Detecting Failures of Backward Induction:
Monitoring Information Search in Sequential Bargaining[*]

Eric J. Johnson

Columbia School of Business
Columbia University
New York NY 10027

Colin Camerer

Division of the Humanities and Social Sciences
California Institute of Technology
Pasadena, CA 91125

Sankar Sen

Fox School of Business and Management, Temple University
Philadelphia, PA 19122

Talia Rymon

Arison School of Business, The Interdisciplinary Center,
Herzliya, Israel

Running Title: Information Search in Sequential Bargaining

Correspondence to: Eric J. Johnson, School of Business, Columbia University, 3022 Broadway, NY, NY 10027, phone: 212-854-5068, email: ejj3@columbia.edu.

Abstract

We did experiments in a three-round bargaining game where the (perfect) equilibrium offer was $1.25 and an equal split was $2.50. The average offer was $2.11. Patterns of information search (measured with a computerized information display) show limited lookahead rather than backward induction. Equilibrium theories which adjust for social utilities (reflecting inequality-aversion or reciprocity) cannot explain the results because they predict subjects will make equilibrium offers to "robot" players, but offers to robots are only a little lower. When trained subjects (who quickly learned to do backward induction) bargained with untrained subjects, offers ended up halfway between equilibrium and $2.11.

Keywords: bargaining, experimental economics, bounded rationality, behavioral economics, behavioral game theory, fairness, limited cognition

JEL classification: C7,C9

1. Introduction

Game-theoretic models of sequential bargaining, in which agents alternate offers and counteroffers, have proved useful in economic theorizing in several ways (e.g., Ståhl [6]; Rubinstein [58]). The models serve as building blocks for theories of imperfectly competitive markets, such as "matching markets" with small numbers of buyers and sellers (see Osborne and Rubinstein [50]). Sequential bargaining models with incomplete information have been useful in explaining facts about labor strikes (Kennan and Wilson [40]). Alternating-offer models also give a procedural underpinning to axiomatic cooperative solution concepts such as the Nash bargaining solution (Binmore, Rubinstein and Wolinsky [4]), unifying different styles of analysis.

While alternating-offer models are analytically useful and realistic, direct experimental tests have typically rejected the equilibrium predictions they make. Offers are closer to equal divisions of an initial "pie size" than to equilibrium predictions, and offers do not change with parameter changes as sharply as theory predicts (Camerer [12] chapter 3). There is now an unresolved debate about whether offers that are too close to equal splits are due to stable preferences for fair division or to players' inabilities to compute equilibrium offers, or some combination.

In our game two players, 1 and 2, bargain over a pie that shrinks in value[1] in each of three periods (reflecting discounting due to impatience). The pie is worth $5 in the first period, $2.50 in the second period, $1.25 in the third period, and nothing after that. Starting with player 1, the players alternate making offers, which the responding player can accept or reject. If an offer is rejected, the pie shrinks and the player who rejected it then makes a counteroffer.

---

1 In Camerer et al. [14] we studied an expanding-loss game which gives the same wealth outcomes as the shrinking-gain games studied in this paper. Offers were more dispersed and rejected more frequently, and players spent twice as long looking at pie sizes. This difference between shrinking-gain and expanding-loss games is consistent with gain-loss distinctions or "framing effects" reported in many choice experiments.

If players are purely self-interested (and believe others are too[2]), perfect equilibrium

divisions can be derived by backward induction, starting with the third-period division and working

backward. The perfect equilibrium is (approximately[3]) for player 1 to offer $1.25 and keep the rest

for himself.

In earlier experiments, subjects typically offered something between the $1.25 equilibrium

and $2.50 (which is an equal split of the first-round pie). Offers average around $2.00. Lower

offers, including offers near the equilibrium prediction of $1.25, are frequently rejected.

1.1 Two theories: Limited cognition and social preference

We compare two general explanations of why these deviations from perfect equilibrium

occur, which we label "limited cognition" (or bounded rationality) and "equilibrium social

preference".

Limited cognition theories contend that players do not intuitively reason game-theoretically;

looking ahead is computationally expensive and unnatural because it means reasoning about events

that probably will not occur. But players who do not look ahead will not initially understand how the

structure of the game conveys bargaining power. In this view, players offer equal splits or a little less

because they aren't sure what else to do or aren't willing to bet too much on the self-interest and

---

2 Of course, the perfect equilibrium prediction in an n-stage game requires n levels of iterated belief in self-interest. For example, in our game player 1 offers $1.25 only if she is self-interested, believes 2 is (which guarantees a minimal ultimatum offer in the third round), and believes 2 believes she (1) is (which guarantees a $1.25 offer to 1 in round 2, leading player 1 to offer $1.26 in round 1),

3 In the third round, player 1 offers $.01 or nothing (depending on whether player 2 will accept nothing) and earns $1.24-$1.25. In the second round, player 2 must offer $1.24-$1.26 to player 1, leaving $1.24-$1.26 to herself. In the first round, player 1 must therefore offer $1.24-$1.27 to player 2. For simplicity, we refer to this range of perfect equilibria as "the" equilibrium at $1.25, recognizing that the equilibrium is not unique, and is subgame perfect rather than only Nash (every offer is consistent with Nash). Note that a short cut to calculating equilibria in games of this form is that the first player earns the sum of the pie sizes in odd-numbered rounds ($5.00 + $1.25), minus the pies values in even-numbered rounds (- $2.50). Alternatively, the first player should offer the third pie minus the

rationality of others. Fair offers may therefore be only temporary-- fairness is a heuristic for generating an offer, not an expression of stable preference-- and may be displaced by strategic behavior as suitable experience accumulates. There is limited evidence consistent with this view from experiments in which subjects get special experience with future subgames (that are otherwise unlikely) and learn to make equilibrium offers (Binmore, Shaked and Sutton [5]; Harrison and McCabe [33][4]; but cf. Carpenter [17]). Given this special experience, subjects' offers converge away from equal splits and toward the subgame perfect equilibrium, although most other experiments (including ours) show little learning of this sort (Ochs and Roth [49]; Spiegel et al. [65]).

To be precise, we define limited cognition as substituting an empirical assumption about others' behavior for the more extensive computation of equilibrium. Specifically, we believe that some subjects (1) are not examining pie sizes in (some) future stages; and (2) assume that other players will behave as they do in an ultimatum game (i.e., making and accepting 40% offers in ultimatum games). This is not a complete theory of limited cognition but it will suffice to organize what we see in three-period games (cf. Costa-Gomes et al. [19]; and Gabaix and Laibson [27], for more ambitious efforts). Following Harless and Camerer [30], El-Gamal and Grether [23] and Costa-Gomes et al. [19], we posit a discrete set of decision rules, identify how often they are used, and test whether classifying subjects by decision rules to predict the offers they make. Like Costa-Gomes et

---

second pie to the second player, which does not even require her to know the first-period pie size.
4 Binmore, Shaked and Sutton [5] studied two-stage bargaining with pie sizes of 100 and 25. The equilibrium offer is 25, and players initially offered between 25 and 50. When player 2's were asked to take the player 1 role, in a (hypothetical) repetition of the two-round game, they offered only 25. Experience in the player 2 role seemed to teach subjects how much bargaining power player 1 has, and they were eager to exploit it in future games. Harrison and McCabe [33] had subjects play a series of three-round games with pie sizes $5, $2.50, $1.25 (as in our initial example). Between each three-round game subjects played a two-round game with pie sizes $2.50 and $1.25. Notice that the two-round game is simply a subgame of the three-round game (corresponding to its second and third rounds). With experience, subjects learned from playing the two-round game what player 2 would get at the second round of the three-round game (around $1.25). Then they gradually converged near to the equilibrium division in the first round of the three-round game. However, Carpenter [17] did not find the same convergence, using different pie

al [19] we use attentional data to validate these strategies. One decision rule is level-0—look <u>no</u> steps ahead, and offer 40% (or $2.00) (cf. Ståhl and Wilson [64], and see Camerer [12] chapter 5). Another rule is level-1— look 1 step ahead, assume the other player will make a 40% offer in the second round (which leaves her $1.50), and backward induct from there, leading to an offer of $1.50.[5] A third rule is perfect equilibrium, which requires players to look at all the pies in all three rounds.

Equilibrium social preference theories account for departures from perfect equilibrium by positing a "social utility" for others' payoffs or differences in payoffs (Loewenstein, Thompson and Bazerman [41]; Fehr and Schmidt [25]; Bolton and Oeckenfels [9]; Charness and Rabin [18]; Camerer [12] chapter 3), a preference for reciprocating fairness and unfairness (Rabin [54]), or an unobserved component of payoffs which appears like noise or a mistake to an outside observer (McKelvey and Palfrey [44]; Goeree and Holt [28]; Weiszacker [67]). The clearest evidence of social preference comes from dictators and ultimatum games (see Camerer and Thaler [15]). In ultimatum games, offers are usually much more generous than predicted by equilibrium (40% on average) and low offers are often rejected (20% offers are rejected half the time), though results vary interestingly across cultures (see Henrich et al. [34]). According to social preference theories, these offers are consistent with an equilibrium in which players make offers knowing that other players may reject allocations which appear unfair, in a way the theories make precise.[6] Social preference

---

sizes.

5 These limited look-ahead rules are similar to rules used in artificially-intelligent game-playing programs. Such programs usually look a small number of steps ahead to avoid combinatorial explosion in ``bushy trees", and substitute an "evaluation function" at each truncated node for a formal analysis of the value function at that node. Our evaluation function is based on the regularity of 40% splits in these games.

6 For example, in Fehr and Schmidt's [25] theory, the responding player has a utility of $x-\alpha(S-x-x)$ for a stake of S and an offer of $x<S/2$. The weight α represents the strength of disutility for envy (the difference between the responder's payoff x and the proposer's share $S-x$) or disadvantageous inequality. Responding players will reject an offer x if they have $\alpha>x/(S-2x)$. Denoting the cdf of α in the population of responding players by $F(\alpha)$, the chance of

theories have been successful at finding a single social utility function, parameterized by one or two

free parameters (and permitting individual differences), which can explain a broad range of data

from many different games, including apparent reciprocation in trust games (e.g., Berg, Dickhaut

and McCabe [2]), contributions and punishment of noncontributors in public goods games (Fehr and

Gachter [24]), and the willingness of proposers to give away virtually everything, and responders to

accept anything, in the face of competition (Roth et al. [57]). Equilibrium social preference theories

all assume players derive utility from payoff <u>vectors</u> across players, or vectors and choices (rather

than just to their own payoff), and then do standard kinds of game-theoretic reasoning.

Our paper compares the limited cognition and equilibrium social preference explanations in

three experimental studies. (Binmore et al. [3], also study backward induction and related principles,

experimentally, but less directly than we do.) In the first baseline study, players bargain against one

another in three-period alternating-offer bargaining (using the parameter values in the example

above). Offers average \$2.11, and offers below \$1.80 are rejected half the time (replicating earlier

results). We also collected measures of information search with a "Mouselab" system used

extensively to study individual decision-making (Payne, Bettman and Johnson [52]). The

information measures show that players do not tend to look at future rounds and reason backwards,

as backward induction implicitly assumes; in fact, they do not even glance at the second and third-

round pie sizes in 19% and 10% of the trials, respectively. And classifying subjects by the amount

of look-ahead they do is strongly predictive of the offers they make.

The second study "turns off" social preferences (and beliefs that other players express social

---

having an offer of x accepted is $F(\alpha)$, so the proposer should choose $x^*=\text{argmax }(S-x)F(x/(S-2x))$. Simple calculations in early studies showed that offers are close to x^* for empirically plausible beliefs about $F(\alpha)$, but more powerful recent analyses show that offers are too generous, even controlling for risk-aversion (see Costa-Gomes and Zauner [20]; Henrich et al. [34]).

preferences) by having human subjects bargain with robot players who play subgame perfectly and maximize their own earnings, and believe the humans will too. The robot condition removes the influence of social preference. The subject's bargaining partner does not receive money (so the subject should feel neither envious nor guilty), and the robots don't care about being treated fairly.[7] Since equilibrium social preference theories do not explicitly admit limited cognition or heuristic processing, when social preferences are turned off human subjects should be able to calculate a payoff-maximizing perfect equilibrium offer ($1.25). In fact, the human players offered $1.84 on average, and frequently rejected the robots' opening offers of $1.25. While players did not learn to make and accept equilibrium offers from brief experience (eight trials), they could be quickly taught: Brief instructions about backward induction caused them to make offers to the robots within a couple of pennies of the self-interest equilibrium prediction.

Our third study addresses a different question: What happens when subjects trained in backward induction, who had played previously against robots, are mixed together with untrained subjects? The answer is that there is a tug-of-war: Trained subjects offer less and untrained subjects initially reject those low offers. Over time, offers by trained subjects rise (because their initial low offers got rejected) but offers by untrained subjects fall (as they learn that trained subjects will accept less). After eight periods, offers average around $1.60, roughly midway between the equilibrium offer and the empirical average offer from untrained subjects. Thus, the presence of trained subjects does draw a group of untrained subjects toward equilibrium, but not all the way.

---

7 The robot treatment also sharply penalizes deviations from perfect equilibrium behavior. For example, if a subject in the player 2 role rejects a near-equilibrium robot offer of $1.50, thinking it unfair, she will end up making less since a second-stage counteroffer giving her more than $1.25 will be automatically rejected by the robot. Harrison and McCabe [33] also used robots in ultimatum games (unbeknownst to their subjects) to see whether repeated

2. Studying Cognition with Measures of Information Acquisition

A novel feature of our approach is that we recorded the information subjects looked at, and for how long, as they bargained (see Payne, Bettman and Johnson [52]; Costa-Gomes, Crawford and Broseta [19]). These information-gathering data illuminate the heuristics subjects are using and enable an indirect test of whether they are computing perfect equilibria.

Insert Figure 1 (display) here

Figure 1 shows the computer display subjects saw, created using a software system called MOUSELAB that has been used in individual choice research.[8] The computer screen has 6 boxes. Hidden behind each box is the amount of the pie to be split in a round (left-hand column boxes) or the role of the subject in a round (right-hand column boxes), for each of the three rounds of the game. To see what is in a box, subjects used a mouse to move a cursor into the box. Once the cursor enters the box, the box automatically opens and reveals its contents. The box closes automatically when the cursor is moved outside it. By recording which box is open every 16 milliseconds, we collect a fine-grained measure of the time when each box is entered and exited, allowing us to study the order in which boxes were opened, and for how long each box was opened.

---

exposure to game-theoretic play would reduce subjects' offers. (It did.)

8 Applications include preference reversals (Johnson, Payne and Bettman [36]; Schkade and Johnson [62], reactions of decision-makers to time pressure (Payne, Bettman and Johnson [51]), and differences in risk attitudes across different response modes (Johnson and Schkade [37]). See Payne, Bettman and Johnson [52] for a thorough review.

Computerized measurement of information acquisition has several advantages over other methods. For instance, subjects can be tape-recorded giving on-line or retrospective explanations of their thinking ("verbal protocols"). But protocols are often too coarse to capture subtleties in thinking, they can distort the process being measured because speaking affects thinking, and subjects may be unable to explain why they make the choices they do. Another method is photographing the eyes of subjects while they look at a display of information. The pattern of eye movements reveals what information is being looked at, for how long. (Eye movements are widely used in studies of reading comprehension.) But eye movement machines are very expensive and the method is uncomfortable. Compared to these methods, MOUSELAB is easier because the mouse is a quick, low-cost pointing device which is easy to learn, unobtrusive, and more adaptable to group decision making.

MOUSELAB is also easily adapted to group situations, unlike verbal protocols and eye movements. The recording of information acquisition is automatic, so experiments can be run on networked computers without intervention by experimenters.

In Figure 1 the cursor is in the upper left-hand box, so the player sees that the pie in the first round

is $5.00.[9]

Players make an offer by moving the cursor to the dollar scale at the bottom of the box. As

the arrow moves along the scale, the numerical value the arrow is pointing to is shown to the left of

the response scale ($1.29, in Figure 1). Players click on the scale to make an offer.

We use information acquisition to test game theoretic principles as if they are predictions

about cognition. This runs counter to an important convention in economics--theories should be

judged only by their implications, because models may still make good predictions even if their

assumptions are false. However, game-theoretic sequential bargaining models do not make good

predictions about initial offers. Finding out which underlying assumptions are false is the next

logical step. Direct tests of information search are an efficient way to do so.[10]

An analogy can be made to revealed preference. Choices between objects (such as

commodity bundles or gambles) reveal unobservable preferences; studies of industrial output reveal

unobserved production functions. Asking subjects to `choose' information provides similar clues to

---

9 An important assumption in our analysis is that information search is closely linked to information usage (i.e.,
subjects do not memorize the box contents). This assumption is plausible for four reasons. First, after a little
practice using the mouse is easier than conscious memorization. Second, players opened some boxes several times
before making a decision, suggesting that they did not memorize the boxes' contents (and we did not observe
subjects taking notes, although they could). Third, we used slightly different pie sizes in each trial to inhibit
memory. The pie sizes were actually $5+a, $2.50+b, and $1.25+c (where a, b, and c ranged from -$.25 to +$.25).
By choosing a-b+c=0, the pie sizes changed each period but the equilibrium share for player 1 was held fixed.
Fourth, the results showed that the total amount of time all boxes are open did not decline much over trials, as it
would if they were memorizing.
10 There is ample evidence from cognitive psychology that the backward induction assumptions underlying the
theory are likely to be false. Search through a problem space is usually limited to a few levels and a subset of
possible actions (Newell and Simon [47]). Anderson et al. [1] found that computer programmers had terrible
difficulty learning the LISP language when it involved recursion like backward induction. Backward induction
places unusual constraint on working memory similar to embedded sentences like "Sarah bought the dog Rick bred
that Phil's friend liked so much," which are notoriously difficult for subjects to remember. In experimental
economics, Forsythe, Palfrey and Plott [26], Eckel and Holt [22], and others have found that people do not backward
induct instinctively, but can learn to, in markets and games. It is notable that Cox and Oaxaca [21] report good
performance by subjects in a job-search experiment which requires backward induction (but for which heuristic

unobservable thinking patterns. By analogy, educated guesses about a factory's production process can be made by observing the flow of inputs into the factory, length of processing time, and final output without ever looking inside the factory. Similarly, the information search data can be used to make educated guesses about how subjects are thinking (i.e., the production process in their "thinking factory"). For example, if a subject does not open the third box at all, showing the third-stage pie size, she <u>cannot</u> be computing a perfect equilibrium offer just as a car factory which never receives shipments of rustproofing chemicals cannot possibly be rustproofing its cars.

Of course, game theory makes no specific prediction about exactly how information is processed to reach a perfect equilibrium. But an empirical benchmark can be created by observing information-gathering of subjects who are instructed in how to do backward induction (and who do make equilibrium offers). When this is done (see below), it is evident that subjects open the second- and third-round boxes most often and look at them longest, make many transitions between those boxes, and look least at the first-round box.[11] Call a type of player who exhibits this pattern an equilibrium player.

We categorize limited cognition by assuming some players look ahead only one or two steps (creating two decision rules or "types"), and substitute empirical guesses about behavior at truncated nodes for thoughtful analysis (cf. Jehiel [35]). Level 0 players see the first round as an ultimatum game. They ignore future rounds and offer 40% of the first-round pie ($2.00), the average offer in ultimatum games. Level 1 players look one round ahead and act like they are playing a truncated two-stage game (cf. Neelin et al. [46]). They anticipate that player 2 will make an ultimatum offer of

---

search paths approximate optimal paths fairly closely).

[11] 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 than the second round pie, as subjects were told.

the round two pie (40% of $2.50, leaving $1.50 for player 2), so planning ahead one step, they offer $1.50 to player 2 in the first round.

These decision rules are ad hoc, but they do sort subjects into three categories somewhat cleanly according to the amount of time they spend looking at each of the three boxes. Furthermore, the level 0, level 1, and equilibrium players are predicted to offer $2.00, $1.50, and $1.25 respectively, so there is a clear relation predicted between search patterns and offers. The equilibrium social preference theories make no such prediction.

3. Study 1: Bargaining with other players

As in all the subsequent experiments, subjects were Wharton School undergraduates recruited from business or economics classes, or from general sign-up sheets posted around the campus.[12] Groups of 10 students met in the lab. An experimenter read instructions (see Appendix A) aloud as the players read their own copies, to make the instructions public knowledge. Afterwards, they worked through some examples (balanced to avoid biasing the results) and took a quiz to ensure they understood the instructions.

Each of ten subjects played eight three-round alternating-offer bargaining games, reversing roles each period, rematching opponents with no repetition. At the end of the session they were paid in cash according to their performance (half of their dollar earnings). They earned $12.31 on average, including $3 for participating. We conducted three sessions, a total of 120 trials.

We will first discuss offers and rejections, which replicate standard results. Then we turn

---

12 Some subjects were high school juniors attending a summer program at the University of Pennsylvania; they behaved like college students in the lab. Matthew Spiegel et al. [65] found few subject pool differences in

to the information search data.

3.1 Offers, rejections, and counteroffers

Alternating-offer bargaining experiments produced three stylized facts (e.g., Ochs and Roth [49]). First, average offers lie somewhere between equal split and equilibrium. Second, while no offers should be rejected (in equilibrium), some (10-20%) are rejected, and rejections of equilibrium offers are very common. Finally, most counteroffers (about 80%) are "disadvantageous"—that is, they yield less to the person making the counteroffer than the amount that same person previously rejected.

Figure 2 shows a histogram (left) of offers made in the first round (pooling across sessions and trials) and a plot (right) of rejected first round offers vs. the resulting second round counteroffer. The average offer is $2.11, about a third of the way from the equal split point ($2.50) to the equilibrium of $1.25. Most offers (88%) are closer to the equal split than to the equilibrium, but only 13% of the offers are very close to the equal split.

Insert Figure 2 (data summary)here

The shaded portion of each bar in the Figure 2 histogram indicates the number of rejected offers. Offers are rejected 10.8% of the time, a rejection rate comparable to those in earlier research. Equilibrium offers (between $1.20 and $1.40) are rare, and offers below $1.80 are rejected about half the time.

First-round rejections result in a second round of play. Of the offers made in the second round, Figure 2 shows that 85 percent of the offers are disadvantageous, falling above the diagonal line. These leave the proposer (player 2) with less than he or she would have received

---

bargaining game much like ours.

by accepting the first-round offer from player 1. This large percentage of disadvantageous counteroffers is also comparable to prior results.

The solid circles in the right panel of Figure 2 represent second round offers that were rejected. Of all second round offers, 21 percent (3/14) are rejected, resulting in a third round. Two of the three third-round offers were rejected and resulted in no payment for either side.

Given the novel information display we used, it is comforting to note that the results replicate the stylized facts about offers, rejections, and counteroffers observed in other experiments. To be absolutely sure that hiding pie sizes in boxes did not affect the outcome of the bargaining much, we ran another session using the same methods and parameters, but keeping the boxes open (so subjects did not have to move a mouse to acquire information). The results are a close replication: The average first round offer was $2.10, and 7.5% of opening offers were rejected.[13]

3.2 Information search

Player 1: Prior to opening offer.

We start our analysis of information search patterns by studying information search by

---

13 These figures are not significantly different from those in the closed-box study (t(154)=.18, p>.60, two-tailed test for offers, z(156)=.63 p>.52 for proportion of rejections). Forcing subjects to open boxes does not seem to affect the offers they make. Lohse and Johnson [42] examine the impact of monitoring information acquisition more extensively.

player 1 subjects in the first round of each game (pooled across all eight trials of all three experimental sessions). Table I presents these average data, and Figure 3 presents the equivalent icon graph and histograms of the distribution.

<div align="center">Insert Table I here</div>

Table I shows that most of the looking time (12.91 seconds out of 20.82 total seconds) is spent looking at the first-round pie size. Half as much time is spent looking at the second round. Subjects barely glance at the third box, only looking for about one second.[14] Each box is opened about 2-4 times (i.e. the "number of acquisitions") each trial. That means 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 I . Each table entry shows 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 about .88 times per trial.) Contrary to the backward induction prediction, there are always more forward transitions (above the table diagonal) than backward ones (below the diagonal).

<div align="center">Insert Figure 3 (icon graph) here</div>

To depict information acquisition data compactly we use icon graphs (see Figure 3). Each box in the graph corresponds to one of the three payoff boxes on the computer display. The width of each box is proportional to the number of times the box was opened (i.e., the number of acquisitions). The height of the shaded area in each box is proportional to the amount of time spent looking at that open box.   Horizontal lines show box midpoints to guide the eye.

---

14 To distinguish between 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. People do not accurately perceive anything they see that briefly (Card, Moran and Newell [16]).

The arrows represent the number of transitions between boxes, with the thickness of each arrow proportional to the transition frequency (transitions that occurred less than once a trial are omitted to sharpen contrast).

Figure 3 expresses what Table I summarized: Players look most often at the first-round pie size and make forward transitions more than backward ones. The two last columns of Figure 3 show the distribution of the number of acquisitions and total looking times for each box. The difference in both number of acquisitions and looking time among the boxes is quite apparent, as can be seen by comparing the distributions for different rounds within each column. It is particularly interesting that some subjects do not look at the second and third round payoffs at all. The first bars of the acquisition histograms in the middle column show the number of subjects who never open each box (i.e., they open it zero times). While all respondents look at the first round box at least once, 19 percent never look at the round 2 box and 10 percent never look at the round 3 box. These subjects cannot possibly be calculating the perfect equilibrium. A control session in which the pie sizes are presented upside down (i.e., with round 3 at the top) replicates the basic looking pattern.[15]

The relation between information search and first-round offers

It is useful to know whether information processing correlates with offers. We divided

---

15 An alternative explanation for the concentration of attention on first round payoffs is that information at the top of the display will be looked at more frequently and longer than information at the bottom, and the first round happens to be at the top. To control for this effect, we ran another study, identical to study 2, with the same instructions, computer display, pie sizes, and subject pool, except the order in which the pie sizes was displayed was inverted. The top line showed information for round 3, the middle line for round 2, and the bottom line for round 1. There were no significant differences in the proportion of time spent on each round size. However, subjects did have longer looking times for all three boxes, presumably because of the unnatural format, and the average offer was $1.94, significantly less than the $2.11 in study 2. Thus, the most important findings-- average offers are far from equilibrium, and subjects look at the first round pie size most often, and the third round least-- are robust to flipping the box order upside down.

trials into two groups based on first-round offers, corresponding to the offers we expect for

level 0 and level 1 cognition: $2.00 and above (n=84) and those between $2.00 and $1.50

(n=37). (The four offers below $1.50 are omitted.) In Figure 4, each column represents average

information search for trials in which player 1's make high (above $2) or medium ($1.50-2)

offers, corresponding roughly to the prediction of level 0 and 1 decision rules. The filled box

represents 15.34 seconds of looking time and the widest box represents 5.09 acquisitions

Insert Figure 4 (offer groups) here

Information search in the medium-offer (level 1) group shows more attention to the

second round payoff and more backward induction transitions than the high-offer (level 0) group.

In fact, subjects making level 1 offers open the second-round box more often than the first-round

box and move between the second and third round boxes more than once each trial.[16] Thus,

there is a strong relationship between offers closer to equilibrium, and search for information in

ways more consistent with backward induction. We return to this sort of player-type analysis

after study 2, in which we create enough observations about players who make low offers and

exhibit equilibrium patterns of information search. Then we will put players from both studies

together, divide them into level 0, level 1, and equilibrium groups by information measures, and

see what average offers look like in the three groups.

Finally, a natural hypothesis is that players learn to make perfect offers from experience.

---

16 Hypothesis tests were conducted using an Analysis of Variance (ANOVA) for each of the process measures, and subsequent t-tests to test for differences among the two response groups. The process measures are significantly different across the two groups. The only exception to this pattern is the time spent on the first round payoff, which does not differ between the two groups. In all significance tests we report, differences among groups is tested by an **F** test that the group means differ, **p** < .05. Specific differences between pairs of groups are tested by **t** tests, **p** < .05 for a

However, the average offer changes little over time, and information-processing measures do not change either.[17]

Information Search by Player 2.

So far we have only discussed first-round offers made by player 1. What happens after[18] player 2 receives an offer?

Insert Figure 5 (player 2) here

Figure 5 is an icon graph of the information search by player 2s, broken down by whether offers were rejected or accepted. Acceptances are grouped by the size of the offer accepted. Overall, the search patterns look much like the patterns for player 1, with attention concentrated on the first round payoff. Recall that almost all (12/14) of the rejected offers were below $2.00. Therefore, a telling comparison is between the "rejected" (first) column in Figure 5 and the accept "<2.00" (second) column. Those who accepted low offers looked at the second-round pie value almost twice as long as those who rejected the proposal; they also made more transitions back and forth between the first and second round pie boxes.[19] While looking ahead more is correlated with accepting a low offer, we do not know the direction of causality.

---

priori hypotheses, and Bonferroni confidence intervals for post-hoc tests.

17 ANOVA's on each of the process measures show no differences among the trials for any of the process measures, except the number of transitions between the first and third round payoff. Those transitions are rare (occurring less than once per trial) and of no theoretical significance.

18 Player 2s could not open boxes until they received an offer, so their lookup behavior is conditioned on the offer they received.

19 However, these differences are not significant because there are so few rejections (14 of 120) that the comparison lacks power.

4. Study 2: Bargaining with Robots and Instructions

In study 1, subjects may know how to reason game-theoretically but not do so because they care about fairness or believe other subjects are concerned with fairness. (Indeed, offering $1.25 is not profit-maximizing given the rejection rates observed.) In study 2, we control for these possibilities by matching players with robot players. To disable fairness concerns, our instructions provided public knowledge that the computers simulated rational players. Subjects were explicitly informed (see Appendix C) that they play against a computer that behaved as if it cared only about its own payoff, believed other players behave the same way, and believed that other subjects also believe also all subjects do the same.

Subjects who only made generous offers in study 1 because they cared about behaving fairly toward others, or who feared rejection of equilibrium offers from fair-minded players, should offer $1.25 to robots. The robots also may help the human players learn because they make equilibrium offers (and accept them). Since the study 2 procedure is otherwise identical to study 1, we can assess the relative impact of fairness vs. limited cognition on departures from subgame perfect play by comparing the results of the two studies.

Thirty subjects, in groups of 10, first played eight trials, with the same equilibrium as in study 1. After finishing the first eight trials, subjects were given explicit instructions (see Appendix B) on using backward induction in these games. Two numerical examples were given, and the subjects were walked through a third example where they had to figure out the optimal offer using the backward induction strategy, and received feedback. Then subjects were instructed to use this strategy when playing the remaining eight trials. They were again reminded that they were playing against the computer and the trials were independent of each

other.

In a second set of sessions, the first four post-training trials had the same equilibrium as in the first eight trials (and all previous studies). Then they played four trials with a different equilibrium, to assess whether they could transfer or generalize the backward induction strategy to a different set of parameters. Payment procedures were the same as in Study 1. Subjects earned $16.24 on average, including $3 for attending.

4.1 Results

Offers, Counteroffers, and Rejections.

If concerns for fairness are all that is preventing subgame perfect play among people, offers to robots should be near the subgame perfect equilibrium of $1.25. If generous offers to people come from limited cognition, we should see generous offers to robots too.

Table II shows the offers and rejections from study 2, aggregated over the 30 human players. Each column represents the mean offer and the fraction of rejections (out of a possible 30) for each pair of trials, in which each human player either made the first round offer or decided to accept or reject the offer made by the robot.

The offers in trials 1-8 are significantly lower than those made in study 1, averaging $1.84 instead of $2.11, but are far from the game-theoretic prediction and do not converge toward it across trials. Rejections of first-round robot offers are common: Half (61/120 or 51%) are rejected, although the human players reject more in the first two trials (83%) than in trials 3-8 (40%).

Instruction causes a dramatic change. The average offers are near the equilibrium, $1.22,

and the rejection rate falls to 7%.[20] When presented with a set of pies with different

parameters, in trials 13-16, players make offers which are close to the new equilibrium and they

rarely reject the game theoretic offers. The players have learned a general principle (the

backward induction strategy) and transferred it to a structurally-similar game with a different

equilibrium.

While playing robots did not quickly produce equilibrium play among uninstructed

players, it **did** lower offers, from $2.11 to $1.84. Loosely speaking, concern for fairness toward

(or from) other human players accounts for only $.27 of the $.86 gap between observed offers

and equilibrium. However, the gap disappears rapidly after instruction in backward induction.

Insert Table II here

4.2 Information Search

Figure 6 shows the icon graphs for all eight trials in which human subjects were players

1. These data corresponds to the offers reported in the second row in Table II. The looking time,

acquisitions, and pattern of search for the first four offers look much like that from study 1

(except the third trial, where they look ahead more often).

After instruction (between the 4th and 5th columns), the first round box is open much

less, and attention is shifted to the second and third round payoff boxes, coinciding with the shift

in offers in Table II. Across the 16 trials, looking ahead more and offering less are modestly

related because the correlation between average third-round box looking time and average

---

20 The discrepancy between the average response $1.22 and the equilibrium ($1.25) is probably due to the
MOUSELAB response scale, which recorded responses to the nearest nickel. All differences due to instructions that
we report were tested with an Analysis of Variance and reached or exceeded conventional (.05) levels of

deviation from equilibrium is -.49 (p<.05).

Insert Figure 6 (robot icon graphs) here

Pooling the observations of trained players who make low offers in study 2 with the variety of offers from study 1 now gives enough variation to powerfully correlate players' types (sorted by information measures) with their offers. We define level 0 cognition as opening the second-round box strictly less often than the first-round box. Level 1 cognition is the opposite, looking at the second round weakly more often than the first. We define equilibrium play as looking at the third round box weakly more often than the first and second. (Observations that satisfy both level 1 and equilibrium criteria are classified as equilibrium.) These definitions are easy to use because the number of acquisitions is discrete. And since acquisitions, looking times, and transitions are all highly correlated, a more complex discriminant analysis using other statistics added nothing substantive. Note also that the unit of analysis is a trial, not a player. A player-specific analysis paints a similar picture.

Insert Table III and Figure 7 here.

Table III and Figure 7 present the results of this sorting. While we categorized observations based purely on simple search criteria, the average offers in each category are significantly different by ANOVA. The averages in the level 0, level 1, and equilibrium categories of $2.07, $1.71, and $1.44, are above the predictions of $2, $1.50, and $1.25 but they are ranked correctly and simply serve as a reminder that the decision rules we posited could be refined. Looking times and transitions also plausibly correspond to decision rules. Note in Figure 7, that there are many more second- and third-round transitions corresponding to backward

significance.

induction in the equilibrium type observations, although this is not strictly required by the pattern of acquisitions.   Finally, half the trials classified as reflecting equilibrium cognition come after training.

 Study 3: Mixing trained and untrained subjects

Studies 1-2 raise an interesting question: What happens when trained and untrained subjects are mixed together? To find out we did one more study. In this study 3, five subjects first played four trials against robots without instruction, then played four trials after instruction. In trials 9-16, these five trained subjects were matched with five untrained human players. Neither type of player was told anything specific about the other type's training. However, when instructions were read to the untrained subjects, the trained subjects were present and they went through the instructions (for a second time) along with the untrained subjects. We ran two sessions, with 10 subjects in each, for 16 trials. Of course, there are many possible design configurations for this kind of experiment, varying what subjects knew about each other, the proportion of trained subjects, and so on. This study is simply meant to be the beginning of exploring trained-untrained population mixtures.

What might happen when subjects who are good at game-theoretic reasoning (after their training), but may still care about fairness toward other people (or think others care), are combined with others who do not instinctively reason game-theoretically? Who "teaches" whom? Perhaps trained subjects convey what they have learned by making low offers to the untrained subjects, who learn to accept low offers and almost make lower offers themselves.

(Note that this outcome is not likely, a priori, because untrained subjects did not learn to make equilibrium offers after playing "trained" subjects-- the robots— in study 2.) Or the opposite outcome could occur— even though they know how to compute equilibrium offers, if untrained subjects persistently reject their low offers the trained subjects may be forced to revert to making generous offers which reflect (perceptions of) social preference. Or the result may lie between these extremes.

Results

Results are summarized in Table IV and Figure 8. The average offer in pretraining trials 1-4 was $1.83, close to the $1.92 in the first four trials of study 2, and rejection rates are comparable (75% vs. 60%).   Training had a significant effect in trials 5-8, lowering offers from $1.83 to $1.46 and rejections from 75% to 35%.[21]

Insert Figure 8 (graph of offers & rejections, study 3) here

Figure 8 shows the average offers by players 1, and the proportion of offers rejected (as player 2) for each pair of trials and player type, in the mixed-population trials 9-16. (Recall that since players alternated between the two roles, each four-trial block generated only two periods of player 1 and two periods of player 2 observations.) Vertical hash marks separate trials before and after training, and before and after the introduction of untrained players. The horizontal line shows the subgame perfect equilibrium.

The mixed-population trials 9-16 resemble a 'tug-of-war' in which trained players make

---

21 Note that training did not reduce offers and rejections as close to equilibrium as in study 2, either because subjects had less post-training practice (four trials rather than eight) or due to two outlying subjects who did not

low offers that are repeatedly rejected. The overall rejection rate is 30% (compared to 11% in study 1) but there is a more striking difference in the rate of complete breakdown: rejections of third-round offers, resulting in no earnings for either player. In study 1, only 1.7% of trials ended in breakdown but in study 3 that figure was 12.5%. Over time, both sides converge slightly from their initial positions, although the untrained subjects lower their high initial offers more than the trained subjects raise their offers. The average offer after eight periods of mixed play is around $1.60. This intermediate figure is substantially above the subgame perfect equilibrium of $1.25 and the level (about $1.40) of offers made by trained players after instructions, but is well below the average offer of $2.28 made by untrained subjects in the first two trials.

To examine the process used to calculate and evaluate these offers, we again look at information search. Figure 9 is an icon graph summarizing the attention measures. When playing game-theoretic robots in rounds 1-4, the initially untrained subjects look more often and for longer periods of time at the first box and pay little attention to the second and third round payoffs. As in study 2, the brief instruction in backward induction produces dramatic changes in looking time. The last two columns contrast the search patterns on the now trained subjects with those who have just been introduced to the game, when they bargain together. The trained subjects spend a bit less time looking at all boxes than when playing robots, but their patterns still look very different from the untrained subjects.

Insert Figure 9 (icon graph, tug of war study 3) here

6. Implications and Conclusions

We ran experiments on three-round alternating-offer bargaining over a "shrinking pie".

---

learn as well.

Offers were scattered between an equal split of the first-round pie ($2.50) and the perfect

equilibrium offer $1.25, averaging $2.11. Offers could be well above equilibrium because

subjects think strategically but have a social preference for equal offers or fair treatment (or

believe others have such a preference), or because subjects use only limited cognition. (And of

course, both hypotheses could be true.) We tested these competing explanations in two ways.

First, we conducted sessions in which players bargained with self-interested robots.

Subjects who can do backward induction instinctively but are constrained by social preferences

when they bargain with people should offer $1.25 to the robots. They offered only $1.84 on

average, which clearly shows that most subjects were not using game-theoretic reasoning.

However, they could be taught to make such offers with very brief instruction.

Second, we measured patterns of information search by showing the three rounds of pie

sizes in boxes that could only be opened by moving a cursor into a box. These data show that

most subjects did not look at the pie sizes in the correct order, and for the length of time,

necessary for backward induction (compared to a control group trained in backward induction).

Instead, most subjects concentrated on the current round when making decisions and looked

ahead insufficiently. In 19% and 10% of the trials, respectively, they did not even open the

second and third-round boxes. Importantly, looking patterns are related to behavior: Subjects

could be classified using the information measures into those who looked only at the first round

(level 0), those who looked one step ahead (level 1) and those who did equilibrium reasoning.

These classifications are highly predictive of average offers; those who looked ahead further

offered less.

The results show that social preference and limited use of backward induction both play a

role in bargaining. Social preferences matter because about a third of the gap between the prediction and the mean offer disappeared when subjects switched from bargaining with people to bargaining with robots. Furthermore, experiments on very simple games show that social preferences are being expressed even when no strategic thinking is required (e.g., evaluating an ultimatum offer or deciding how much to allocate in a dictator game).

At the same time, the results show clearly that subjects do not think as strategically as most game theory applications presume. However, it is crucial to note that subjects can be taught to think strategically very rapidly, and their learning transfers to a bargaining game with different pie sizes. Backward induction is not computationally complex; it is simply not natural, at least in an abstract, multi-stage game like this. It is like learning to lean back on a windsurfer and forward on your skis, knowing that investments earning a compound interest rate of r doubles in $.72/r$ years, or knowing that doubling the tax on a restaurant bill in a state with 8% sales tax will come close to the amount you should tip. Backward induction is a neat trick that is not computationally hard to execute (at least in these simple games) once it is taught, but it seems to be difficult to induce from raw experience.

In our study 3, we combined subjects who had learned to make self-interested equilibrium offers to robots with untrained subjects. The two types of subjects engaged in a tug-of-war resulting in a mean offer after eight trials was about $1.60, a little less than halfway between the perfect equilibrium and the level of offers the untrained subjects make at the start. The looking patterns show that the trained subjects continued to "use" their training, looking more often at future pie sizes, although they looked ahead less often than they did when playing robots.

One interpretation of our results is that we measure strategic heuristics that were evolutionarily adapted to facilitate social exchange in primitive environments and are encoded in specialized cognitive modules. Presumably, this adaptation did not equip people to do backward induction, but did equip them to react hostilely to unfair offers, and to anticipate hostile reactions of others. Samuelson [61] formalizes a sophisticated form of this argument.. Our research and evolutionary interpretations are complementary. We are simply trying to understand how subjects think when they bargain, by measuring what information they look at it, to understand the offers they make and accept. The evolutionary interpretation leapfrogs further back in time, asking where the bargaining heuristics we observe might have come from.

Our work leaves much more to do for both empiricists and theorists. For empiricists, measurement of information processing could be applied to many other games— and already has been. Costa-Gomes, Crawford and Broseta [19] test whether players look at game payoffs in ways consistent with iterated eliminated or dominated strategies and other decision rules. Cachon, Camerer and Johnson [11] show that violations of forward induction are common (see Cachon and Camerer [10]) and can be predicted by how frequently player 2 looks back at player 1's foregone payoff (which is required for player 2's forward induction). Salmon [60] uses information acquisition to distinguish different learning theories.

For theorists, the challenging of deriving plausible decision rules from other principles (and data) looms large. We presented above a plausible model in which subjects make heuristic judgments about which future nodes to ignore, then reason game-theoretically about the simplified game, using empirical judgment about the value of the continuation game from truncated nodes. (Note that subgame perfection also eliminates future nodes, but only after

figuring out which nodes are likely to be reached by working through all of them.) A

complete theory could endogenize which nodes are truncated, and how values are imputed at

those nodes, tying these steps to parameters of the game. Necessary statistical power for

distinguishing models will probably come only from combining choices and cognitive measures

as dependent variables (see Costa-Gomes et al. [19]). Syntheses of this sort should take the

cognitive content of actual strategic thinking seriously, while preserving the emphasis on

statistical care and inferences from the choices people make which has come to be a useful

standard in recent studies.

Our research helps develop behavioral game theory by using experimental measurement

to understand the roles of thinking and learning in producing both successes and failures of game

theoretic predictions, and to shape new theory. This is part of a sensible shift away from

interpreting game-theoretic equilibria as solutions that brilliant players figure out, toward

thinking of equilibria as resting points which are the result of evolution or learning by players of

limited rationality (e.g. Kandori, Mailath and Rob [39]; Binmore and Samuelson [7]; Roth and

Erev [56]; Camerer and Ho [13]). This emerging approach also follows earlier studies of

individual choice that examine how heuristic processes can be procedurally rational (Payne,

Bettman and Johnson [52]).

References

1.   J. R. Anderson, P. Pirolli, and R. Farrell, The nature of expertise, *in* "Learning to program recursive functions" (M. Chi, R. Glaser, and M. J. Farr, Eds.), Earlbaum, Hillsdale, NJ, 1988.

2.   J. Berg, J. Dickhaut, and K. McCabe, Trust, reciprocity and social history, *Games Econ. Behav.* **10** (1995), 122-142.

3.   K. Binmore, J. McCarthy, G. Ponti, L. Samuelson, and A. Shaked, "A backward induction experiment," September, 2000. http://www.ssc.wisc.edu/~larrysam/papers.htm

4.   K. Binmore, A. Rubinstein, and A. Wolinsky,. The Nash bargaining solution in economic modelling, *RAND J. Econ.* **17** (1986), 176-188.

5.   K. Binmore, A. Shaked, and J. Sutton, Testing noncooperative bargaining theory: A preliminary study, *Amer. Econ. Rev.* **75** (1985), 1178-1180.

6.   K. Binmore, A. Shaked, J. Sutton, A further test of noncooperative bargaining theory: Reply, *Amer. Econ. Rev.* **78** (1988), 837-839.

7.   K. Binmore and L. Samuelson, Drift, *Europ. Econ. Rev.* **38** (1994), 859-867.

8.   G. E. Bolton, A comparative model of bargaining: Theory and evidence, *Amer. Econ. Rev.* **81** (1991), 1096-1136.

9.   G. E. Bolton and A. Ockenfels, ERC: A theory of equity, reciprocity and competition, *Amer. Econ. Rev.* **90** (2000), 166-193.

10.   G. Cachon and C. F. Camerer, Loss-avoidance and forward induction in experimental coordination games, *Quart. J. Econ.* **444** (1996), 165-194.

11.   G. Cachon, C. F. Camerer, and E. J. Johnson, Information processing and forward induction, *Department of Decision Sciences working paper,* University of Pennsylvania (1994).

12.   C. F. Camerer, "Behavioral Game Theory: Experimental studies of strategic interaction," Princeton University Press, Princeton, 2001.

13.   C. F. Camerer and T. H. Ho, Experience-weighted attraction learning in normal-form games, *Econometrica* **67** (1999), 827-874.

14.   C. F. Camerer, E. J. Johnson, S. Sen, T. Rymon, Cognition and framing in sequential bargaining for gains and losses, *in* "Frontiers of Game Theory" (K. Binmore, A. Kirman,

and P. Tani, Eds.), MIT Press, 1993.

15.   C. F. Camerer and R. H. Thaler, Anomalies: Ultimatums, dictators, and manners, *J. Econ. Perspect.* **9** (1995), 209-219.

16.   S. K. Card, T. P. Moran, and A. Newell, The Psychology of Human-Computer Interaction, Earlbaum, Hillsdale, NJ, 1983.

17.   J. P. Carpenter, Bargaining outcomes as the result of coordinated expectations, *working paper,* Middlebury College (October, 2000).

18.   G. Charness and M. R. Rabin, Social preferences: A model and new evidence, *Department of Economics working paper*, Berkeley (2000)

19.   M. Costa-Gomes, V. Crawford, and B. Broseta, Cognition and behavior in normal-form games: An experimental study, *Econometrica*, forthcoming (2001).

20.   M. Costa-Gomes and K. Zauner, Ultimatum bargaining behavior in Israel, Japan, Slovenia, and the United States: A social utility analysis, *Games Econ. Behav.* **34**, (2001), 238-269.

21.   J. C. Cox and R. L. Oaxaca, Laboratory experiments with a finite-horizon job - search model, *J. Risk and Uncertainty* **2** (1989), 301-329.

22.   C. Eckel and C. A. Holt, Strategic voting in agenda-controlled committee experiments, *Amer. Econ. Rev*. **79** (1989), 763-773.

23.   M. El-Gamal and D. M. Grether,   Are people Bayesian? Uncovering behavioral strategies, *J. Amer. Statistical Assoc.* **90** (1995), 1127-

1145.

24.   E. Fehr and S. Gachter, Fairness and retaliation: The economics of reciprocity, *J. Econ. Perspect.*, **14** (Summer 2000), 159-181.

25.   E. Fehr and K. Schmidt, A theory of fairness, competition, and cooperation. *Quart. J. Econ*. **114** (1999), 817-868.

26.   R. Forsythe, T. Palfrey, and C Plott, Asset valuation in an experimental market, *Econometrica* **50** (1982), 537-567.

27.   X. Gabaix and D. Laibson, Bounded rationality and directed cognition. October, 2000. http://www.mit.edu/xgabaix/www/broct16.pdf.

28.   J. K. Goeree and C. A. Holt: "Stochastic game theory: For playing games, not just for doing theory," *Proceedings of the National Academy of Sciences,* **96** (1999) 10564-10567.

29.   W. Güth and R.Tietz,. Ultimatum bargaining behavior: A survey and comparison of experimental results, *J. Econ. Psych.* **11** (1990), 417-449.

30.   D. W. Harless and C. F. Camerer, The predictive utility of generalized expected utility theories, *Econometrica*, **62** (1994), 1251-1290.

31.   G. W. Harrison, Theory and misbehavior of first-price auctions, *Amer. Econ. Rev.* **77** (1989), 749-762.

32.   G. W. Harrison and K. McCabe, Testing bargaining theory in experiments, *in* "Research in Experimental Economics*"* (R. M. Isaac, Ed.), Vol. 5 1992.

33.   G. Harrison and K. McCabe, Expectations and fairness in a simple bargaining experiment, *Int. J. Game Theory* **25** (1996), 303-327.

34.   J. Henrich, R. Boyd, S. Bowles, H. Gintis, C. Camerer, E. Fehr, and R. McElreath, In search of homo economicus: Experiments in 15 small-scale societies. *Amer. Econ. Rev*. May, 2001.

35.   P. Jeheil, Predicting by analogy and limited foresight in games, *working paper* CERAS and University College London (1999). http://else.econ.ucl.ac.uk/papers/analog.pdf.

36.   E. J. Johnson, J. W. Payne, and J. R. Bettman, Information displays and preference reversals *Organizational Behavior and Human Decision Processes* **42** (1988), 1-21.

37.   E. J. Johnson and D. A. Schkade, Heuristics and bias in utility assessment: Further

evidence and explanations, *Management Science* **35** (1989), 406-424.

38.   D. Kahneman and A. Tversky, Prospect theory: An Analysis of decision under risk, *Econometrica* **47** (1979), 263-91.

39.   M. Kandori, G. Mailath, and R. Rob, Learning, mutation, and long-run equilibria in games, *Econometrica* **61**, (1993) 29-56.

40.   J. Kennan and R. Wilson, Can strategic bargaining models explain collective bargaining data? *Amer. Econ. Rev.* **80** (1990) 405-409.

41.   G. Loewenstein, L. Thompson, and M. Bazerman, Social utility and decision making in interpersonal contexts, *J. Personality and Social Psychology* **53** (1989), 426-441.

42.   G. Lohse and E. J. Johnson, A comparison of two process tracing methods for choice tasks, *Organizational Behavior and Human Decision Processes*, **68** (1997), 28-34.

43.   R. D. McKelvey and T. R. Palfrey, An experimental study of the centipede game, *Econometrica* **60** (1992), 803-836.

44.   R. D. McKelvey and T. R. Palfrey,. Quantal-response equilibrium, *Games Econ. Behav.* **10** (1995), 6-38.

45.   R. D. McKelvey and T. R. Palfrey, Quantal response equilibria for extensive form games, *Experimental Economics* **1** (1998), 9-41.

46.   J. Neelin, H. Sonnenschein, and M. Spiegel, A further test of noncooperative bargaining theory: Comment, *Amer. Econ. Rev.* **78** (1988), 824-836.

47.   A. Newell and H. A. Simon, "Human Problem Solving," Prentice Hall, Englewood Cliffs, NJ, 1972.

48.   R. E. Nisbett and T. D. Wilson, Telling more than we can know: Verbal reports on mental processes, *Psychological Review* **84** (1977), 231-259.

49.   J. Ochs J. and A. E. Roth, An experimental study of sequential bargaining, *Amer. Econ. Rev.* **79** (1989), 355-384.

50.   M. J. Osborne and A. Rubinstein, "Bargaining and Markets," Academic Press, San Diego, CA, 1990.

51   J. W. Payne, J. R. Bettman, and E. J. Johnson, Adaptive strategy selection in decision making, *Journal of Experimental Psychology: Learning, Memory, and Cognition* **14** (1988), 534-552.

52.   J. W. Payne, J. R. Bettman, and E. J. Johnson, "The Adaptive Decision Maker," Cambridge University Press, Cambridge, 1993.

53.   D. P. Porter and V. L. Smith, Futures contracting and dividend uncertainty in experimental asset markets, *J. Bus.* **68** (1995), 509-541.

54.   M. Rabin, Incorporating fairness into game theory and economics, *Amer. Econ. Rev.* **83** (1993), 1281-1302.

55.   A. E. Roth, Bargaining experiments, *in* "Handbook of Experimental Economics" (G.E. Roth and J. Kagel, Eds.), Princeton University Press, Princeton, 1995.

56.   A. E. Roth and I. Erev, Learning in extensive form games, experimental data and simple dynamic model in the intermediate term, *Games Econ. Behav.* **8** (1995), 164-212.

57.   A. Roth, M Okuno-Fujiwara, V. Prasnikar, and S. Zamir, Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh and Tokyo: An experimental study, *Amer. Econ. Rev.* **81** (1991), 1068-1095.

58.   A. Rubinstein, Perfect equilibrium in a bargaining model, *Econometrica* **50** (1982), 97-109.

59.   J. E. Russo, A software system for the collection of retrospective protocols prompted by eye fixations, *Behavior Res. Methods & Instrumentation* **11** (1979), 177-179.

60.   T. C. Salmon, Evidence for learning to learn behavior in normal form games, *working paper* Florida State University (1999). http://garnet.acns.fsu.edu/~tsalmon/

61.   L. Samuelson, Analogies, adaptations, and anomalies, *J. Econ. Theory*, in press.

62.   D. A. Schkade and E. J. Johnson, Cognitive processes in preference reversals, *Organizational Behavior and Human Decision Processes* **44** (1989), 203-231.

63.   M. Shubik, The dollar auction game, *J. Conflict Resolution* **15** (1971), 109-111.

64.   D. Ståhl and P. Wilson, On players' models of other players: Theory and experimental evidence, *Games Econ. Behav.* **10** (1995), 218-254.

65.   M. Spiegel, J. Currie, H. Sonnenschein, and A. Sen, Understanding when agents are fairmen or gamemen, *Games Econ. Behav.* **7** (1994), 104-115.

66.   I. Ståhl, Bargaining theory, *Swedish J. Econ.* **75** (1973), 410-413

67. G. Weizsacker, Ignoring the rationality of others, *working paper* Harvard Business School (2000).

Appendix A: Instructions

   This is an experiment about bargaining. You will bargain using the computer in front of you, in several sessions. In each session you will either be a BUYER or a SELLER. Each BUYER will be randomly paired with a SELLER in a session. You will be paired with a different person in each session.

               Please go to the next screen

      At the end of the experiment, you will be paid half of the total amount you actually earned bargaining, plus $3 for attending.
    Please do NOT talk to other subjects at any point during the experiment. Direct any questions you may have to us.

               Go to the next screen....

 This is how the bargaining works: In each session the bargainers will divide an amount of money, which we will call a "pie". The pie is initially worth some amount ($4.00, for example). Each session has three rounds. In the FIRST round, the SELLER will offer part of the pie to the BUYER; the rest of the pie is left for the SELLER. If the SELLER's offer is accepted by the BUYER, the session ends and the bargainers earn their shares of the pie.

               Go to next screen....

    If the SELLER's offer is rejected by the BUYER, the session continues to a SECOND round. The SECOND round is like the FIRST except the pie is smaller (for example, $2.00 instead of $4.00) and the roles of BUYER and SELLER are reversed: the original BUYER makes an offer to the original SELLER. If the BUYER's offer is accepted by the SELLER, the session ends. If the offer is rejected, the session continues to a THIRD round.

               Go to the next screen....

 In the THIRD round, the pie will get even smaller (for example, $1.00 instead of $2.00) and the roles are reversed again: the original SELLER makes an offer to the original BUYER. If the SELLER's offer is rejected in the THIRD round, the session ends (there is no fourth round), and both bargainers get zero in that session.

   Notice that the pie in the SECOND round will be split only if the FIRST round goes on. And similarly, the pie in the THIRD round will be split only if the SECOND round goes on.

               Go to the next screen....

   Now we will illustrate how the bargaining works through an example. It is shown on a sheet

in front of you. Please wait here until we go through the example.

Thank you in advance for your participation.

Now, we will walk you, screen by screen, through an EXAMPLE session using the computer and the mouse.

All the sessions in the actual experiment will have the same structure as this EXAMPLE

Go to next screen....

The EXAMPLE is now over. The experiment starts with the next session.

Please don't feel that any particular kind of behavior is expected of you. There is no correct way of bargaining in this experiment. All your responses will be coded by a subject number, so your responses will be anonymous to us.

So that you can keep track of the bargaining process in each session, please record offers and accept/reject responses and calculate your share of the pie in all three rounds of each session on the worksheet provided.

However, please do not write down total pie values on your worksheet.

Remember: Do NOT talk to other subjects at any point during the experiment. Direct any questions you may have to us.

Appendix B

The first part of the experiment is now over. Please wait here until the experimenter tells you to go to the second part.

One way of bargaining in the situation you were just in is to think first about what would happen in the third (last) round of the session. Imagine that you are a seller in the third round. You would make an offer to the buyer. If she or he agreed to the split, they would get what you offered them and you would get the remainder of the pie.

If the buyer were trying to make the most money and knew that you were trying to do the same, then the buyer should settle for any amount greater than zero (which is what they would receive for that session if they rejected your third round offer). The buyer will know that the most he can get in the third round is a penny, so you can get the third round amount minus a penny.

Now, in the second round, the buyer should know what you could make in the third round. In the second round, he could offer you one penny more than what you could make in the third round, and you would accept it. The buyer would get the rest of the second round pie (the total pie minus the amount you would accept).

Now imagine that you are in the first round of the session. You know what the buyer can get in round 2. In round 1, you should offer him or her one penny more than what she or he could make in round 2.

An example may help in understanding the bargaining strategy that was just described to you. Imagine that in a particular session the first round pie value is $4, the second round pie value is $2 and the third round pie value is $1. If you were making a first round offer, what would your offer be?

If both you and the buyer were trying to make the most money, then in the third round you would offer the buyer one cent and keep 99 cents for yourself. The buyer would accept the one cent because it is greater than nothing which is what they would receive for that session if they rejected your third round offer.

Likewise, in the second round, the buyer would offer you one dollar (99 cents plus an additional cent) and keep the remaining dollar for him or herself. You would accept this offer because it would be one cent more than what you could get if you rejected this offer and went to the third round of the session.

In the first round, then, you would offer the buyer $1.01. They should accept this since it is one cent more than what they could make in round 2.

Here's another example. Imagine that the first round pie value is $3, the second round pie value is $1 and the third round pie value is 75 cents. If you were making a first round offer, what would your offer be?

In the third round you would offer the buyer one cent and keep 74 cents for yourself.

In the second round, the buyer would offer you 75 cents (74 cents plus an additional cent) and keep the remaining 25 cents.

In the first round you would offer the buyer 26 cents. You would keep $2.74 for yourself.

Before we start the second part of the experiment, we will walk you through an EXAMPLE session. Try to figure out the optimal offer using the strategy just described.
The EXAMPLE is now over. The optimal offer in this example was $2.01. If you have any questions about the example, please raise your hand and the experimenter will come to you.

The second part of the experiment starts with the next session. In the next few sessions, use the strategy we just explained to you in generating your offers and in deciding whether to accept or reject offers in each session. As before, you will be playing against a computer, and the sessions are independent of each other.

Appendix C

In generating your offers, or deciding whether to accept or reject offers, assume the following:

1. You will be playing against a computer which is programmed to make as much money as possible for itself in each session. The computer does not care how much money you make.

2. The computer program expects you to try to make as much money as you can, and the program realizes that you have been told, in instruction (1) above, that it is trying to earn as much money as possible for itself.

3. The sessions are independent of each other. The computer program does not store information about the previous sessions; it uses the same program in each session.

Round	Number of Acquisitions	Total Time Examining Payoff	Transitions		
			1	2	3
Round 1 ($5.00)	4.38	12.91	--	2.55	.65
Round 2 ($2.50)	3.80	6.67	2.10	--	1.24
Round 3 ($1.25)	2.12	1.24	.50	.88	--

Table II: Results in training study 2

	No Instruction, Equilibrium= $1.26				Post Instruction, Equilibrium= $1.26		Post Instruction, Equilibrium= $1.51	
Round	**1-2**	**3-4**	**5-6**	**7-8**	**9-10**	**11-12**	**13-14**	**15-16**
Average Offer	1.83	2.02	1.67	1.82	1.21	1.23	1.51	1.54
Number of Rejections (out of 30)	25	11	12	13	1	4	6	0

Table III: Results and Process Measures by Strategy, Studies 1 and 2. Standard Deviations are in Parentheses, and all means that do not share a subscript are significantly different p < .05.

	N of trials	Average Offer	Percent Trained	Time 1st Payoff	Time 2nd Payoff	Time 3rd Payoff	Acq. 1st Payoff	Acq. 2nd Payoff	Acq. 3rd Payoff
Zero	136	2.07 (.43)	5.1	12.44_a (9.8)	2.27 (8.0)	1.13 (1.2)	3.40 (2.3)	1.51 (2.1)	1.40 (1.20
One	107	1.71 (.49)	21.4	11.60_a (14.4)	14.67_a (15.4)	5.82 (9.07)	4.63 (3.1)	6.33 (4.0)	3.32_a (2.4)
Equilib-rium	57	1.44 (.44)	52.6	7.09 (10.2)	14.89_a (19.7)	10.04 (13.2)	2.02 (1.7)	3.46 (2.7)	3.80_s (3.0)

Table IV: Offers and rejection frequencies in trained-untrained mixture study 3

	Trials 1-4 (pre-instruction)	Trials 5-8 (post-instruction)	Trials 8-16	
			Trained	Untrained
Amount offered by player 1, round 1	$1.83	$1.46	$1.49	$1.90
Percentage rejections by player 2, round 1	75%	35%	28%	33%

Figure 1 MOUSELAB computer display
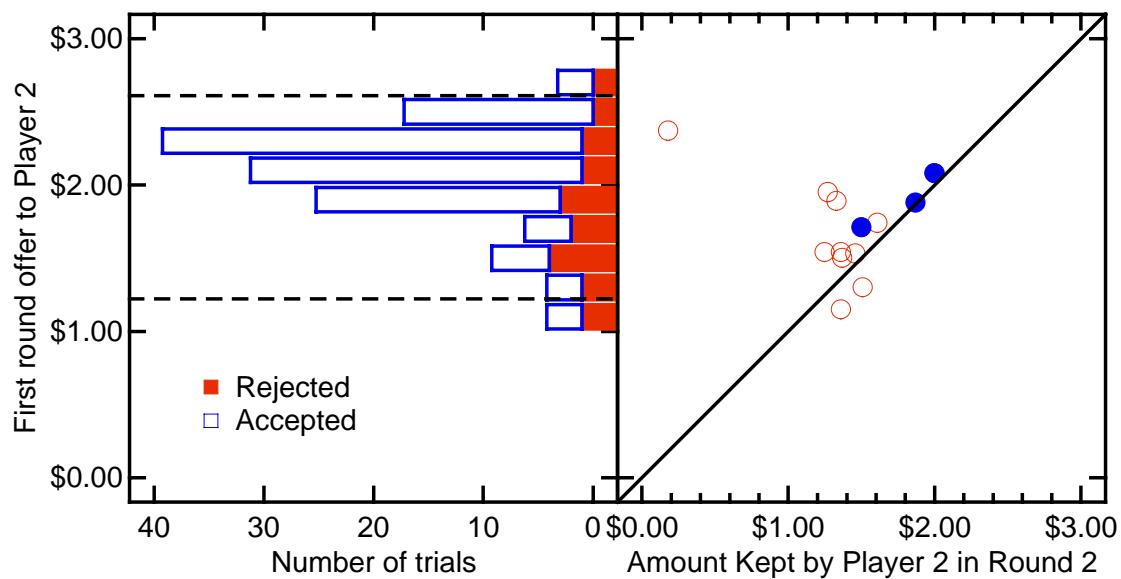
Figure 2: Distribution of offers to player 1 in round 1 and amount counter offered in round 2.
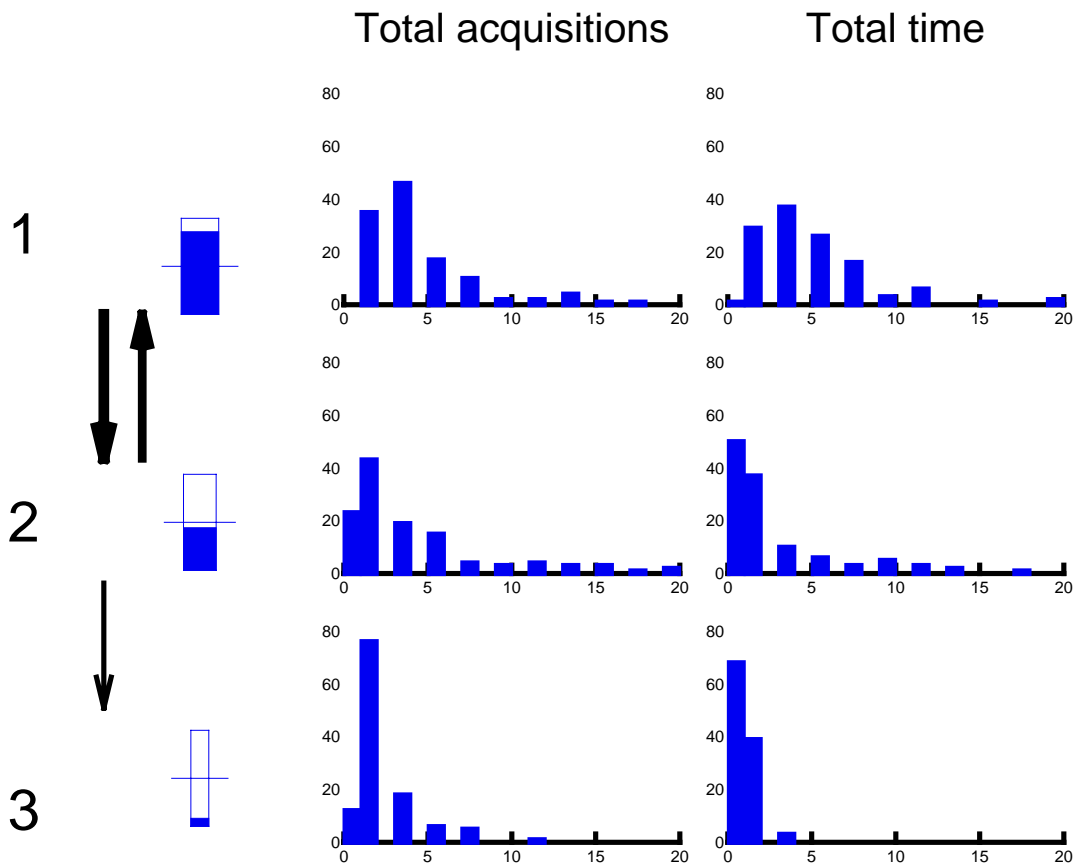
Figure 3 Icon graphs and histograms of acquisitions and looking time, study 1

Level 1   Level 0
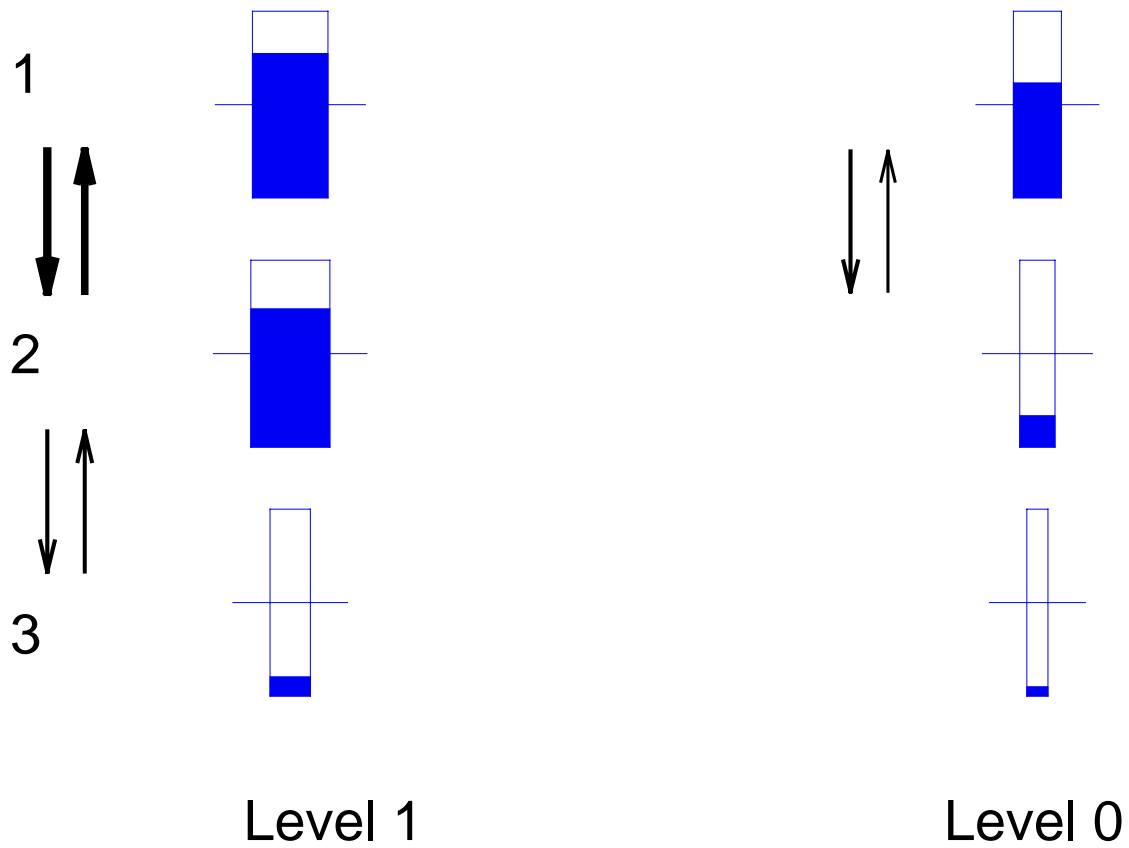
Figure 4 Icon graphs for player 1 by type inferred from first-round offers.
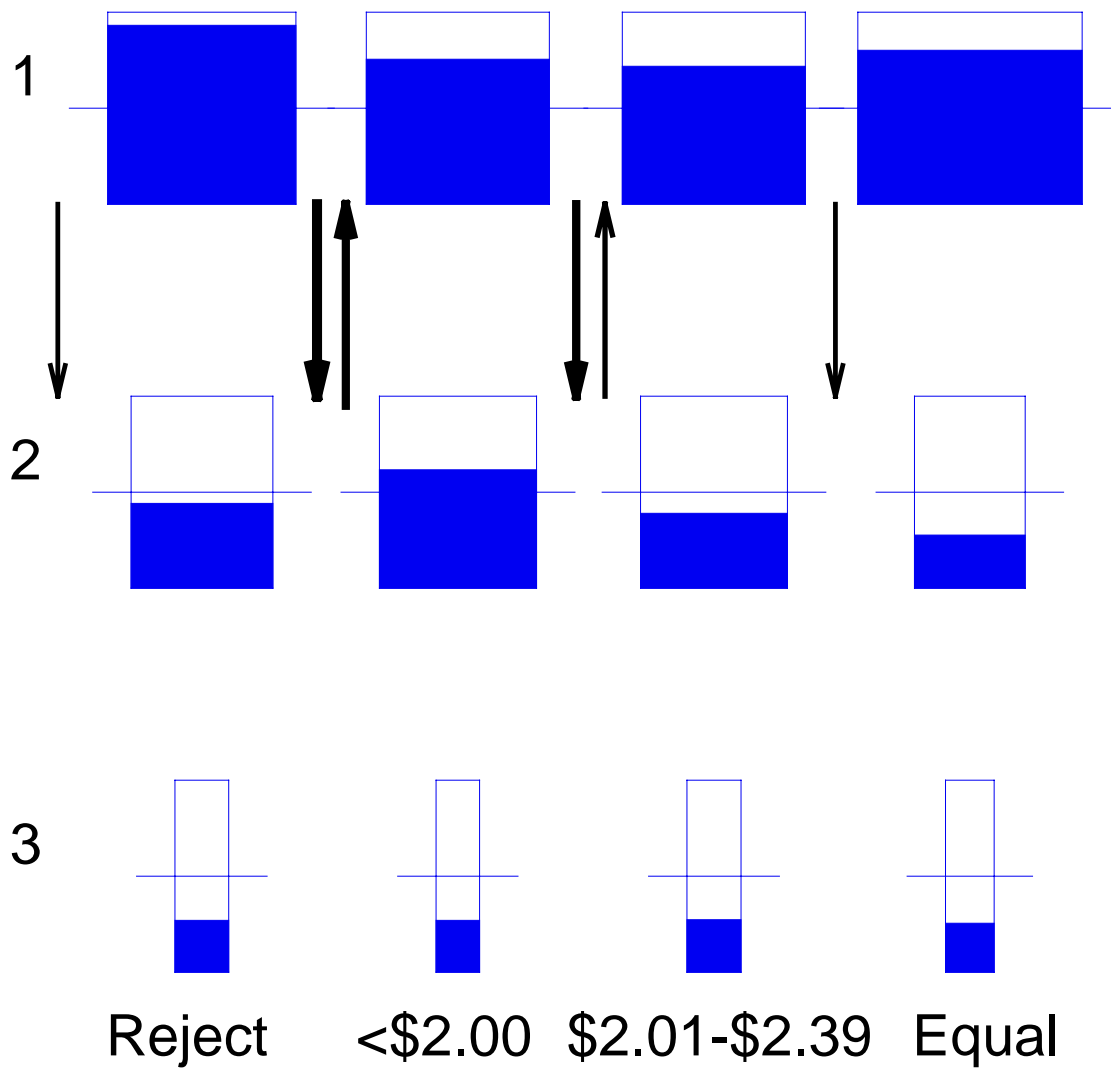
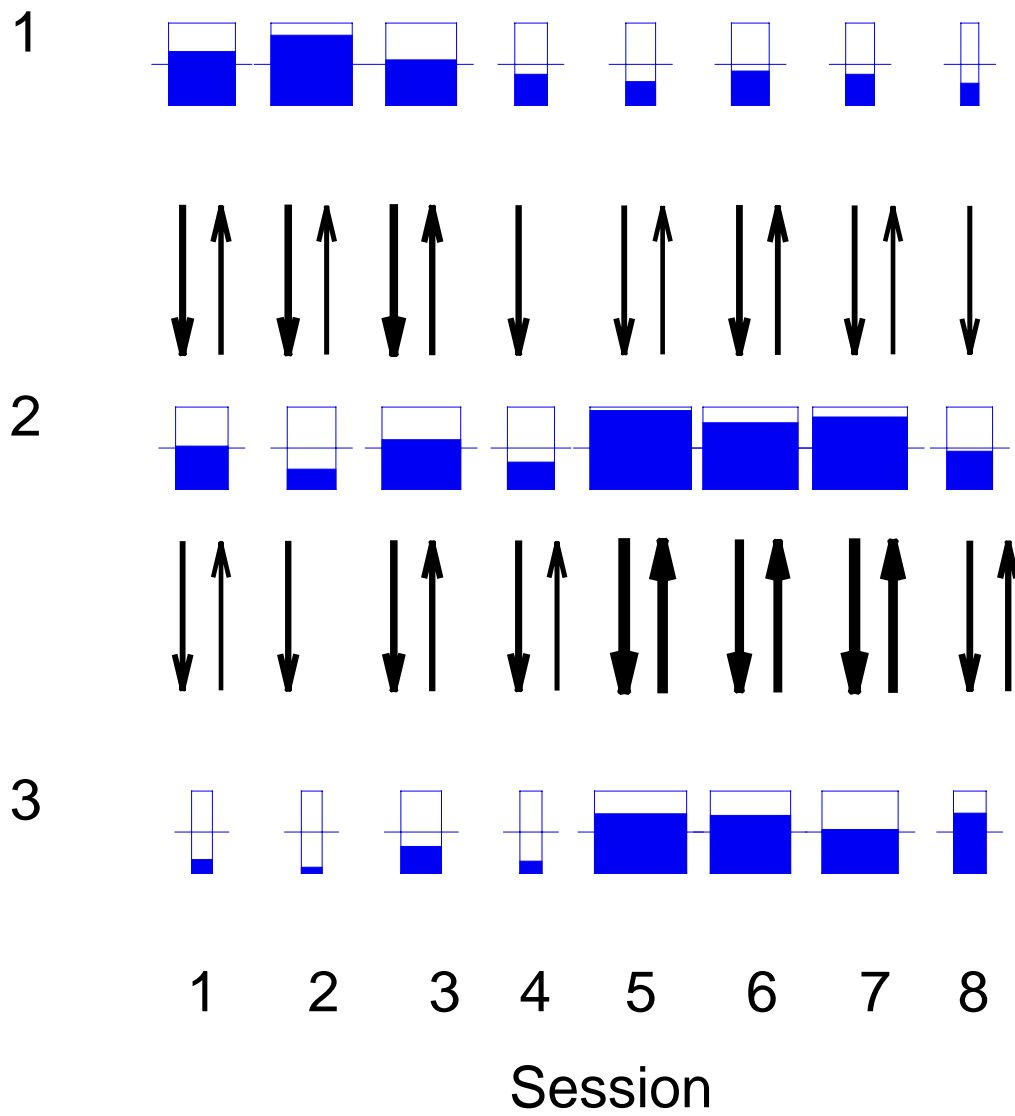**Figure 5 Icon graphs for player 2 by rejection or size of offer accepted**

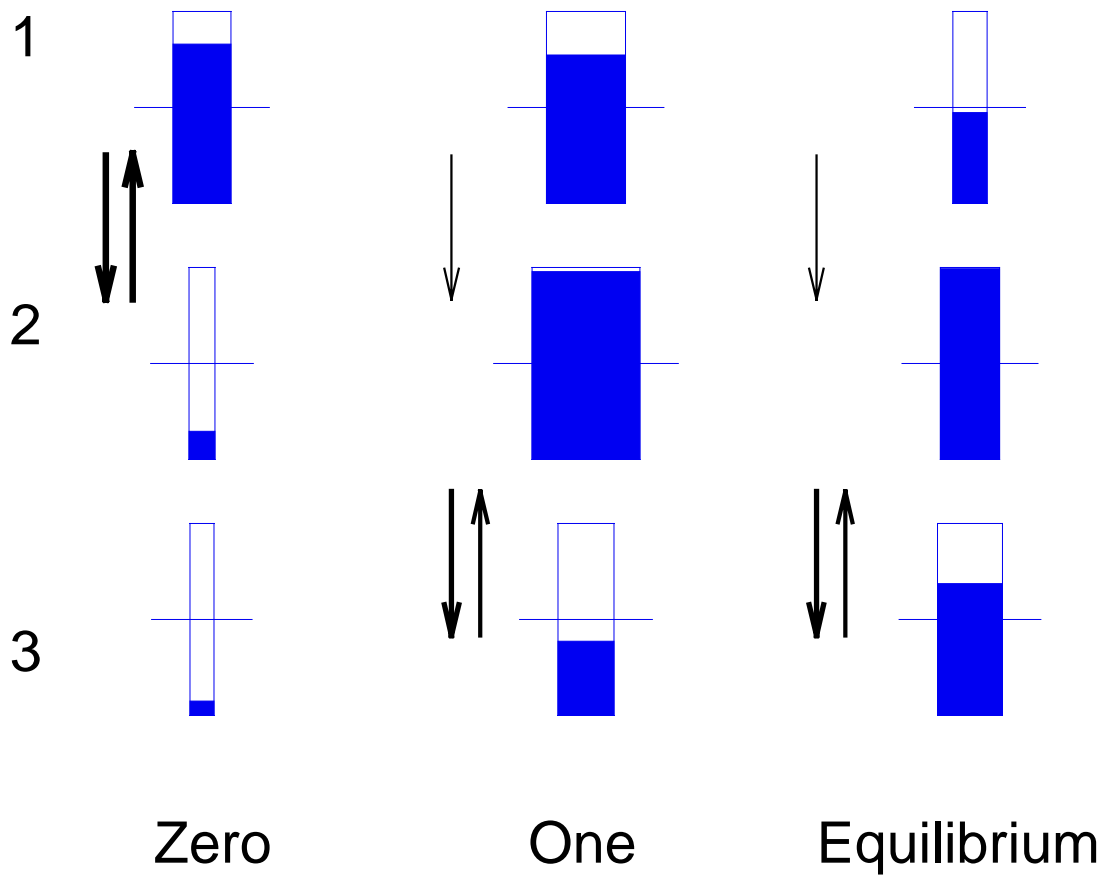**Figure 6 Changes in information processing before training (Sessions 1-4) and after (Sessions 5-8)**

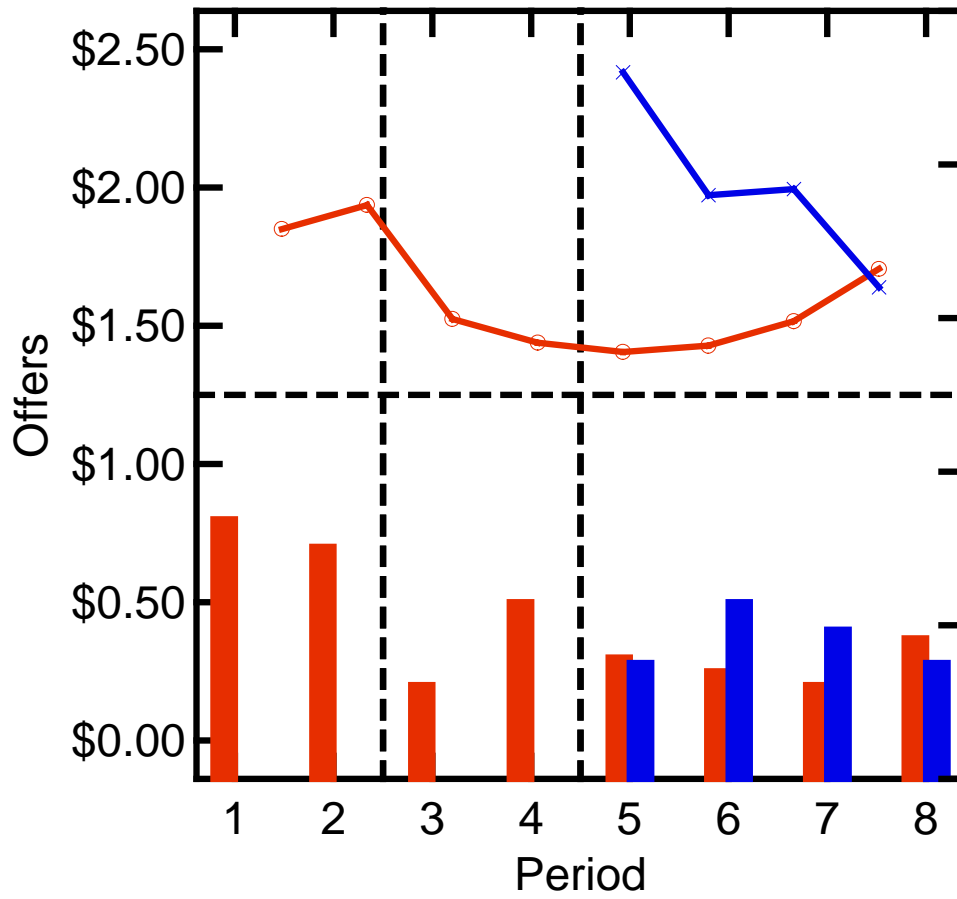**Figure 7 Icon graphs of information processing (time, number of acquisitions and transitions) by type.**

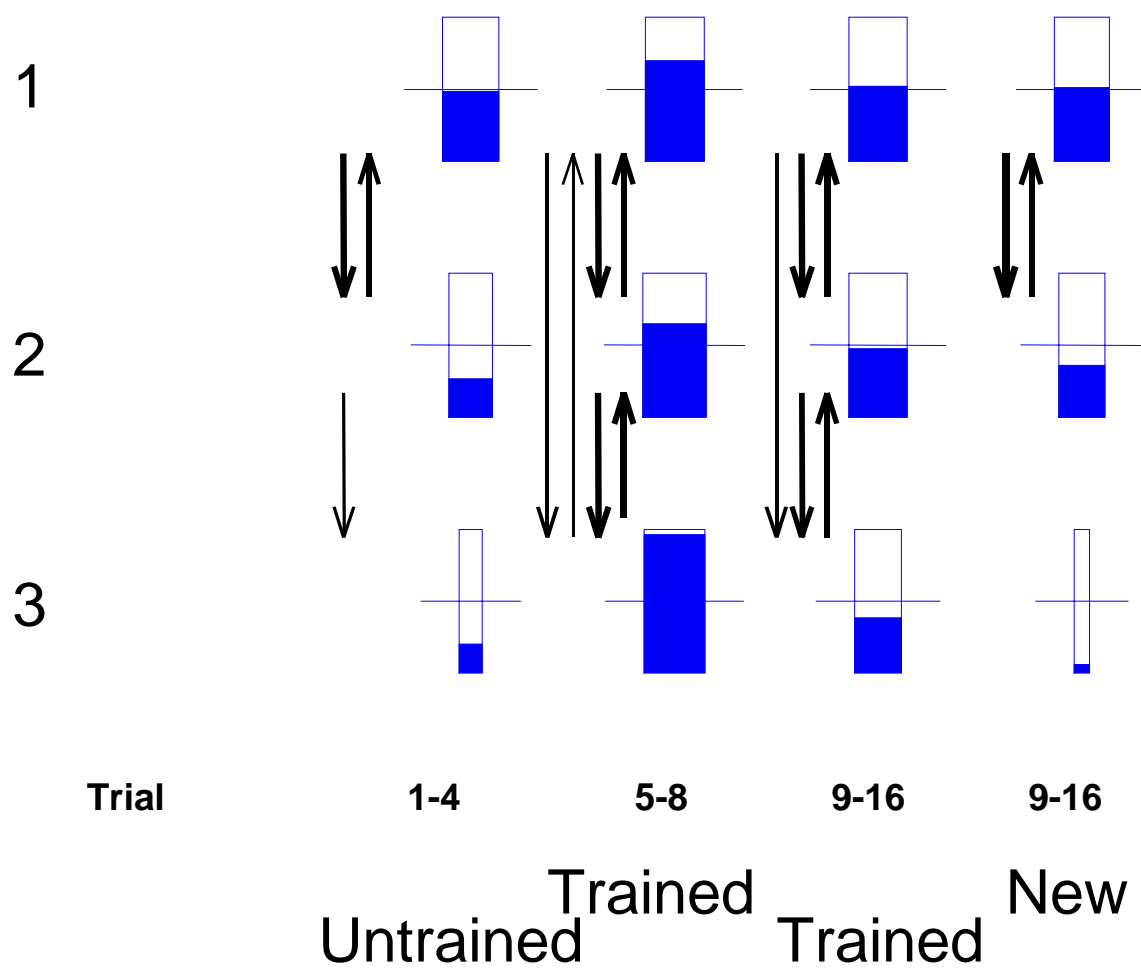**Figure 8 Mean offers and rejection proportions by period and player type, study 3**

Figure 9 Icon graph, looking time, acquisitions and transitions by training and session, study 3