23$), middle (between \$2.01 and \$2.39, $n = 36$), and equal-split (\$2.40 or above, $n = 61$). Figure 1.4 shows an icon graph of information processing measures for trials which fall into each of the three groups. Information processing is similar in each of the

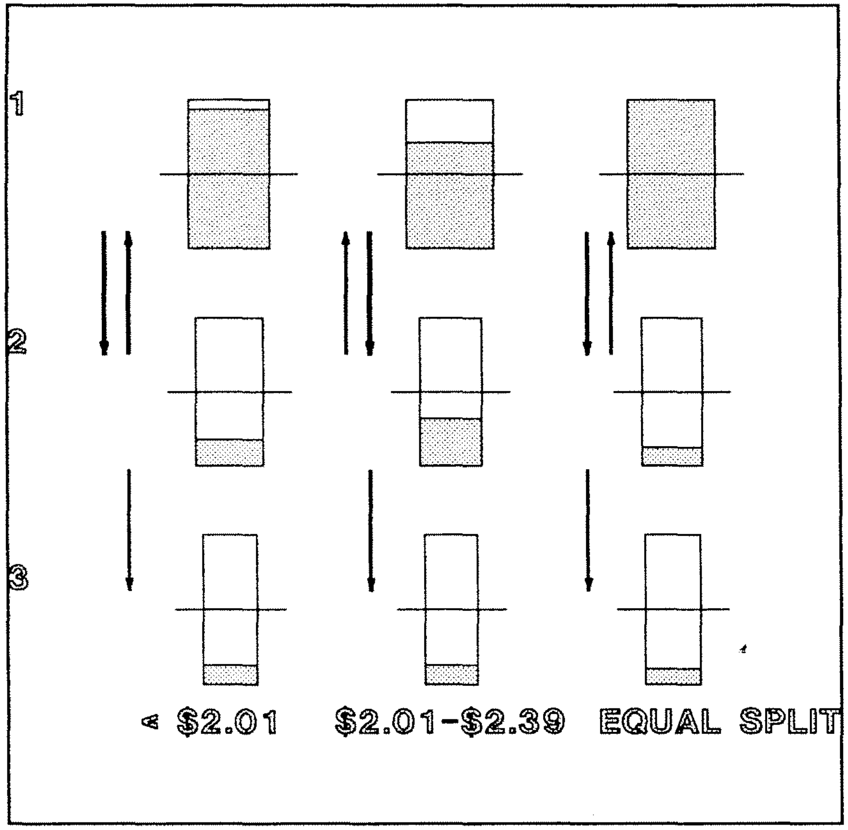

Figure 1.4
Player 1 information search measures classified by first-round offer

three groups. By contrast, in bargaining over gains, subjects who looked further ahead in the game also tended to make lower "near-perfect" offers (see Johnson et al. 1991).

Information Research by Player 2

So far we have concentrated on first-round offers made by player 1. What about information search by player 2s as they consider the offers? (Note that player 2s were not allowed to open boxes until they got a specific offer from player 1.) Figure 1.5 shows icon graphs of information search by player 2s, classified by whether offers were rejected and by the size of offers that were accepted. Generally, search patterns by player 2s look

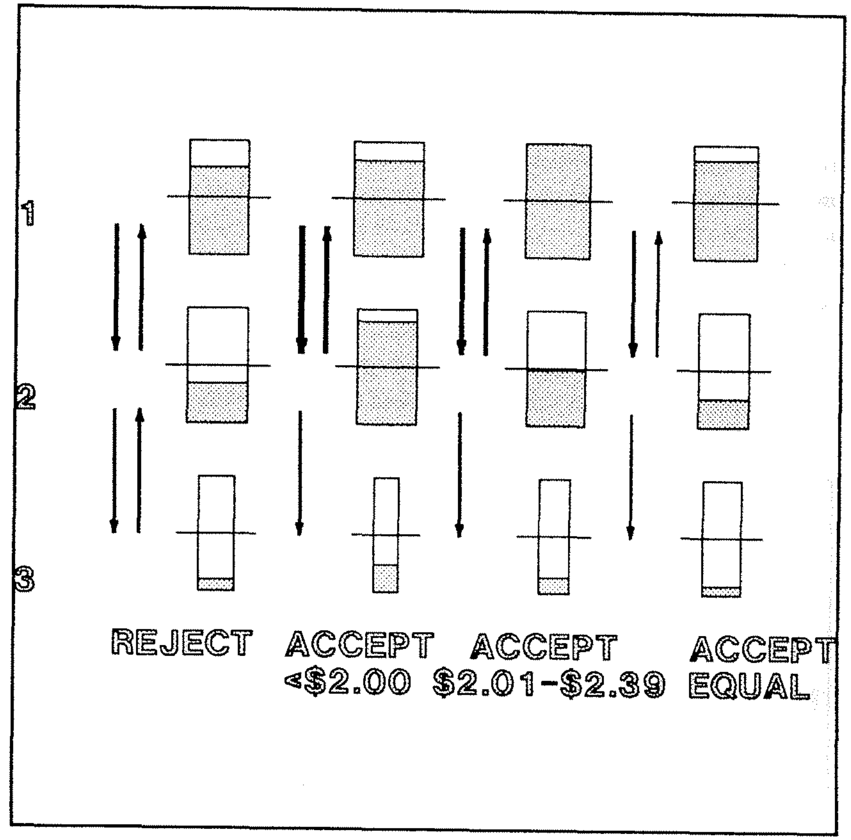

Figure 1.5
Player 2 first-round information search measures classified by rejection or acceptance and first-round offer

much like the patterns for player 1: search is concentrated on the first-round payoff, the third-round box is rarely opened, and transitions tend to be forward (from earlier- to later-round boxes) rather than backward. The most striking difference is between the first column (representing processing by players who reject offers, which tend to be below $2.00) and the second column (players who accept roughly the same offers, below $2.00). Compared to players who reject low offers, players who accept low offers look about twice as long at the second pie size, and make more transitions between the first and second boxes (shown by thicker arrows in the second column). Subjects who accept low offers appear to think through the consequences of rejecting an offer better than rejecting subjects do. A similar difference is evident in bargaining over gains (Johnson et al. 1991).

Learning and Experience

In our experiments each subject participated in eight trials, four as player 1 and four as player 2. Learning might be manifested by changes in information processing over the trials. For example, if subjects learn backward induction by sheer repetition of bargaining, then there should be a shift in information processing from round one to future rounds, across trials in an experiment. Figure 1.6 shows icon graphs of information processing measures over the four trials in which a subject was player 1. There appears to be no substantial effect of repetition on information processing.

Simply repeating the same game does not seem to teach subjects backward induction. But learning might occur when subjects have an offer

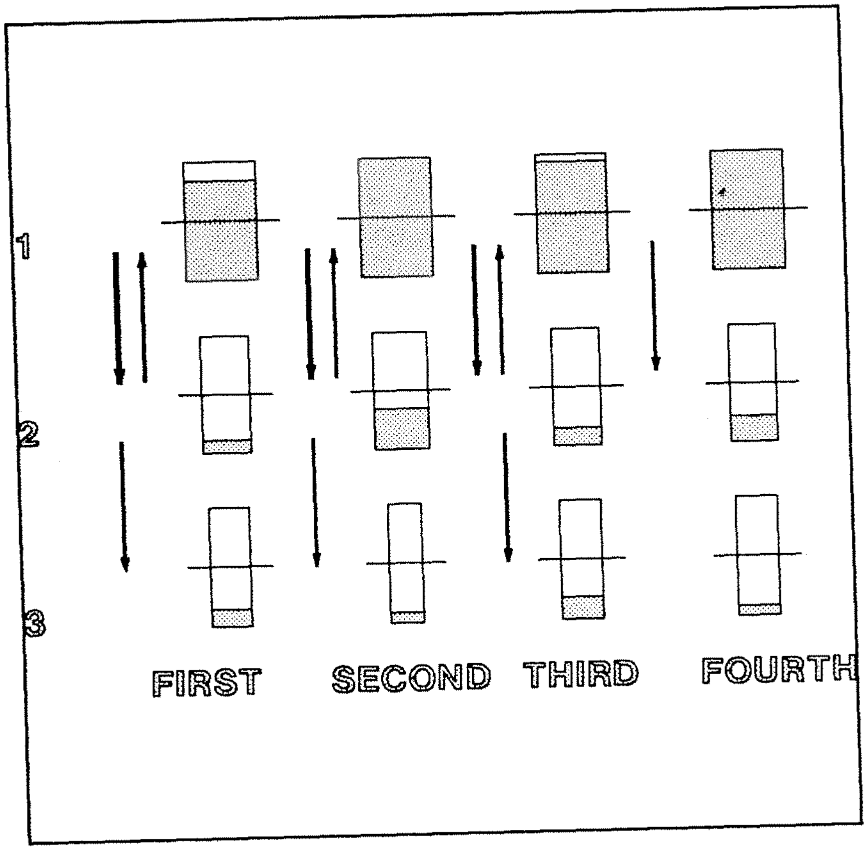

Figure 1.6
Player 1 first-round information search measures across four trials

rejected, and bargaining proceeds to the second round. Then subjects directly experience the consequences of a first-round rejection, which might affect the way they think about their subsequent first-round offer. (Harrison and McCabe's (1992) design mimics this kind of learning by forcing subjects to play subgames between each trial of three-round bargaining.)

To test the effect of a first-round rejection on subsequent play, we compare measures of information processing by a player 1 subject at four points in time: (1) the trial in which a first-round offer was rejected ("1st reject"); (2) the second round of that same trial ("next round"); (3) the "next trial"; and (4) the next trial in which the subject had the role of player 1 again ("next as 1"). Figure 1.7 shows an icon graph comparing information processing measures at these four points. The only striking difference is

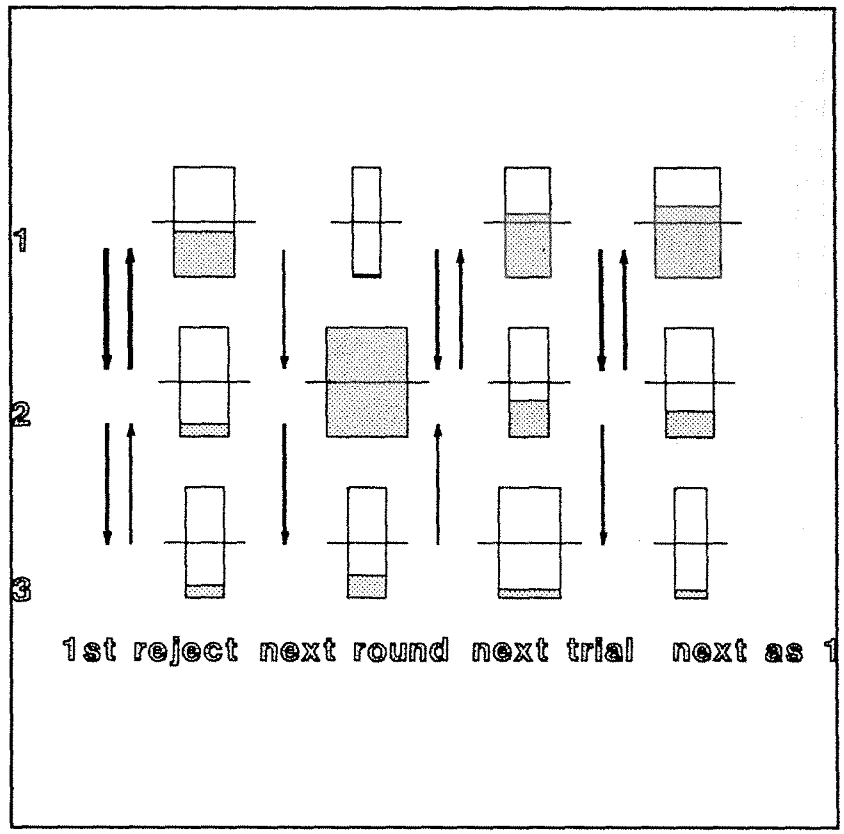

Figure 1.7
Player 1 information search measures classified by four kinds of experience

that in the second round after a first-round rejection (marked "next round"), subjects look much more frequently at the second-round pie size, as one would expect. Experiencing a rejection in the first-round of one trial does *not* cause players to look further ahead in the first-round of the next trial.

Implications and Conclusions

We studied offers and patterns of information search in three-round alternating-offer bargaining over expanding *losses*, and compared the results to a shrinking-*gain* game studied by Johnson et al. (1991). As in most other experiments on bargaining with discounting, offers were scattered between an equal split of the first-round pie and the perfect equilibrium offer (derived from backward induction). About 10 percent of the offers were rejected in gains and 23 percent were rejected in losses.

Patterns of information search showed that most subjects did not look at the pie sizes in the correct order, and for the length of time, necessary for backward induction. Instead, subjects concentrated on the current round when making decisions. In the first round they looked mostly at the first-round pie size and looked only briefly at the third-round pie. Subjects who looked ahead were no more likely to make offers closer to the equilibrium prediction, but they were more likely to accept low offers. Repeating the game, or having a first-round offer rejected, had no discernible impact on information processing.

Since subjects do not appear to be using backward induction to compute offers, it would be helpful to know what they are thinking. We think that subjects try to approximate the profit-maximizing offer, and their attention is guided by (at least) two heuristics: the *size* heuristic directs attention to the largest-sized pies; and the *distance* heuristic directs attention to current-round pies, leading people to ignore distant rounds that are unlikely to occur. A theory that weaves these heuristics into a cognitive account of how subjects choose offers in sequential bargaining would be an extremely useful alternative to the perfect equilibrium theory.

While the gain-and-loss conditions were designed to be identical in terms of subjects' final wealth, there were several notable differences in behavior. In the loss condition subjects opened the boxes more frequently, and looked at the pie sizes about twice as long. Although the average offer was similar for gains and losses, offers were more widely dispersed in the loss condition and were rejected twice as often (see Bazerman 1983).

Recall that two features have been proposed to explain anomalies in sequential bargaining: fairness and learning. Some fairness-based models suggest that people play sophisticatedly but care about their opponent's payoff (e.g., Bolton 1991 in a two-round game). Our information-processing evidence suggests that *any* model that assumes players backward induct is a poor descriptive account of how people think. Such a model might make reasonable predictions, but models that incorporate better representations of human cognition could make even better predictions.

Our data support the initial premise of the learning-based approaches—that players do not begin by reasoning game-theoretically—but we found little evidence that any learning actually occurred. Either learning requires more sheer experience than the eight trials our subjects engaged in, or the experience must be of a different sort. The effect of different amounts and kinds of experience on learning is an important empirical question that our system is well-equipped to answer.

Notes

The financial support of the National Science Foundation to the authors Camerer and Johnson (grants NSF 88-09299 and NSF 90-23531) is gratefully acknowledged. Helpful comments were received from audiences at Harvard, Penn, Cornell, MIT, NYU, Penn State, the Bonn Mathematical Economics Workshop (BoWo IV), and the International Conference on Game Theory (Florence, June 1991).

1. The perfect equilibrium is calculated as follows (assuming players reject offers they are indifferent toward and the smallest unit of currency is $.01). In the third round, an ultimatum game, player 1 offers $.01 and keeps $1.24. Anticipating that, player 2 offers $1.25 in round 2 and keeps $1.25. So player 1 offers $1.26 in round 1 and keeps $3.74. If players accept offers they are indifferent toward, there are a small band of perfect equilibria from $1.24–$1.27. We refer to "the" equilibrium at $1.25 for simplicity.

2. Our focus on deviations as a source of understanding about typical reasoning is not unusual in science. Researchers often study unusual departures from everyday patterns—the Great Depression, supernovae, volcanic eruptions—in order to learn more about those patterns.

3. In these experiments subjects alternated roles across rounds, but the software was designed to be flexible enough to study nonalternating-offer bargaining, too.

4. Of course, there are many qualifications to the stylized facts. The most important is that learning may create convergence to equilibrium, reducing the number of rejections and disadvantageous counteroffers (as in Binmore, Shaked, and Sutton 1985; and Harrison and McCabe 1992). But the empirical evidence of learning is mixed.

5. To distinguish players actually examining information from those who opened boxes accidentally (usually while moving from one box to another), we filtered out all information acquisitions lasting less than .18 seconds. Psychologists have found

that people do not accurately perceive anything they see that briefly (Card, Moran, and Newell 1983).

6. Data presented in Johnson et al. 1991 show that differences in processing *were* correlated with outcomes in bargaining over gains. Therefore, the idea that offers are independent of information search patterns, so that equilibrium offers could result from any search pattern, is wrong empirically.

7. Of course, to calculate an equilibrium offer to player 2, player 1 does not actually need to look at the first-round payoff at all, as long as she knows the first round pie is bigger (smaller for losses) than the second-round pie.

References

Bazerman, Max H. 1983. "Negotiator Judgment." *American Behavioral Scientist* 27:211–228.

Binmore, Ken, Avner Shaked, and John Sutton. 1985. "Testing Noncooperative Bargaining Theory: A Preliminary Study." *American Economic Review* 75:1178–1180.

Bolton, Gary E. 1991. "A Comparative Model of Bargaining: Theory and Evidence." *American Economic Review* 81:1096–1136.

Card, Stuart K., Thomas P. Moran, and Alan Newell. 1983. *The Psychology of Human-computer Interaction*. Hillsdale, N.J.: Erlbaum.

Forsythe, Robert, Joel Horowitz, N. Savin, and Martin Sefton. 1988. "Replicability, Fairness, and Pay in Experiments with Simple Bargaining Games." Working paper 88-30, University of Iowa.

Harrison, Glenn, and Kevin McCabe. 1992. "Testing Bargaining Theory in Experiments." In R. M. Isaac, ed., *Research in Experimental Economics*. Vol. 5. Greenwich, Conn.: JAI Press.

Hoffman, Elizabeth, Kevin McCabe, Keith Shachat, and Vernon L. Smith. 1991. "Preferences, Property Rights, and Anonymity in Bargaining Games." Working Paper. University of Arizona. Department of Economics.

Johnson, Eric J., Colin Camerer, Sankar Sen, and Talia Rymon. 1991. "Behavior and Cognition in Sequential Bargaining." Working paper, University of Pennsylvania Department of Marketing.

Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47:263–291.

Kennan, John, and Robert Wilson. 1990. "Can Strategic Bargaining Models Explain Collective Bargaining Data?" *American Economic Review* 80:405–409.

Maule, A. J. 1989. "Positive and Negative Decision Frames: A Verbal Protocol Analysis of the Asian Disease Problem of Tversky and Kahneman." In H. Montgomery and O. Svenson, eds., *Process and Structure in Human Decision-making*. New York: Wiley.

Neelin, Janet, Hugo Sonnenschein, and Matthew Spiegel. 1988. "A Further Test of Noncooperative Bargaining Theory: Comment." *American Economic Review* 78: 824–836.

Ochs, Jack, and Alvin E. Roth. 1989. "An Experimental Study of Sequential Bargaining." *American Economic Review* 79:355–384.

Osborne, Martin J., and Ariel Rubinstein. 1990. *Bargaining and Markets*. San Diego: Academic Press.

Rapoport, Amnon, Eythan Weg, and Daniel S. Felsenthal. 1990. "Effects of Fixed Costs in Two-person Sequential Bargaining." *Theory and Decision* 28:47–71.

Rubinstein, Ariel. 1982. "Perfect Equilibrium in a Bargaining Model." *Econometrica* 50:97–109.

Spiegel, Matthew, Janet Currie, Hugo Sonnenschein, and A. Sen. 1990. "Fairness and Strategic Behavior in Two-person, Alternating-offer Games: Results From Bargaining Experiments." Working paper, Columbia University.

Stahl, Ingolf. 1972. *Bargaining Theory*. Stockholm: Stockholm School of Economics.

Thaler, Richard H. 1988. "Anomalies: The Ultimatum Game." *Journal of Economic Perspectives* 2:195–206.

Thaler, Richard H., and Eric J. Johnson. 1990. "Gambling with the House Money and Trying to Break Even: The Effect of Prior Outcome on Risky Choice." *Management Science* 36:643–660.

Weg, Eythan, and Rami Zwick. 1991. "On the Robustness of Perfect Equilibrium in Fixed Cost Sequential Bargaining under an Isomorphic Transformation." *Economics Letters* 36:21–24.