BARGAINING OUTCOMES AS THE RESULT OF COORDINATED EXPECTATIONS: AN EXPERIMENTAL STUDY OF SEQUENTIAL BARGAINING

by

Jeffrey Carpenter

April, 2002

MIDDLEBURY COLLEGE ECONOMICS DISCUSSION PAPER NO. 02-04

DEPARTMENT OF ECONOMICS
MIDDLEBURY COLLEGE
MIDDLEBURY, VERMONT 05753

http://www.middlebury.edu/~econ

BARGAINING OUTCOMES AS THE RESULT OF COORDINATED EXPECTATIONS: AN EXPERIMENTAL STUDY OF SEQUENTIAL BARGAINING*

April 22, 2002

Jeffrey P. Carpenter
Department of Economics
Middlebury College
Middlebury, VT 05753
(jpc@middlebury.edu)

ABSTRACT: Experimental studies of two-person sequential bargaining demonstrate that the concept of subgame perfection is not a reliable point predictor of actual behavior. Alternative explanations argue that 1) fairness influences outcomes and 2) that bargainer expectations matter and are likely not to be coordinated at the outset. This paper examines the process by which bargainers in two-person dyads coordinate their expectations on a bargaining convention and how this convention is supported by the seemingly empty threat of rejecting positive but small subgame perfect offers. To organize the data from this experiment, we develop a Markov model of adaptive expectations and bounded rationality. The model predicts actual behavior quite closely.

JEL CLASSIFICATION NUMBERS: C72, C78, C91, D84

KEYWORDS: Sequential Bargaining, Experiment, Convention, Fairness, Finite Markov Chain, Bounded Rationality

---

* I am grateful to Herbert Gintis, Kevin McCabe and Vernon Smith for their comments on the experimental design. I would also like to thank Sam Bowles, Corinna Noelke, John Stranlund, and Dale Stahl. This project is funded by the National Science Foundation (SBR 9730332 and SES-CAREER 0092953).

If the phenomenon of "rational agreement" is fundamentally psychic - convergence of expectations - there is no presumption that mathematical game theory is essential to the process of reaching agreement, hence no basis for presuming that mathematics is a main source of inspiration in the convergence process. (Schelling [1960], pp.114)

1 INTRODUCTION

Experimental studies of two-person sequential bargaining have documented two behavioral regularities - subjects do not end up in subgame perfect equilibria (SGPE), even with experience, and observed outcomes typically diverge from the SGPE toward an equal split of the surplus.[1] Concentrating for now on the first result, it is clear that subgame perfection can only be supported if bargainers use SGPE strategies and expect that their opponents will also use the SGPE strategy. In this sense, subgame perfection requires common knowledge of the ability to do backward induction. Experiments have been conducted that indirectly test if subgame perfection is behaviorally important by analyzing whether subjects make backward induction calculations to inform decisions. These studies find little support for a general capacity to do multiple levels of backward induction or iterative dominance.[2] As such, it is unlikely that the key to understanding how subjects reason in sequential bargaining lies in the mechanics of backward induction.

On the other hand, expectations have proven to go a long way in explaining behavior.[3] Experiments that analyze expectations demonstrate that two agents will be able to maintain an efficient (conflict minimizing) convention only when they come to anticipate the response of their partner. In situations such as in the experimental lab that are devoid of the kind of social history that establishes prevailing behavioral conventions, subjects are likely to initially rely on social heuristics. Social heuristics are general behavioral rules developed outside the lab which, presumably, are shared by all participants from a common culture (Roth et al. [1991], Henrich [2000], and Henrich et al. [2001]). In this way, subjects rely on heuristic rules as a benchmark from which they explore alternative strategies within the setting of the experiment. The exploration process consists mainly of forming expectations about the future success of various strategies based on what has worked previously in the current population of bargaining partners. Put another way, bargaining conventions, and adaptive social norms in general, are important, not necessarily because they dictate behavior, but because they coordinate agents' expectations. This paper explores how

---

1 For a review of the ultimatum game see Camerer and Thaler [1995]. For sequential bargaining studies see Gueth and Tietz [1988], Neelin, Sonnenschein and Spiegel [1988], or Ochs and Roth [1989].

2 For experimental studies of backward induction see McKelvey and Palfrey [1992]. See Nagel [1995] and Ho, Camerer and Weigelt [1998] on iterative dominance.

3 Roth and Schoumaker [1983] demonstrate the importance of expectations in determining bargaining outcomes in an experiment using a two-stage version of the Nash demand game. Additionally, Harrison and McCabe [1996] demonstrate that, with exposure to subgame perfect play, subjects eventually coordinate on the SGPE. For other experimental studies where expectations are found to be important see Ochs [1995] on coordination problems and Sunder [1995] on speculative asset market bubbles.

the initial expectations of subjects formed early in bargaining can evolve into a convention.

Returning to the stylized fact that bargaining surpluses tend to be split equally, two theories that predict the equal split find empirical support. One theory states that bargainers who are placed in context-free laboratory rely on social heuristics founded in equity norms in one-shot games. A second theory predicts that bargainers make 'fair' offers because they are sure such offers will be accepted. A naive version of the first theory predicts that bargainers will maintain the expectations they bring to the lab which implicitly are that the surplus should be split equally. The second theory emphasizes the idea that subjects must experience the game and its incentive structure to form the expectations that will dictate which offer maximizes expected payoff.

Carpenter [2000] finds support for both theories, but more importantly for the current discussion, the results demonstrate that only egoists (those who behave as payoff maximizers) use expectations about their opponent's social orientation (e.g. altruistic) to inform proposing decisions. This supports a theory of bargaining outcomes (conventions) wherein some agents enter the lab with preexisting expectations of how to divide a pie, and others start with a blank slate, but form expectations in the process of bargaining. As such, conventions that develop in this setting are likely to be skewed toward the behavior of those who don't deviate from the heuristic rules they bring to the lab. The end result is a hybrid in the sense that participants with particularly salient heuristics dictate the benchmark and those who try various strategies to find one that maximizes payoffs cause the resulting convention to drift away from the benchmark. Additionally, it is reasonable to think that the amount of drift will be a function of the off-perfect equilibrium incentives of the bargaining institution.[4]

This paper reports on a repeated sequential bargaining experiment that supports this theory of bargaining conventions. A brief summary of the results is as follows. To begin, subgame perfection is not supported as a predictor of outcomes. Rather, a convention develops in the early stages of bargaining (despite various treatment conditions that alter the strategic incentives of the game) wherein the bargainer making a proposal gets 55% of the current pie. This convention is supported by the theoretically incredible threat to reject subgame perfect offers. Further, a model of adaptive expectations does well to explain the evolution of this convention, given the starting distribution of bargaining strategies.

2 The Experimental Design

The experiment described below is essentially a simplification of the design used by Harrison and McCabe [1992]. The Harrison and McCabe design gives subjects

[4]Such a simple theory can explain behavioral changes such as those documented in Prasnikar and Roth [1992] which highlight the difference in behavior when the off-perfect equilibrium incentives are changed from the ultimatum game where more generous proposers are rewarded to the best-shot game where generous first movers are taken advantage of.

repeated experience in both a three-stage sequential bargaining game and the two-stage subgame of the larger game. This design is used to show that experience in the subgame is enough to coordinate the expectations of subjects in a way that mimics the notion of common knowledge mentioned above. Given this design, the authors find support for subgame perfection as the limiting outcome when subjects gain sufficient experience, and therefore, have coordinated their expectations.

The current experiment exploits the elements of the Harrison & McCabe design that are useful for examining expectations (i.e. repeated experience in the subgame), but takes cues from experiments that have been run which demonstrate that experimenters should not expect subjects to accomplish more than two levels of backward induction (see footnote 2). To account for this stylized fact, the current experiment is a simplified (one and two round) version of the earlier study to assure subgame perfection is given a fair chance of working. Figure 1 summarizes the repeated sequential bargaining environment that subjects faced.

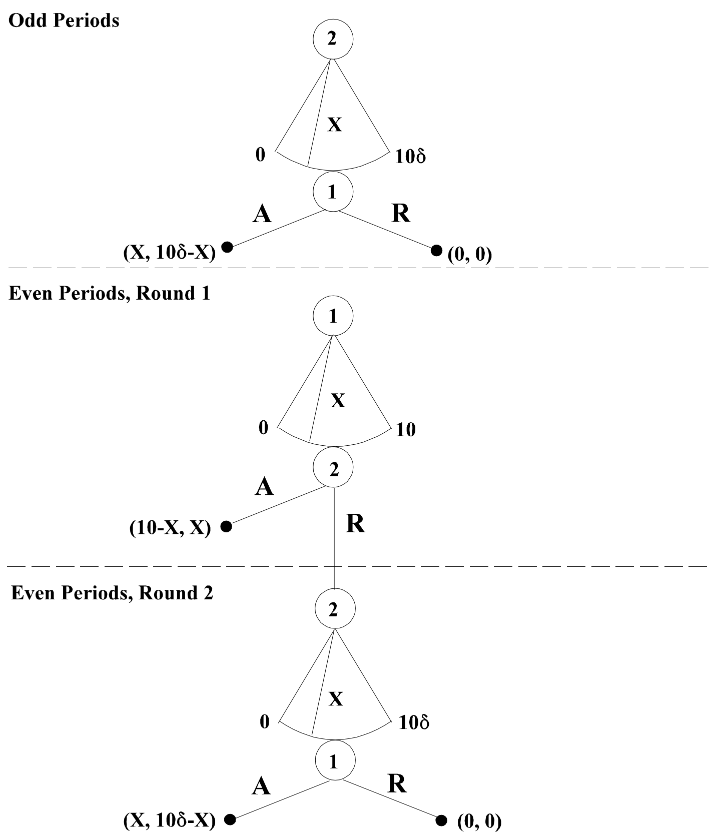

The Bargaining Institution

Subjects were randomly assigned to roles as either player one or player two. There were 15 periods of bargaining. All the odd periods were one-round ultimatum games and all the even periods were two-round games. In periods 1,3,...15 player two proposed a division, X, of a pie of size 10δ where $0<\delta<1$ (in the experiment δ=.25 or .75). Next, player one decided to either accept or reject this proposal. If the proposal was accepted, player one received X and player two received 10δ-X. If the proposal was rejected, both players received 0 for the period.

In periods 2,4,...14 player one made an opening offer over a pie that was initially 10 experimental francs. If player two accepted player one's proposal, the period was over, player two received X and 10-X went to player one. However, if player two rejected player one's round one proposal, then the two subjects moved to round two - the subgame - and the pie shrank to 10δ. By design, the resulting subgame was identical to the ultimatum game played in odd periods.

All participants were given a worksheet to assure that they understood the structure of bargaining. The worksheet required subjects to keep record of proposals sent, received and responses made for each period. As a result, the worksheet clearly laid out the structure of bargaining and therefore reinforced the consequences of moving to the subgame. Also, by filling out the worksheet, each subject had in front of him or her a complete history of prior proposals and responses. This was done to facilitate strategic thinking and give subgame perfection it's best shot.

The subgame perfect equilibrium for each period is calculated with the help of Figure 1. Start with the subgame played in each odd period. In addition, assume that bargainers have standard preferences for monetary outcomes and have common expectations that everyone else has similar preferences, then the SGPE outcome occurs where two offers one the smallest unit of account, ϵ, and one accepts because ϵ is better than nothing. This is true for each odd period and thus forms the expectation of what is likely to occur if bargaining in even periods moves to the subgame. Given this expectation about subgame play, in even periods one will offer two what two expects to receive if bargaining moves to the second round, namely 10δ. Faced with this offer, two accepts as she cannot possibly do better by rejecting and forcing the interaction to round 2. This pattern will repeat itself regardless of how subjects are matched.[5]

This bargaining environment was chosen for two reasons. First, repeated bargaining is used because it provides the kind of experience that might lead subjects to eventually experiment in the direction of the SGPE. Secondly, this institution

[5]Notice that even repeated bargaining between players matched with the same partner for the duration of the game cannot support any other Nash equilibrium. For example an early rejection by either one or two can only sustain a more favorable series of future proposals if the other player is uncertain about the preferences of the rejector. The assumption of common knowledge guarantees that any rejection is treated as a mistake and therefore the SGPE prediction is also robust with respect to 'trembling hands.' In addition, the game is finitely repeated which precludes any folk theorem results.

forces subjects, who otherwise would tend to settle in round one of the two-stage game, through the subgame. The reasoning behind this follows Harrison and McCabe [1992] who posit that pairs of subjects that agree in the first round of a sequential game do not necessarily have common expectations about acceptable outcomes because of a lack of experience in the subgame. Without subgame experience, these subjects have no way of forming expectations about what they will get out of the second round.

		Matching Rule	
		Same Pairing	Random Repairing
Discount Factor	$\delta=0.25$	sessions 1, 2, 3 15 bargaining pairs	sessions 8, 9, 10 17 bargaining pairs
	$\delta=0.75$	sessions 4, 5, 6, 7 17 bargaining pairs	sessions 11, 12, 13 16 bargaining pairs

Table 1 Experimental Design

Within the experiment two treatment variables were manipulated: the degree to which the pie shrank in the bargaining game, δ, and the rule that matched bargainers at the beginning of each period. Table 1 summarizes the design. For half of the sessions the discount factor was .25 favoring player one and for the other half the discount factor was .75 favoring player two.[6] Also for half the sessions subjects were matched with the same partner for all fifteen periods and for the other half, subjects were randomly rematched at the beginning of each period. The matching rule was explicitly mentioned in the instructions.

3 Experimental Results

A total of 13 sessions were run using undergraduate subjects at the University of Massachusetts. The entire experiment (from instructions to payment) lasted slightly less than an hour and the average earnings of a subject, including a five dollar show-up fee, was $14.43. To increase the number of pairs in each cell, the current results were pooled with those of a series of experiments run at the University of Arizona. There are two main differences in the experiments. The Arizona experiments were preceded by a preference revelation mechanism and were run for only 10 periods.[7] However, t-tests of mean behavior and Kolmogorov-Smirnov tests of distributional differences did not suggest that behavior was significantly different between

[6] Remember player two is expected to receive the lion's share of the subgame (ultimatum game) pie. Hence, as δ increases player two is the expected beneficiary.

[7] The preference revelation mechanism was essentially a series of bilateral dictator allocation decisions with variable pie sizes. Participants didn't know the outcome of the preference exercise until the end of the experiment For a more in depth analysis of the Arizona experiments see Carpenter [2000].

experiments. After pooling the data, there are 27 (δ=.25, same) pairs, 34 (δ=.25, random) pairs, 34 (δ=.75, same) pairs, and 31 (δ=.75, random) pairs for each of the first 10 periods, and the numbers listed in Table 1 for periods 11 through 15.

The most striking feature of the data is how stationary proposals are in both the ultimatum games and the two-stage games over time. Standard t-tests indicate that mean first period proposals are not significantly different from last period proposals for any of the eight sequences.[8] To illustrate this point Figure 2 plots the amount offered in the first round of the two-stage game played in even periods and proposals made in the ultimatum game played in odd periods and for each cell of the design. Open circles indicate average offers and the lines above and below indicate plus and minus one standard deviation. Further, solid horizontal lines indicate an even split of the surplus and dashed lines mark the SGPE predictions.[9]

Clearly these results are at odds with Harrison and McCabe [1992]. While Harrison and McCabe find that proposals steadily converge towards the subgame perfect prediction over the course of fourteen periods, our results, using a much simpler version of the game, show no movement in that direction. Two differences in the experiments may account for this difference. First, in the more complicated Harrison and McCabe game, the theoretical prediction in the two-stage subgame overlapped with an equal split of the surplus. This feature of the design might have made the theorized subgame result more salient in the minds of the participants compared to the very unequal theoretical outcome in the subgame of the current experiment (recall the subgame of the current experiment is an ultimatum game). If this was true, then it may have been easier for the majority of participants to form a common expectation about what would occur in the subgame. However, this explanation is somewhat unsatisfactory because it is reasonable to think that subjects who rely on fairness to coordinate expectations in the two-round subgame would also do so when making offers in the three round game. An alternative explanation might have something to do with the worksheet that participants filled out during the current experiment. Instead of eliciting strategic behavior by stressing the structure of play and the consequences of rejecting an initial offer, the worksheet might have reinforced fair play because participants often looked at a history of fair offers when recording their responses.

Although the current results differ from Harrison and McCabe [1992], it would be a mistake to conclude that expectations do not matter. Instead, the difference in these two experiments only questions whether strict fairness norms (i.e. the equal split) can easily be displaced by experience. As we will show below, expectations might also account for the current results. In the next section we build a model to argue that proposers quickly formed expectations that were maintained over the course of the experiment and that these expectations did not drive

[8] All tests were two-tailed and no differences were found at the 5% level.

[9] For all ultimatum games the SGPE is zero, for δ=.25 the first round prediction is 2.50F, and for δ=.75, the prediction is 7.50F.

proposals to the levels predicted by subgame perfection nor did they reflect strict fairness.[10] Instead, we see the following hybrid - offers hover just below the equal split reflecting a slight advantage for the proposer.[11] This result is true for both the one-stage ultimatum games and the two-stage shrinking pie games.

It appears that proposers are somewhat affected by the discount factor in the two-stage game. Returning to Figure 2, when δ is .25 and player one (the initial proposer) has the advantage, first round proposals are slightly less than when δ is .75 and player two has the advantage. This supports the hypothesis developed in (Gueth and Tietz [1990]) that players one hide behind fairness and do not offer more than half when the discount factor does not benefit them. We will return to this below when the treatment effects are analyzed more fully. Overall, however, Figure 2 demonstrates that the results can be summarized by a convention wherein the bargainer who is currently proposing gets 55% of the pie and the responder gets the remainder.

{Figure 2 here}

Ochs and Roth [1989] coin the term *disadvantageous counterproposals* (DACs) for second round proposals that leave the proposer with less than they would have gotten if they had accepted the first round offer. They assert that this is an important phenomenon because it describes 81 percent of the rejections found in their study and is also not predicted by subgame perfection. The results of the current experiment reveal a substantial amount of DACs (36 percent) although not as many as in the Ochs & Roth study.[12] Rejected offers, in relative terms so that the data could be pooled across different values of δ, and DACs that followed rejected offers are illustrated in the left panel of Figure 3.

{Figure 3 here}

Figure 3 should be read as follows. The darker region in both panels is a histogram of offers (in relative terms) that were rejected. The lighter region in the left panel is a histogram of relative offers that were rejected and followed by DACs. This panel illustrates the fact that a significant number of relatively fair offers are rejected and followed by DACs. The lighter region in the right panel is a histogram of DACs but in *relative terms* (i.e. counterproposals that gave player two less in relative

[10] One-tailed t-tests reject the hypothesis that any sequence approached subgame perfection.

[11] In all but two instances (δ=.25, same - ultimatum) and (δ=.25, random - ultimatum), one-tailed t-tests demonstrate that offers, pooled across periods, are significantly below the equal split even though in many instances the equal split is continuously within one standard deviation of the mean.

[12] The reason that the amount of DACs is lower in the current study is probably because of the discount factors used. Virtually all counterproposals in the δ=.25 sessions are disadvantageous. However, many of the δ=.75 counterproposals are not disadvantageous because by rejecting a low offer in round one, a player can still receive a substantial amount in round two. By comparison, Ochs and Roth [1989] used δs of .4 and .6.

terms than the first round proposal that was rejected). Bolton [1991] explains DACs by arguing that subjects value both absolute and relative outcomes. That is, a subject may reject a round one proposal and counter by asking for less in absolute terms, however the resulting counterproposal may give the proposer a larger share of the smaller pie. This explanation is supported by the right panel of Figure 3. Here the small lighter histogram illustrates the fact that only 2 of the 105 DACs are disadvantageous in relative terms. That is, in 98 percent of the cases subjects rejected offers and countered by asking for more in relative terms.

An analysis of the actual counterproposals suggests two mechanisms that support the evolution of the 55/45 convention. As just mentioned, bargainers reject offers that are small in relative terms and counter with proposals that ask for a larger relative share. However, how much more (in relative terms) do they ask for? Basically, there are two types of counterproposals. In the first case, player two rejects a low first round offer and counters by returning to the 55/45 convention (i.e. asks for 55% of the second round pie). The second type escalates the aggressive offer of player one by rejecting and responding with an even lower counteroffer. This second type is sort of a tit-for-tat player who punishes departures from the convention by escalating the deviation. The first type simply returns to the convention when deviations are encountered.[13]

The size of the second round pie influences the distribution of these two types of counterproposers in the population. When the pie shrinks a lot between round one and round two (δ=.25), many more player twos respond to low proposals by returning to the norm. However, when the pie does not shrink much (δ=.75), more player twos escalate the deviation from the norm. The treatment effect of δ on counterproposals is demonstrated in the left panel of Figure 4 (in Figures 4 and 5 the first letter in the abbreviation stands for the matching rule, Same or Random, the number corresponds to the discount factor, the abbreviation SR represents the first round offer of the Ståhl-Rubinstein game played in all even periods and UG is obvious). In the δ=.25 case we see that counterproposals are indistinguishable from odd period ultimatum offers.[14] However when δ=.75, counterproposals are on average less than their odd period counterparts which suggests the escalation of low offers in this case.[15]

The heterogeneity in counterproposing behavior is better demonstrated in the scatter plot in the right panel of Figure 4 which plots rejected offers in relative terms against the fraction of the round two pie that was counterproposed. Here we see a large mass of data below the 45 degree line perhaps reflecting the inequality

---

[13]Note, these results are consistent with many distribution and inequality based models of preferences. Examples include Fehr and Schmidt [1999], Bolton and Ockenfels [1999] and Falk and Fischbacher [1998].

[14]Counter proposals are not different from the ultimatum offers in the (δ=.25, same) treatment (t=1.02, p=.15) nor are they different from the (δ=.25, random) treatment (t=1.41, p=.08).

[15]Here, counterproposals are marginally different from the (δ=.75, random) treatment (t=-1.76, p=.04) and highly significantly different from the (δ=.75, same) treatment (t=-5.29, p=0).

aversion of some subjects.[16] Also note the mass of observations between the 40 and 50% counterproposal which illustrates those players who return to the convention. However, we also see a considerable amount of tit-for-tat play to the northwest of the 45 degree line.

The last thing to notice about couterproposals is that they are more likely to be rejected than first round proposals (34% of counters are rejected versus 21% of first round proposals). This observation also suggests that the punishment implicit in a first round rejection is often escalated. Controlling for offer size, pie size, and the subject pairing rule, the following random effects logit regression shows that offers are significantly ($p<0.01$) more likely to be accepted in the odd period ultimatum games than in the second round of the even period Ståhl-Rubinstein games.

$$\underset{}{Response\ (accept=1)} = \underset{(0.41)}{0.75} + \underset{(0.25)}{2.01}Offer - \underset{(0.11)}{0.78}Pie - \underset{(0.28)}{0.20}Pairing + \underset{(0.32)}{1.57}UG$$

Wald $\chi^2=77.06,\ p<0.01$

{Figure 4 here}

Turning to a discussion of the treatment effects, we see that varying the discount factor has a strong effect on behavior while manipulating the matching rule influences behavior to a lesser extent. Figure 5 illustrates the differences between treatments by plotting all eight sequences of average offers by period. One can see the striking difference in first round behavior between the $\delta=.25$ and $\delta=.75$ treatments (compare S25SR to S75SR and R25SR to R75SR).[17] Regardless of the pairing rule, the average offers of the $\delta=.75$ treatment always lie above their $\delta=.25$ counterparts and continue to separate as the experiment progresses. This phenomenon further supports the idea that proposers push harder on the other player when the discount factor favors them, but hide behind fairness when they are in the theoretically disadvantaged position.

The pairing rule has less effect. There were significant differences in behavior in only two instances. First and surprisingly, first round offers in the two-round game ($\delta=.75$) were significantly higher when bargainers were randomly repaired after every round than when bargainers stayed with the same partner for the duration of the experiment.[18] Secondly, and as one would expect if repeat interaction supports sharing, offers in the odd period ultimatum games where $\delta=.75$ were significantly higher when bargainers were paired with the same partner for the entire

[16] For a general theory of inequality aversion see Fehr and Schmidt [1999].
[17] Kolmogorov-Smirnov tests show that the differences between the distributions of offers for both pairing conditions are significant at any level.
[18] KS=.1545, p=.018.

experiment than when they were reshuffled after every period.[19] In the other two cases (first-round proposals where δ=.25, ultimatum proposals where δ=.25) there was no significant effect of the matching rule.

{Figure 5 here}

4 Coordinated Expectations and Bargaining Outcomes

The lack of variation in proposals across treatments with respect to time suggests that subjects quickly agreed on a convention that guided proposing behavior. This convention, though obviously linked to fairness, evolved away from the equal split in response to the institutional rules of the experiment.[20] Additionally, we have seen that the rejection behavior of subjects works to support this convention by punishing deviations from the established norm and by returning to it with counterproposals. However, as an explanation of the evolution of the 55/45 convention, it is still conjecture to say that the driving force is coordinated expectations. In this section, we will build and discuss a model of adaptive expectations. The purpose of the model is to illustrate how, given the initial expectations of subjects and despite the fact that interaction occurred in dyads, the most likely convention to evolve based on the decentralized flow of information between subjects is the one observed - proposers get 55% of the pie.

The structure of the model is adapted from a discussion of the evolution of conventions in Young [1993], Young [1998] and Gintis [2000]. More specifically, we will first develop a deterministic model of adaptive expectations in which a pseudo-Markov transition matrix is constructed based on best-reply dynamics applied to initial expectations about the success of three particular proposing strategies. These three strategies organize all of the first proposals made by participants. For what follows, *bargaining conventions* will be defined broadly as the states of the resulting dynamic system which demonstrate the most attracting power. In the initial model conventions will simply be absorbing states. When we complicate the model later on, conventions in the resulting ergodic system will be the states in which the system spends most of its time.

To begin with, we define three proposing strategies. The first will be called ***Low***, ***L***, and will be defined as making a proposal, x to your counterpart for less than 40% of the pie (i.e. $L=\{x|x<.4(current\ pie)\}$). Notice that in the ultimatum game (periods 1,3,...15) and in even periods when the discount factor is .25, the Low proposing strategy includes behavior that would be expected from subjects offering

[19] KS=.2451, p=0.

[20] Note however, the resulting convention only moves in the direction of the SGPE in the case of the low discount factor. In the high discount factor treatments, the SPGE requires the proposer to offer more than half the pie in two-round games. This suggests that, instead of responding to the theoretical structure of the game, participants associated bargaining power with the role of proposing.

the SGPE. Considering first proposals across treatments (i.e. considering only period one in the case of player twos and period two in the case of player ones) the strategy L accounts for 17% of proposals. The second strategy will be called ***Half***, ***H***, which is played by proposing at least half the current pie to your opponent *(H={x | x ≥.5(current pie)})*. In the first two periods H was played by 49% of player twos and by 47% of player ones. Lastly, define the strategy ***C*** as the ***bargaining convention*** where proposers get approximately 55% of the current pie *(C = {x | .4<x<.5 of the current pie})*. Strategy C was played by 36% of player ones and by 35% of player twos when making their first proposals.

Using these three strategies to organize the data on proposals for the first two periods, we can calculate the average (and expected) payoff of each strategy. We use the data from the first two periods only because we are interested in how behavior and expectations adapt to the results of initial play. Figure 6 presents the expected payoff to each strategy when it meets a proposer who makes a proposal of either L, H or C.[21] The entries of Figure 6 are in relative terms. We will continue to speak in terms of relative proposals so that the data from all four cells can be pooled and because the development of the observed convention appears to be independent of the treatment conditions.

Define an ***expectational equilibrium*** as a Nash equilibrium of the game played between prosposers where the payoffs are expressed as expectations. Clearly, for player one proposing half is dominated in expected payoff by playing the convention. Upon further examination, one can see that there are two expectational equilibria based on play in the first two periods. Both are symmetric and occur where all proposers coordinate on either proposing low or all play the convention. Referring to the definition of Nash equilibria, it should be clear that the expectational equilibria in Figure 6 will be absorbing states in our model of adaptive expectations. In other words, once bargainers transit to one of the equilibrium states, LL or CC (to be read Player one's strategy, Player two's strategy), their expectations will be coordinated in that they have no incentive to play anything else. The question then becomes, which equilibrium are subjects most likely to end up at.

[21]The payoffs in Figure 6 were calculated as the average payoff in period 1 (period 2) for player 2 (player 1) of the encounters between player 2 who played L, H or C when matched with a player 1 who subsequently played L, H or C on the next round and between player 1 who played L, H or C when matched with a player 2 who previously played L, H or C on the round before.

		Player 2		
		L	*H*	*C*
Player 1	*L*	*.52, .67*	*.48, .49*	*.49, .58*
	H	*.42, .43*	*.51, .47*	*.48, .43*
	C	*.50, .32*	*.52, .46*	*.53, .56*

Figure 6 Expected Payoff to Each Strategy

The dynamic we will use is best reply based on one period of recall. Hence, the transition matrix will be deterministic and not stochastic, but as a first step in the analysis we are interested to see where best reply dynamics will drive the data. More specifically, we assume players only remember the last proposal they made and its payoff. Therefore, on average, the dynamics of the population of bargainers can be represented by two agents who play their best reply based on the expected payoffs in Figure 6 and their last encounter. For example, because LL is an absorbing state, the best reply for player one of ending up in the LL cell is to continue to play L, likewise for player two. Similarly, the best reply of player one who's last period outcome was HH is to play C, while player two will stick with H.

We develop a pseudo-Markovian transition matrix for the current model by calculating the probability of transiting from one state to another of the game illustrated in Figure 6. Because we use the simple best reply dynamic and because none of the expected payoffs that need to be compared are equal, the transition probabilities are either 0 or 1 (i.e. there is always a unique best reply to the stated history). Figure 7 presents the best reply transition matrix, *E*, that has been constructed based on the expectations of subjects in the experiment as defined by the average payoff from first proposals.

$$E = \begin{array}{c} \\ LL \\ LH \\ LC \\ HL \\ HH \\ HC \\ CL \\ CH \\ CC \end{array} \begin{array}{c} \begin{array}{ccccccccc} LL & LH & LC & HL & HH & HC & CL & CH & CC \end{array} \\ \left[\begin{array}{ccccccccc} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{array}\right] \end{array}$$

Figure 7 Expectational Transition Matrix

Given we assume that expectations adapt according to the best reply dynamic used to create E, the predicted distribution of states in the second period of the experiment is calculated by multiplying the distribution of starting states, call it S_0, by the transition matrix, E. Effectively, this calculates the best reply of the starting population distribution to the expected payoffs of each strategy in the game depicted in Figure 6. Likewise, to calculate the expected distribution of states in period 3 we would multiply S_0 by E^2. In general, to calculate the distribution of states after n periods we find S_0E^n.

Because the transition matrix has absorbing states (which conform to the Nash equilibria of Figure 6), the process of expectation adaptation is likely to be absorbed by either LL or CC. The question that we are interested in is which state is more likely to become a convention? In other words, what is the long run behavior of the system? To examine limiting behavior of our model of expectation coordination we calculate

$$E^* = \lim_{n \to \infty} E^n$$

In the present case, $E^* = E^k$, $\forall\ k \in \mathbb{N}$ and

$$E^* = \begin{array}{c|ccccccccc|} & LL & LH & LC & HL & HH & HC & CL & CH & CC \\ \hline LL & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ LH & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ LC & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ HL & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ HH & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ HC & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ CL & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ CH & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ CC & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{array}$$

As we see, if the best reply dynamic that underlies E represents the process of expectation coordination, then those pairs which start in state LL will remain there, and those starting in states LH and LC will be absorbed by state LC where player one plays the convention and player two plays low. Further, pairs starting in states HL and Cl will be absorbed by CL and finally, pairs starting in states HH, HC, CH or CC will converge on the convention observed in the data.

Returning to the data on proposals from periods 1 and 2 of the experiment, we find that S_0 = (.06, .06, .05, .06, .27, .13, .06, .12, .19). Multiplying S_0 by the long-run steady state of the expectation adaptation process, E^* we have the predicted end state distribution, S_n = (.06, 0, .11, 0, 0, 0, .12, 0, .71). In other words, the model we have developed predicts that 71 percent of subjects will be in dyads that play the convention. Overall, if the model accurately predicts the process of expectation coordination, 82 percent of the bargainers will be playing the convention, 18 percent will play Low and no one will play Half.

Now let’s compare the predicted distribution of bargaining strategies to the actual end state of the experiment. In period 14 players one were distributed as follows: 16 percent Low, 31 percent Half, and 53 percent Convention. In period 15 players two were distributed 25 percent playing Low, 25 percent playing Half and 50 percent playing the convention. On average, pooling ones and twos, we find 21 percent Low, 28 percent Half, and 51 percent played the convention. Therefore, while the model is not a bad predictor, it over-predicts the absorbing power of the convention and under-predicts the attracting power of proposing half (i.e. it underpredicts the resiliency of the expectations of players who subscribe to an equity norm).

To try to understand the difference between the predicted distribution of proposing strategies and the actual distribution, we can explore relaxing the assumptions we have made about the dynamic that underlies the expectational transition matrix, E. In particular, it is obviously incorrect to assume that bargainers can make best replies when they do not have the information to calculate the expected payoffs for each cell of Figure 6. Therefore, we can examine another more boundedly rational dynamic. Rather than assuming that agents make best replies to their first proposals, we will now assume that they sometimes make errors because they don't know what the best reply is or they can't identify it from their limited sample of play. That is, now let us assume that bargainers are only able to make best replies $(1\text{-}e)$ of the time and with probability, e, bargainers choose their second best reply. To justify this change, we might imagine that because players meet one responder at a time they have not been exposed to enough responders to form the payoff expectations inherent in Figure 6. As a result, with imperfect information, they are forced to make boundedly rational strategy choices. For example, return to Figure 6 where the best reply to ending in state HH for player one is to play C. Now if player one does not experience playing C against H, then the best reply is to stick with H. Transforming the matrix E by incorporating the probability, e of playing the second best reply we get what we call the *Second Best Reply* transition matrix,

$$
E_e = \begin{array}{c} \\ LL \\ LH \\ LC \\ HL \\ HH \\ HC \\ CL \\ CH \\ CC \end{array}
\begin{array}{ccccccccc}
LL & LH & LC & HL & HH & HC & CL & CH & CC \\
(1-e)^2 & 0 & (1-e)e & 0 & 0 & 0 & e(1-e) & 0 & e^2 \\
0 & 0 & 0 & e(1-e) & 0 & e^2 & (1-e)^2 & 0 & (1-e)e \\
e(1-e) & 0 & e^2 & 0 & 0 & 0 & (1-e)^2 & 0 & (1-e)e \\
(1-e)e/2 & (1-e)^2 & (1-e)e/2 & 0 & 0 & 0 & e^2/2 & e(1-e) & e^2/2 \\
0 & 0 & 0 & e^2/2 & e(1-e) & e^2/2 & (1-e)e/2 & (1-e)^2 & (1-e)e/2 \\
e^2/2 & e(1-e) & e^2/2 & (1-e)e/2 & 0 & (1-e)e/2 & 0 & (1-e)^2 & 0 \\
0 & (1-e)e & (1-e)^2 & 0 & 0 & 0 & e(1-e) & e^2 & 0 \\
0 & 0 & 0 & 0 & e^2 & e(1-e) & 0 & (1-e)e & (1-e)^2 \\
0 & e^2 & e(1-e) & 0 & 0 & 0 & 0 & (1-e)e & (1-e)^2
\end{array}
$$

Because there is a positive probability of moving from any state to any other state (not necessarily in one move), E_e is ergodic and we can calculate the predicted long-run distribution directly by finding E_e^* where

$$E_e^* = \lim_{n \to \infty} E_e^n$$

The matrix E_e^* arises after seven iterations and is a 9 by 9 matrix of the same 9 row vectors. The expectational equilibrium of the second best reply dynamic is described by this row vector. The first thing to note is that E_e^* converges quickly as does E^*. In fact, both matrices, remarkably, converge within the time frame of the experiment (i.e. ≤ 15 time periods). Figure 8 illustrates the equilibrium predicted distribution of states for three levels of e. It is clear from Figure 8 that the frequency of bargainers playing Half increases as *e* increases and therefore, as bargainers become more boundedly rational they are more likely to continue playing Half.

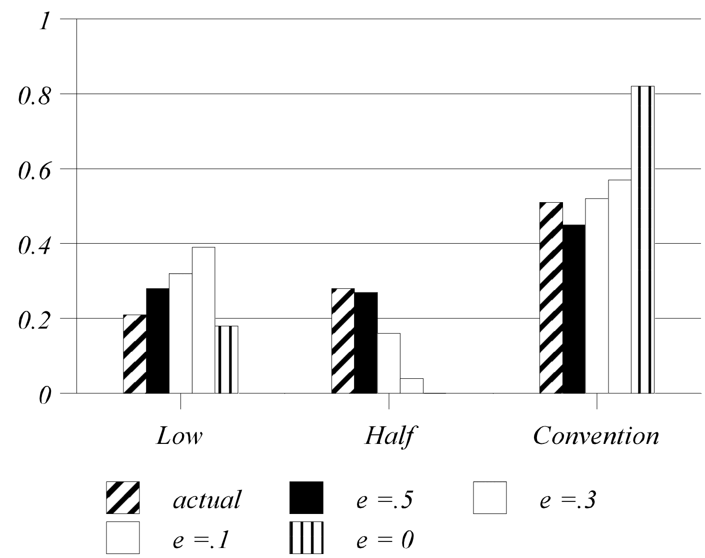

Figure 8 Predicted vs. Actual State Distributions

As Figure 8 demonstrates, the model fits better as we relax the assumption that bargainers play best responses to the expected payoff of the three strategies. When the probability that agents make second best responses rather than best responses is .5, our model of coordinated expectations fits rather closely to the behavior observed in the lab. Summarizing, this suggests is that bargaining conventions may arise as epiphenomenon of the actions of decentralized bargainers who make best responses to available information given they may feel some gravity towards preexisting heuristics. Moreover, it is reasonable to think of bargainers as locally optimizing in that they make best reply comparisons of current payoffs to the outcome of past negotiations. In the process of doing so, a convention arises that is founded on both existing norms (e.g. fairness) and the coordinated expectations that evolve as a by-product of making boundedly rational comparisons within a specific institution.

5 DISCUSSION

Figure 9 is a nice summary of the data generated by the current experiment. Figure 9 reworks Figure 2 so that the vertical axis now measures relative proposals. Plotting relative proposals in both odd period ultimatum games and even period two-round games on the same graph clearly illustrates the convention of the proposer (regardless of player number) getting a little more than half. The sequences of proposals plotted in Figure 9 show that, regardless of the treatment variables which have a slight effect on the level of the convention, proposals are flat with respect to time. This regularity and the corresponding reduction of variance in proposals is evidence that expectations have stabilized. Also, the fact that relative proposals overlap so tightly suggests that the convention resulting from coordinated expectations is established in terms of the proposer's share of the given pie, rather than in absolute payoffs.

{Figure 9 here}

While Figure 9 illustrates that proposals have stabilized, for expectations to be truly coordinated we would also anticipate that the rate of rejection would diminish over the course of the experiment. This would occur as participants started to sort out an agreeable allocation for each role. Figure 10 plots the rejection rates for each role in the experiment. The first thing to notice is that the rejection rates for the two roles cycle starting relatively low, increasing and then falling again. Also notice the cycles are staggered by one period. After player ones reject more (less) player twos increase (decrease) their likelihood of rejecting in the next period. The path dependant nature of this cycle continues over the course of the experiment demonstrating that spite might be affecting responses (recall the regression results). If it is spite causing the cycles in responses, it also prevents the rejection rates from systematically diminishing over time. The figure leads to an interesting hypothesis worthy of further study. Namely, if bargainers coordinate their expectations about who should get what in a particular bargaining institution, then we would expect that by the end of a series of negotiations most offers would be accepted. However, just as general norms of fairness might influence the direction and speed of convergence to a convention within a specific institution, spite might interfere and hinder the process. Here spite, triggered by having one's last offer rejected, disrupts convergence by causing bargainers to reject offers that might have otherwise been acceptable.

{Figure 10 here}

Overall, the current results seem to be driven both by distributional concerns and by expectations formed early in the game. As a result, a nonadaptive, norm-driven explanation such as that originally offered in Gueth, Schmittberger and Schwarz [1982] can not fully explain these results because it does not predict the slight but robust deviation from the equal split. At the same time however, an explanation based solely on eventually attaining common expectations which reflect the strategic incentives built into the game are also unable to explain these results. Such an approach can not account for the fact that in this experiment participants never reach the SGPE as they did in Harrison and McCabe [1992]. Instead, both explanations seem to be partially true. Hence, a reasonably parsimonious model of the current data would posit agents who enter the experiment with prior expectations that are initially anchored to a distributional norm, but adapt to the current institutional arrangement (i.e. the incentives and rules of the game) and the history of play.[22]

As a first attempt at creating such a model this paper has developed a simple model of adaptive expectations that has been calibrated by the expected payoffs faced by bargainers early in the experiment. Overall, the model approximates the behavior we see in the lab in that it predicts expectational equilibria that arise on a time scale similar to the experiment length. However, the fit was drastically improved by weakening the assumptions underlying the best reply dynamic used. The altered model allows for boundedly rational agents who make best replies to available information. The main contribution of the experiment and the analysis presented herein is that we have provided evidence supporting the idea that, more than the underlying logic of strategic interaction, initial conditions and expectations based on the history of play are the driving force behind bargaining outcomes.

Admittedly the model presented above is just a first attempt at creating a dynamic explanation of bargaining conventions based on experimental data. The list of interesting extensions and modifications of the current model is long. For example, currently spite only enters the model by affecting the expected payoff of making an offer. Judging by Figure 10, the model might be improved by modeling spite more systematically because the current method can not account for the escalation-reduction cycles seen in the data. In another rather expensive variation of the current methodology one could abandon the best reply dynamics we have used to motivate the adaptation process and run enough sessions to create a stochastic transition matrix. Here the matrix would be based on estimates of actually transiting from one state to another. Finally, one might also consider reworking the types of errors that have been used to create drift in the model. Currently, agents are assumed to lack enough information to always find their best response, but another reasonable approach would models errors that cause deviations both for informational reasons and for preference reasons. This might incorporate the models of nonstandard preferences developed in Fehr and Schmidt [1999], Bolton and Ockenfels [1999], Falk and Fischbacher [1998] or Rabin and Charness [1999].

[22] Note however, such a model would not necessarily predict that the resulting dynamic system would be driven towards the SGPE. Things such as the spite witnessed in the current experiment might prevent such convergence.

6 WORK CITED

Bolton, Gary [1991], "A Comparative Model of Bargaining: theory and evidence," *American Economic Review* 81,1096-1136.

Bolton, Gary, and Axel Ockenfels [1999], "ERC: A Theory of Equity, Reciprocity and Competition," *American Economic Review* 90,166-93.

Camerer, Colin, and Richard Thaler [1995], "Anomalies: ultimatums, dictators and manners," *Journal of Economic Perspectives* 9,209-219.

Carpenter, Jeffrey [2000], "Is Fairness Used Instrumentally? Evidence from Sequential Bargaining," mimeo.

Falk, A., and U. Fischbacher [1998], "A Theory of Reciprocity," mimeo.

Fehr, Ernst, and Klaus Schmidt [1999], "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics* 114,769-816.

Gintis, H. [2000], *Game Theory Evolving*, Princeton University Press: Princeton.

Gueth, Werner, Rolf Schmittberger, and Bernd Schwarz [1982], "An Experimental Analysis of Ultimatum Bargaining," *Journal of Economic Behavior and Organization* 3,367-88.

Gueth, Werner, and Reinhard Tietz [1988], "Ultimatum Bargaining for a Shrinking Cake: an experimental analysis," in *Bounded Rational Behavior in Experimental Games and Markets*, edited by R. Tietz, W. Albers, and R. Selten, Springer: Berlin.

Gueth, Werner, and Reinhard Tietz [1990], "Ultimatum Bargaining Behavior: a survey and comparison of experimental results," *Journal of Economic Psychology* 11,417-49.

Harrison, Glenn, and Kevin McCabe [1992], "Testing Noncooperative Bargaining Theory in Experiments," Pp, 137-69 in *Research in Experimental Economics*, JAI Press Inc.

Harrison, G., and K. McCabe [1996], "Expectations and Fairness in a Simple Bargaining Experiment," *International Journal of Game Theory* 25,303-27.

Henrich, Joe [2000], "Does Culture Matter in Economic Behavior? Ultimatum Game Bargaining Among the Machiguenga Indians of the Peruvian Amazon," *American Economic Review* 90.

Henrich, Joseph, Robert Boyd, Samuel Bowles, Colin Camerer, Ernst Fehr, Herbert Gintis, and Richard McElreath [2001], "In Search of Homo Economics: Behavioral Experiments in 15 Small-Scale Societies," *American Economic Review* 91,73-78.

Ho, Teck-Hua, Colin Camerer, and Keith Weigelt [1998], "Iterated Dominance and Iterated Best Response in Experimental "p-Beauty Contests"," *American Economic Review* 88,947-969.

McKelvey, R., and T. Palfrey [1992], "An Experimental Study of the Centipede Game," *Econometrica* 60,803-36.

Nagel, Rosemarie [1995], "Unraveling in Guessing Games: an experimental study," *American Economic Review* 85,1313-26.

Neelin, Janet, Hugo Sonnenschein, and Matthew Spiegel [1988], "A Further Test of Noncooperative Bargaining Theory," *American Economic Review* 78,824-36.

Ochs, J. [1995], "Coordination Problems," in *The Handbook of Experimental Economics*, edited by J. Kagel and A. Roth, Princeton University Press: Princeton.

Ochs, Jack, and Alvin Roth [1989], "An Experimental Study of Sequential Bargaining," *American Economic Review* 79,355-384.

Prasnikar, V., and A. Roth [1992], "Considerations of Fairness and Strategy: Experimental Data From Sequential Games," *Quarterly Journal of Economics* ,865-888.

Rabin, Matthew, and Gary Charness [1999], "Social Preferences: Some Simple Tests and a New Model," mimeo.

Roth, Alvin, Vesna Prasnikar, Masahiro Okuno-Fujiwara, and Shmuel Zamir [1991], "Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh and Tokyo: an experimental study," *American Economic Review* 81,1068-1095.

Roth, A., and F. Schoumaker [1983], "Expectations and Reputations in Bargaining: An Experimental Study," *American Economic Review* 73,362-72.

Schelling, T. [1960], *The Strategy of Conflict*, Harvard University Press: Cambridge.

Sunder, S. [1995], "Experimental Asset Markets," in *The Handbook of Experimental Economics*, edited by J. Kagel and A. Roth, Princeton University Press: Princeton.

Young, H. Peyton [1993], "An Evolutionary Model of Bargaining," *Journal of Economic Theory* 59,145-168.

Young, P. [1998], *Individual Strategy and Social Structure*, Princeton University Press: Princeton.

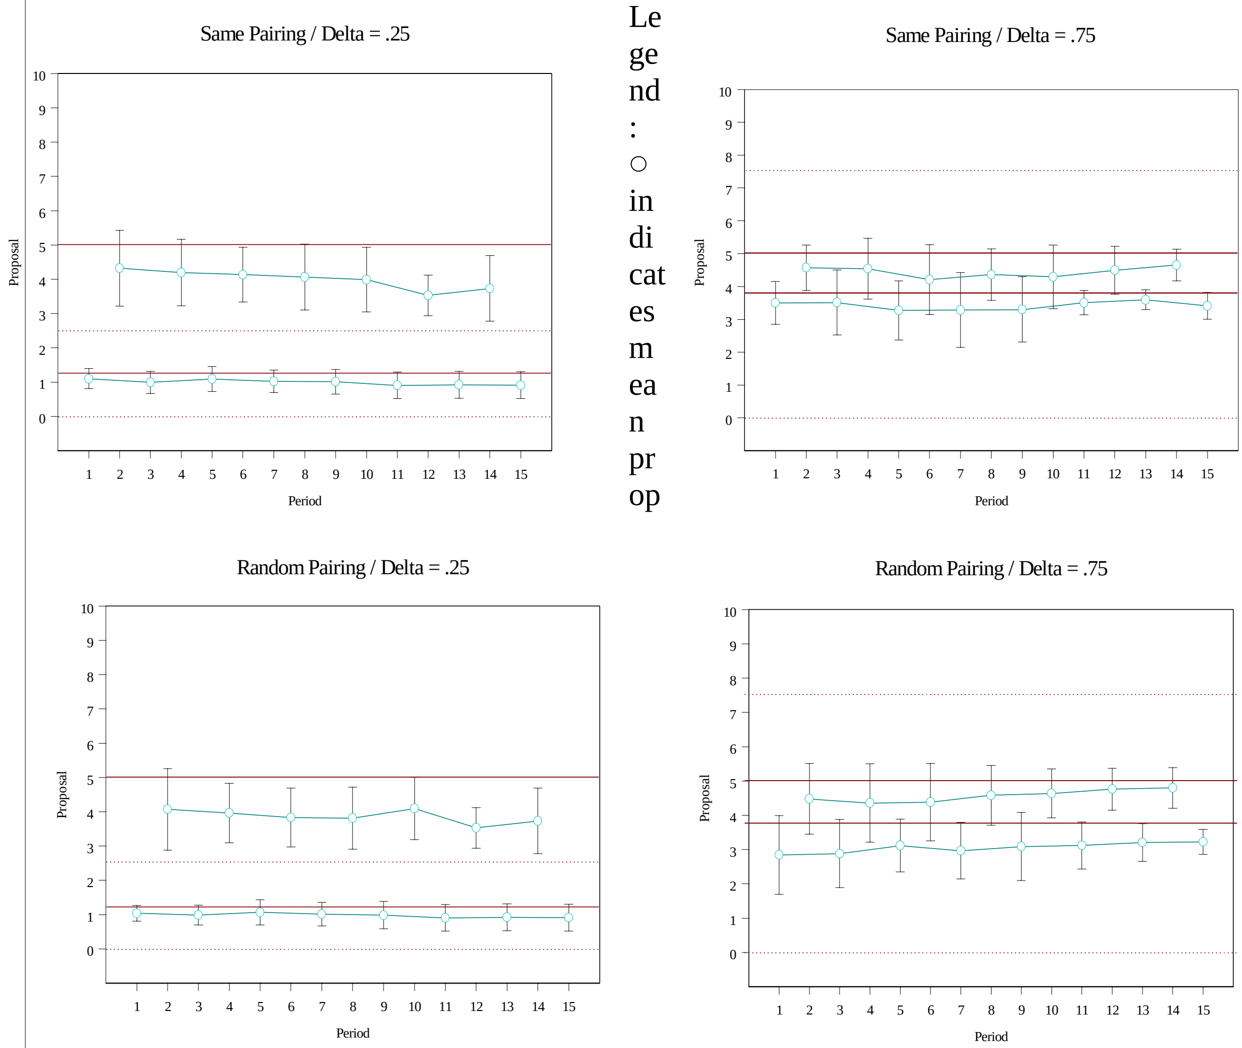

Legend: ○ indicates mean proposal, ⊤ and ⊥ are error bars, ___ indicates and equal split, and ...... indicates the SGPE.

Figure 2 Average First Round and Ultimatum Offers by Period (+/- one deviation)

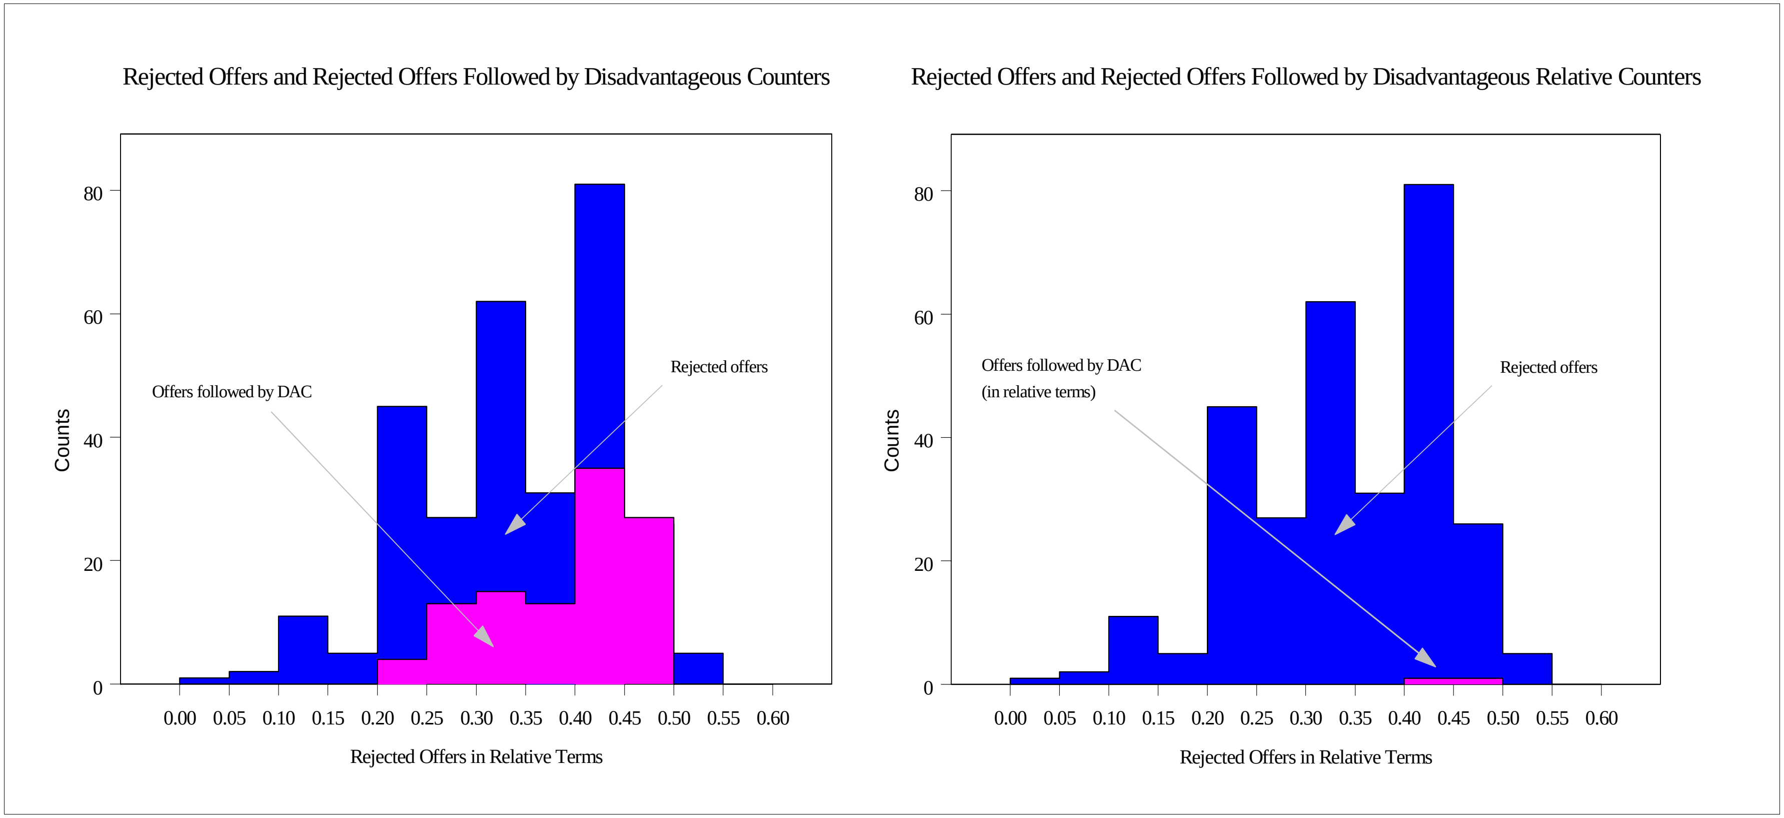

Figure 3 Rejected Offers and Disadvantageous Counterproposals

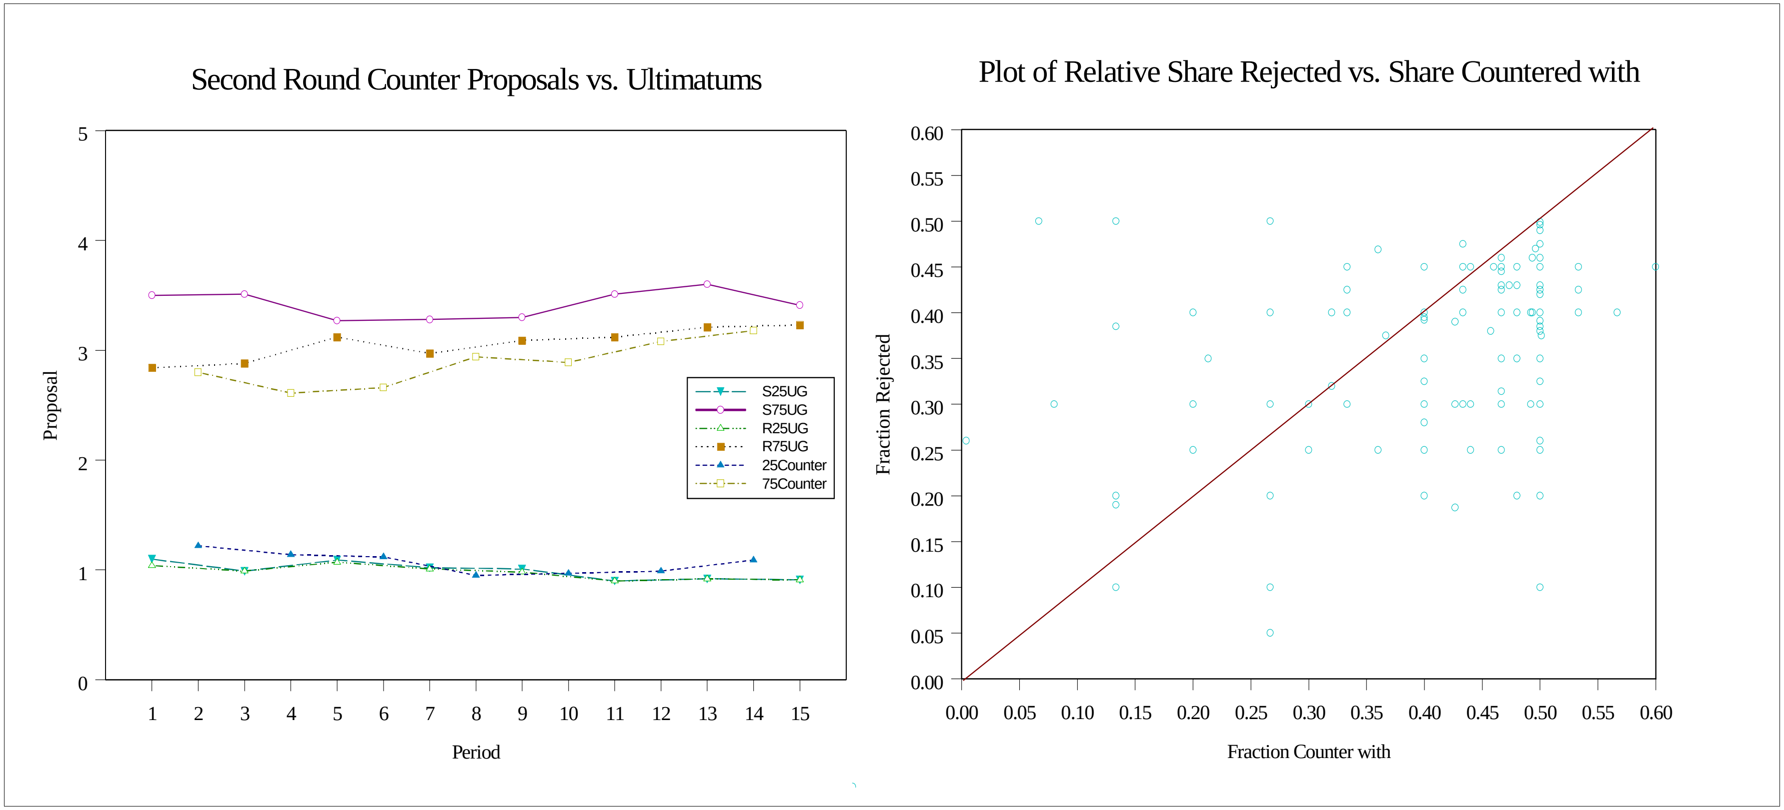

Figure 4 Second Round Counterproposals

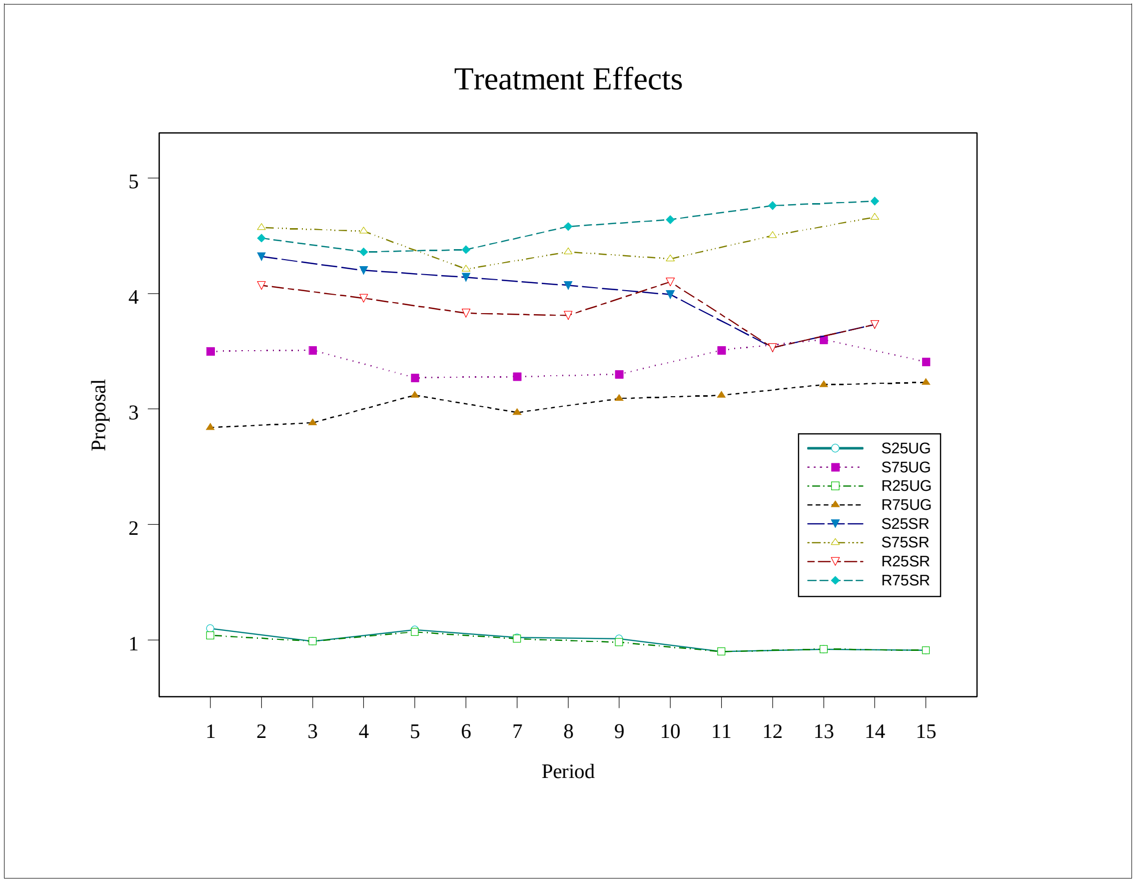

Figure 5 Treatment Effects

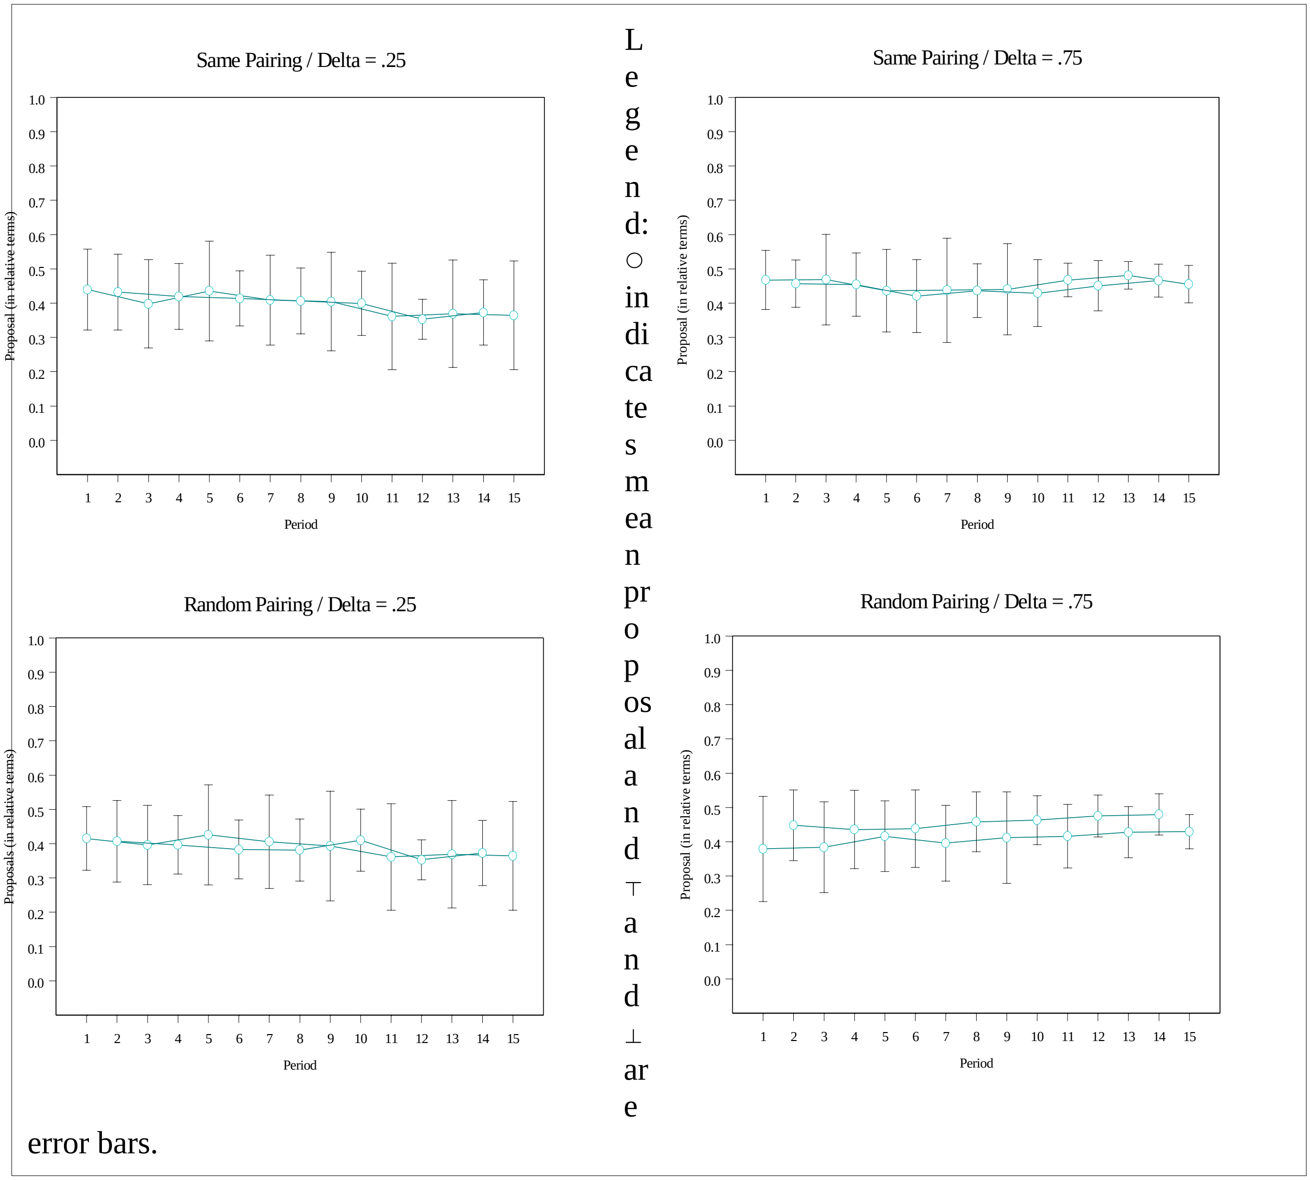

Figure 9 Coordinated Expectations and Relative Proposals

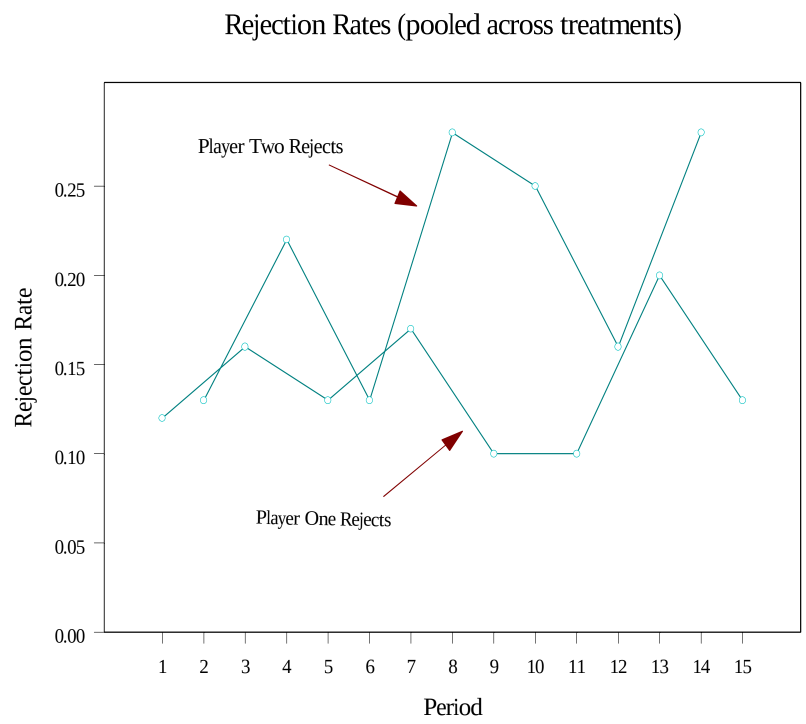

Figure 10 Rejection Rates

APPENDIX - INSTRUCTIONS FOR PARTICIPANTS (random matching treatment)

This experiment is about two-person bargaining. The experiment consists of 15 periods of bargaining between you and another player in the room. All Participants are currently reading the same instructions. At the beginning of each period you will be randomly matched with another player and therefore the likelihood of you being paired with the same player twice is small. You and the person you are matched with will bargain over how to split a sum of Experimental Francs (F) called the 'PIE'. The exchange rate between Experimental Francs (F) and dollars is 1F equals 20 cents.

Each period consists of either 1 or 2 rounds. All ODD periods contain only 1 round and all EVEN periods contain 2 rounds. A round consists of one party's making an offer and the other party accepting or rejecting it. Therefore, in ODD periods (1,3,5,7,9,11,13,15) one party will make an offer and the other party will decide to accept or reject the offer. If the offer is accepted, your final payoff and the final payoff of the other player will increase by the negotiated split of the pie. If the offer is rejected, then both you and the other player will receive 0 francs for this round. Once the second party makes this decision we will wait until all other pairs of subjects have made their choices and then move on to the next period.

All EVEN periods (2,4,6,8,10,12,14) consist of 2 rounds. In the first round one player will make an offer and the other will decide to accept or reject the offer. If this player accepts the offer you will move on to the next period. If this player rejects the offer, the second player will have the opportunity to make a counterproposal in the second round. In the second round, the player who has just rejected an offer will make a counterproposal and the player who made the original offer will be faced with the decision to accept or reject the counterproposal. Additionally, in the second round the SIZE OF THE PIE WILL SHRINK. Therefore, if bargaining in even periods moves to the second round then both parties incur a penalty. Once both players have made their choices in the second round, we will wait for all the other participants and then move to the next period.

When bargaining begins, the half of the screen to the right of these instructions will be filled with buttons, message boxes and information. The message box at the top of the screen will inform you whether you are to make an offer or to wait for an offer. Also this box will tell you the status of the offer you have sent to the person you are paired with. Below this box, are two boxes telling you what period and round it is. Below these boxes, is a frame that appears in yellow that displays the offer that is being proposed to you. You will see both, how much you will get and how much the other player will get if you accept the offer. You will notice that the sum of what you get and what the other player gets always equals the current PIE size.

The current pie size is always displayed below the offer frame. In addition to the current size of the pie, you will see information about the size of the pie last round and next round (if there is a next round).

If the period is odd (1,3,5,7,9,11,13,15) then only this period's pie will be displayed because there is only one round in odd periods. If the current period is even (2,4,6,8,10,12,14) and it is ROUND ONE then you will see the size of the pie this round and the size of the pie next round after accounting for the penalty. If the period is even and it is round two then you will see the current value of the pie and the value of the pie last round. When it is your turn to make a proposal to the other player, you will see a MAKE PROPOSAL button, another message box and a SEND PROPOSAL button. The message box at the top of the screen will prompt you to make a proposal. To make a proposal, click on the MAKE PROPOSAL button. An input box appears asking you how much you would like to propose that the other player gets. You will enter an amount between 0 an the current pie size. THIS IS THE AMOUNT THAT THE OTHER PLAYER WILL RECEIVE.

When you click OK, a message appears in the textbox to the right of the MAKE PROPOSAL button that states the terms of the proposal you are offering. If this is what you want to propose, then send it to your partner by clicking SEND PROPOSAL. IF YOU WANT TO READJUST YOUR PROPOSAL, CLICK THE MAKE PROPOSAL BUTTON AGAIN. If you are in the position to receive an offer in the current round, then you will be told to wait for the other player to send a proposal. When the proposal arrives, it will be displayed in the yellow frame and the buttons to ACCEPT PROPOSAL or REJECT PROPOSAL will be activated. The text boxes next to these two buttons tell you the consequences of accepting or rejecting an offer. When the current period is even and it is round one, if you reject a proposal then you will have the opportunity to make a counterproposal over the pie displayed as 'Value of Pie NEXT ROUND'. If you accept any offer, you will move to the next period.

Your total payoff for the experiment will be the sum of all the francs that you negotiate in the 15 periods. You have been provided with a worksheet to keep track of your earnings for this segment. Please fill out the worksheet as bargaining proceeds. If you have any questions, please raise your hand now. Otherwise, click the FINISHED button to let us know that you have completed reading the instructions. Once bargaining has begun, IT IS VITAL THAT YOU MAKE YOUR

DECISIONS SILENTLY. A summary of the instructions will always appear in this textbox once we have begun the bargaining.