Adjustment Patterns and Equilibrium Selection in Experimental Signaling Games

Jordi Brandts

Charles A. Holt

ABSTRACT

This paper examines the relation between adjustment patterns and equilibrium selection in laboratory experiments with two types of simple signaling games. One type of game has two Nash equilibria, of which only one is sequential. The other type has two sequential equilibria, only one of them satisfying equilibrium dominance. For each type of game, the results show that variations in the payoff structure, which do not change the equilibrium configuration, generate different adjustment patterns. As a consequence, the less refined equilibrium is more frequently observed for some payoff structures, while the more refined equilibrium is more frequently observed in others.

Jordi Brandts
Departament d'Economia i d'Història Econòmica
Universitat Autonoma de Barcelona
08193 Bellaterra, Barcelona, SPAIN

Charles A. Holt
Department of Economics
University of Virginia
Charlottesville, Virginia 22901, USA

ADJUSTMENT PATTERNS AND EQUILIBRIUM SELECTION IN EXPERIMENTAL SIGNALING GAMES

by Jordi Brandts and Charles A. Holt*

1. INTRODUCTION

Even in simple signaling games there are often multiple Nash equilibria. One way to select one of the equilibria is to rule out those that are based on beliefs which, according to some theoretical reasoning, are not sensible. This selection procedure involves considering a particular equilibrium and then determining whether the beliefs that sustain it are reasonable or not. This has been the approach of the refinement literature.[1] In Brandts and Holt (1992) we explore the selection of equilibria from another angle. In experiments with signaling games, subjects gained experience with different partners. This experience during the adjustment to equilibrium determines the equilibrium that is selected. In some of these games, the effect of experience was to reinforce beliefs that are ruled out by the refinement approach, and behavior converged to the *less* refined equilibrium. These results were obtained in two, quite similar, signaling games with two sequential equilibria, only one of them satisfying equilibrium dominance and stronger refinements. In contrast, Banks, Camerer, and Porter (1990) report data from experiments with different signaling games, and their results are supportive of a series of refinements, including sequentiality and equilibrium dominance.

Our objective in this paper is to analyze more systematically how adjustment patterns determine the equilibrium selected in actual play in a sequence of games with different partners. Given the Banks, Camerer and Porter (1990) results, it seems appropriate to explore the validity of the insights obtained in our previous work in the context of their games. We begin by reexamining some of the payoff structures used by Banks, Camerer and Porter in which actual

* We acknowledge financial support from the Spanish CICYT (PB 89-0294 and PB 90-0132), the U.S. NSF (SES-87-20105), and the CIRIT, Generalitat de Catalunya. We thank Jeffrey Banks, José Canals, Ron Harstad, Isabel Sánchez, and two anonymous referees for helpful suggestions.

1 For a discussion of this approach, see Kohlberg and Mertens (1986) and Cho and Kreps (1987).

behavior is consistent with standard refinements. For these games, we argue that experience gained during out-of-equilibrium play tends to confirm the out-of-equilibrium beliefs that support the equilibrium selected by the refinement approach. Next we consider what might happen if behavior in early matchings works against the refinements. We show that a simple modification of the payoff structure can, for a certain type of naive behavior by subjects, generate outcomes that are ruled out by several of the most widely used theoretical refinements.

The way in which subjects gain experience in a series of matchings with different partners in the laboratory is not unlike the strategic interaction envisioned in the literature on evolutionary behavior and learning dynamics. In this literature, it is typically assumed that players do not systematically attempt to influence other players' future actions, and the distribution of decisions changes gradually over time. Rational learning in this context can result in behavior that converges to a Nash equilibrium (Kalai and Lehrer, 1991, and Jordan, 1991). Theoretical work in this area is surveyed in Friedman (1991), Mailath (1992), and Marimon and McGrattan (1992). Experimental work that explicitly studies learning is still rather rare.[2]

Section 2 describes the designs that are used in the experiments reported here. Laboratory procedures are summarized in section 3, and the data are analyzed in section 4. The final section contains a conclusion.

2. Experimental Design

Two types of signaling games will be considered. First, we analyze behavior in games with two Nash equilibria, of which only one is sequential. Second, we analyze games with two sequential equilibria, of which only one satisfies the notion of equilibrium dominance. In each case, we compare the predictions of standard refinements with the implications of a naive adjustment process for a laboratory setup in which subjects learn from experience as they are matched with different partners in a series of two-person games. This adjustment process is based on a consideration of the behavioral patterns in signaling games reported in Brandts and

[2] Selten and Stoecker (1986) use a learning-theoretic approach to analyze experiments with finite repetitions of prisoner's dilemma games. Friedman (1992) reports results from some normal-form game experiments designed to test some elements of evolutionary game theory. Crawford (1991) provides an evolutionary interpretation of some coordination game experiments.

Holt (1992), and it was used as a guide in our design of the experiments to be reported below. Although the adjustment that we describe here is not a precise model of individual decisions, it is roughly consistent with the data to be reported.

Nash versus Sequential

The first game to be considered is game 1:

Game 1

	response					response		
message N	C	D	E		message S	C	D	E
type A	15, 30	30, 30	0, 45		type A	45, 15	0, 0	30, 15
type B	30, 30	15, 60	45, 30		type B	30, 30	0, 0	30, 15

In this game, player 1's type, A or B, is randomly determined. Both types are equally likely. Player 1 of either type can send one of two messages, N or S. Player 2 observes the message but not the type, and can answer either message with C, D, or E. The payoffs for each type/message/response combination are shown in the table, with the payoff for player 1 listed to the left of the comma. This game has the same payoff structure as game 3 in Banks, Camerer and Porter (1990), with the exception that we deleted a third message that player 1 would not send in either of the two equilibria.[3]

Game 1 has two Nash equilibria. In the nonsequential Nash equilibrium, both types of player 1 send message N (for Nash), and player 2 answers both N and S messages with response D. Hence the payoffs for this equilibrium are 30 for the type A player 1 and 15 for the type B player 1. Since the two types are equally likely ex ante, the expected payoff for player 2 is the average of 30 and 60. In the sequential equilibrium, both types of player 1 send message S (for sequential), with message S receiving response C and message N receiving response D. Even

[3] Since our objective was to study how some Nash equilibria are selected over others, we decided to remove the third message in order to reduce the incidence of nonNash outcomes. Moreover, the message that we deleted was never used by any subject in game 3 in the Banks, Camerer, and Porter (1990) experiment.

though both types of player 1 are better off at the sequential equilibrium than at the other equilibrium, player 1 will choose message N if he or she believes that player 2 will respond to S with the response D. Since player 2's response D is dominated, the Nash equilibrium involving N is not sequential. The Nash equilibrium involving S *is* sequential, since the response D to the deviation N is a best response for some well-defined beliefs, i.e., the probability is at least 1/3 that the deviant is of type B. Notice that in game 1, the Nash equilibrium outcome is consistent with out-of-equilibrium beliefs for the sequential equilibrium, but the reverse is not true. That is, the sequential equilibrium outcome does not correspond to the out-of-equilibrium beliefs for the Nash equilibrium. These features somewhat favor the sequential equilibrium.

The results reported in Banks, Camerer and Porter (1990) give rather strong support to the sequential equilibrium in the variation of game 1 that includes the third message. Their experiment involved a group of subjects who were randomly rematched with different partners in a sequence of 10 signaling games.[4] Play with little experience (1 to 5 periods of play) already leads to 63% of the observed outcomes corresponding to the sequential equilibrium and 12% to the Nash equilibrium. With more experience (6 to 10 periods of play), 73% of the outcomes are sequential and 22% Nash, with fewer disequilibrium outcomes. The data show a clear preponderance of the sequential equilibrium outcome, even with little experience.

One possible interpretation of these results is that the selection of the more refined equilibrium in actual play is due to the logic of the refinement argument itself. For game 1 this means that subjects recognize that the threat, which supports the out-of-equilibrium message for the nonsequential Nash equilibrium, is not reasonable. This recognition may be helped by the focality of the zero payoffs that correspond to the threat. As a consequence, both types of player 1 choose S, which leads to the sequential outcome at which both types are better off. But this is not the only interpretation of the results obtained by Banks, Camerer, and Porter. Another possibility is that subjects follow some reasonable but naive thought process that leads them to the sequential outcome, but for the "wrong" reasons.

[4] Banks, Camerer, and Porter also report some 20-period sequences. They conclude that there is little additional convergence during the final 10 periods.

What naive thought process could explain the results for game 1? In the first period, subjects with the player 1 role have no experience on which to base an assessment of which responses to the messages are more likely. A plausible approach would be to behave as if all responses were equally likely and to choose the message with the highest row sum of payoffs. Note that the sum of a type A player's possible payoffs from playing S is larger than from playing N ($45 + 0 + 30 > 15 + 30 + 0$), whereas for a type B the sum of payoffs is larger from playing N. If subjects with player 1 roles believe that each of the three responses is equally likely, independently of the message sent, then the type A player will choose message S, and the type B player will choose message N. Since the two types, A and B, occur with the same probability, the analogous naive behavior for player 2 would be to assume that each message is equally likely to have come from a type A or a type B player 1. Given these beliefs, the best response for player 2 is to select the decision with the greatest column payoff sum, i.e. to answer message S with C and message N with D. For this game, these responses turn out to be both the equilibrium responses for a sequential equilibrium, and the best responses to the naive pattern of messages sent by player 1, i.e., to an S message from type A and an N message from type B. But as subjects with the player 1 role gain experience with the response pattern, type A will not want to switch away from the S message, and type B will come to realize that the payoff can be increased from 15 to 30 by switching to the S message. In this manner, play converges to the sequential equilibrium.[5] According to this explanation, the selection of the sequential equilibrium in actual play is the consequence of the particular adjustment pattern triggered by some of the features of the payoff structure for game 1.

The question that arises from the above description of a possible naive adjustment pattern is whether a change in the payoffs can produce a different pattern of adjustment that leads to the *nonsequential* Nash equilibrium. In particular, consider the following game, which is the result of payoff changes that do not alter the configuration of the two equilibria.

---

[5] Think of each player 1 as having his or her type determined randomly at the start of each period in which a different player 2 is encountered. In any given period, a player 1 only sees the response to one of the two messages, and a player 2 only sees (ex post) the type that sent a particular message. The argument that players switch to best responses leaves unanswered the question of how learning occurs when players do not see aggregate behavior patterns. The learning necessary to form a best response to aggregate behavior presumably requires several periods.

Game 2

message N	C	D	E	message S	C	D	E
	response				response		
type A	75, 30	45, 30	75, 45	type A	60, 15	0, 0	0, 45
type B	75, 30	30, 75	75, 30	type B	45, 60	0, 0	45, 15

It is easy to check that the equilibria in this game are the same as in game 1. Note that, as before, the payoffs for each type of player 1 in the sequential equilibrium (S,C) outcomes are greater than those in the nonsequential Nash equilibrium (N,D) outcomes. For both types of player 1, however, the sum of payoffs across responses is larger when the message N is sent. For this reason, naive behavior by subjects in the role of player 1 could lead to both types initially sending message N. Since the two types, A and B, are equally likely, naive behavior for player 2 would again be to assume that each message is equally likely to have come from a type A or a type B player. Given these beliefs, the best decision for player 2 is to answer message N with D, as was the case for game 1, and this would lead to convergence at the nonsequential equilibrium.[6] A feature of this adjustment process is that message S is never used by either type, and therefore, players do not get experience with that part of the game as their types change from period to period. This may lead to subjects not realizing that the response D to message S, which sustains the nonsequential equilibrium, is dominated. In contrast with game 1, game 2 is designed in a way that experience in the adjustment process may lead to the less refined, nonsequential equilibrium.[7]

[6] The Nash equilibrium D response to the N message provides the lowest of the three payoffs for each type of player 1 that could result from that message. As a player 1 becomes more confident that the N message will receive a D response, the alternative S message becomes more attractive, and we might expect some "leakage" to S in later periods.

[7] In a bargaining context, Harrison and McCabe (1992) show that experience in all parts of the game can be critical. In their experiment, they forced subjects to play all of the subgames before playing the complete game from the beginning. They report that forced experience in the subgames tends to increase the drawing power of the subgame-perfect equilibrium prediction.

Sequential versus Intuitive

The nature of the adjustment process could also be the reason for the behavior observed by Banks, Camerer and Porter in a related experiment designed to distinguish between the sequential equilibrium and the more refined, intuitive equilibrium. Consider Game 3, which has the same payoff structure as game 4 in Banks, Camerer and Porter, after the elimination of a third message that does not lead to either of the equilibrium outcomes:

Game 3

	response					response		
message I	C	D	E		message S	C	D	E
type A	45, 30	15, 0	30, 15		type A	30, 90	0, 15	45, 15
type B	30, 30	0, 45	30, 15		type B	45, 0	15, 30	30, 15

In the sequential equilibrium, both types of player 1 send S, which is answered with C, and a deviation message of I is answered with D. The D response is best for player 2 if the respondent believes that the deviation was very likely to have come from a type B player. According to Cho and Kreps (1987), these beliefs are unreasonable in the sense that the type B player earns 45 in the sequential equilibrium, and a deviation to message I can never yield a payoff above 30. Therefore, the equilibrium involving the S messages is sequential but not intuitive. In the intuitive equilibrium, both types send message I, which is answered with C, and the out-of-equilibrium message S receives the D response that is optimal if the deviant is believed to be of type B. These beliefs are not unreasonable, in the sense that the equilibrium payoff of 30 for a type B who sends the I message could conceivably be increased to 45 by a deviation to message S.

The results in Banks, Camerer, and Porter again give strong support to the more refined equilibrium. The intuitive outcome is clearly more frequent than the nonintuitive equilibrium, regardless of the degree of experience. In sessions with 10 consecutive periods, the frequency of the intuitive outcome was 53% for periods 1 to 5, and 68% for periods 6 to 10. In contrast, the frequency of the nonintuitive equilibrium outcome was only 13% in periods 1-5, and 3% in periods 6 to 10.

In game 3, subjects' naive reasoning of the type explained above could again be the explanation for the strong preponderance of the intuitive equilibrium. Notice that, for a player 1 of type A, the sum of possible payoffs corresponding to message I is higher than for S, whereas for a player 1 of type B the sum of payoffs corresponding to S is higher than for I. Initial choices by player 1 could lead to "type dependence", with type A sending message I and type B sending message S. Similarly naive behavior by player 2, who initially expects that a message is equally likely to have come from either type, will lead to the C response to either message, since this response provides the maximum column sum. This C response should not alter the relation between types and messages for player 1, since each type is getting its maximum payoff of 45. But as in game 1, subjects with the player 2 role will find out that the I message comes from type A and the S message comes from type B. The best response to this type dependence is to answer I with C, which is the response in the intuitive equilibrium, and to answer S with D. This pattern will lead type B to switch to I and causes convergence to the intuitive equilibrium.

Game 4

	response				response		
message I	C	D	E	message S	C	D	E
type A	30, 30	0, 0	50, 35	type A	45, 90	15, 15	100, 30
type B	30, 30	30, 45	30, 0	type B	45, 0	0, 30	0, 15

Game 4 above has the same equilibrium configuration as game 3 but might involve a different adjustment process and, therefore, lead to a different equilibrium. It follows from a naive belief in equally likely responses that message S is more attractive for a player 1 of type A, and message I is more attractive for a type B. In this way, we reverse the correlation between messages and types that was predicted for naive players in game 3. A naive player 2, who initially believes that either message is equally likely to have come from either type, will respond to each message with C. This naive response gives the type B player 1 an incentive to switch from I to S, and it gives the type A player 1 an incentive to stay with S. Hence play converges to the sequential equilibrium. Moreover, player 2 should remember that an I message was previously sent by a type B player, and therefore, any deviation to an I message should be met with the D response that supports the sequential equilibrium. This D response to an I message is optimal if the deviant is expected to be of type B, a belief that is consistent with previous experience. Nevertheless, this belief is inconsistent with the intuitive criterion, since the type B player earns 45 in the sequential equilibrium, which is greater than any of the possible payoffs that could result from a deviation to the I message. To summarize, we designed game 4 so that a series of matchings with different players could generate a pattern of adjustment that determines the beliefs about what would happen off of the equilibrium path once an equilibrium is reached. In particular, these beliefs are inconsistent with the intuitive criterion.

Before proceeding, it is useful to compare games 3 and 4. In game 3, the initial type dependence is for a naive type A to send message I and for a naive type B to send message S. This type dependence pattern is reversed in game 4, with type A sending message S and type B sending message I, which could lead to a different, less refined equilibrium outcome. But in moving from game 3 to game 4, the reversal of the adjustment process necessitated changes in the equilibrium payoffs. Unlike game 3, the sequential (S,C) outcome is best for both player 1 types in game 4. This observation leads us to consider game 5, which has the same initial "reverse type dependence" as game 4 and the same equilibrium payoffs as game 3:

Game 5

	response					response		
message I	C	D	E		message S	C	D	E
type A	45, 30	0, 0	0, 15		type A	30, 90	30, 15	60, 60
type B	30, 30	30, 45	30, 0		type B	45, 0	0, 30	0, 15

Observe that now, naive expectations of equally likely responses cause a type A player 1 to choose message S and cause a type B player 1 to choose message I, as in game 4. Player 2's best response to this behavior is to answer S with C and I with D, which will lead type B to switch to S, and play could converge to the nonintuitive outcome. But this story is not exactly consistent with the adjustment processes described earlier, since player 2's initial behavior, based on a belief that a message is equally likely to have come from either type, will lead to the C response to either message, and the best response of player 1 is to go back to the type dependence pattern (type A sends I, type B sends S) that was predicted for game 3. Given this ambiguity, we defer a discussion of adjustment for game 5 until the data have been presented.

3. PROCEDURES

Games 1-5 were used in experiments with groups of subjects, who were recruited for two-hour sessions from undergraduate economics classes, either at the Universitat Autonoma de Barcelona or at the University of Virginia.[8] Subjects were told that they would receive an initial payment in addition to all cash earnings obtained during the session, which began with a reading of the instructions in the appropriate language. The appendix contains a copy of the instructions, with payoffs for game 1.

There were 10 sessions, two for each game. Each session involved 12 participants and one "monitor", who was one of the subjects selected for this task by the throw of dice at the beginning of the session. The monitor's role was to throw the die that determined the preference types and to observe and ensure that the session was being conducted in accordance with the procedures specified in the instructions. The subjects were placed in two adjoining rooms, with half of them in one room having the role of player 1, the "proponent", and the other half having the role of player 2, the "respondent".

An experimental session consisted of a series of matchings, with each subject being paired with a different partner in the other room in each matching. At the beginning of each matching, the monitor would throw a six-sided die in the proponent's room, separately for each subject. The throw of the die would determine the type: type A for three of the outcomes and type B for the other three. Each proponent then selected his or her message, which was privately communicated to the paired respondent in the other room. Then each respondent would choose

[8] For games 1, 2, and 4, half of the sessions were conducted at each location. All sessions using games 3 and 5 were conducted in Barcelona.

a response, which was again privately communicated to the appropriate proponent in the other room. Finally, each respondent was informed of the proponent's type, and all subjects calculated their payoffs with the tables provided.

The payoffs from the tables were called "points" and were used to determine monetary earnings via a binary lottery procedure: For each subject, two 10-sided dice with markings from 0 to 9 were thrown to determine a random number between 0 and 99. If the random number was less than the point total obtained during the period, then the subject would earn a high amount ($2.00 in Virginia or 200 pesetas in Barcelona).[9] Otherwise, the subject received the low payoff (50 cents or 50 pesetas). In this manner, the points earned from the payoff table are essentially probabilities of getting the high prize, and it is well known that this procedure induces risk neutrality as long as subjects are expected utility maximizers.[10]

There were no public announcements of results; this served to maintain the 2-person nature of the game. After each period, subjects were rematched in a way that ensured that no person would meet the same partner twice in the same role. Moreover, the matching rotation was devised to ensure that there were no indirect links, i.e. there was "no contagion", as in the design used by Cooper et al. (1991). Subjects were told initially that there would be six matchings with different partners using the same payoff table, after which instructions for "another experiment in decision making" would be read. After the sixth matching, new instructions and decision sheets were distributed, with the same payoff table, but with different identification numbers and rotation schedules for 6 more matchings. In addition, the player 1 and player 2 roles were reversed. The objective of this deterministic matching procedure was, specifically, to preserve

---

[9] During the period in which the experiments were run, the exchange rate was approximately $1.00 for 100 pesetas.

[10] The binary lottery procedure has the disadvantage of adding an extra layer of procedural complexity. Although subjects had no trouble understanding the procedure after going through it once, it did require more time. The primary reason that we used this procedure is that our analysis of the equilibrium and of the adjustment process is based on an assumption of risk neutrality. Without a procedure for inducing risk neutrality, we would be open to the additional criticism that subjects become less risk averse as they earn more money in the course of the session, and therefore, that adjustment patterns may be due to changes in risk aversion.

Table 1. Proportions of Outcomes, by Refinement

Game (periods)	Refinements				
	Non-Nash	Nash	Sequential	Intuitive	Number of Matchings
1* (1-5)	.25	.12	.63		60
1* (6-10)	.05	.22	.73		60
1 (1-6)	.10	.32	.58		72
1 (7-12)	.03	.19	.78		72
2 (1-6)	.28	.50	.22		72
2 (7-12)	.12	.50	.38		72
3* (1-5)	.34		.13	.53	90
3* (6-10)	.29		.03	.68	90
3 (1-6)	.34		.19	.47	72
3 (7-12)	.27		.06	.67	72
4 (1-6)	.35		.51	.14	72
4 (7-12)	.26		.72	.02	72
5 (1-6)	.41		.46	.13	72
5 (7-12)	.42		.50	.08	72

* Data for games 1* and 3* in the table are taken from Banks, Camerer, and Porter (1990).

the one-period nature of the game.[11] The procedure cannot protect against the possibility that one player's actions influence those of another player in a later matching. Therefore, the results of two-person games taken from the same session are not independent observations.

[11] The use of a preannounced, deterministic rotation schedule with no contagion severely limits the number of matchings that can be done, even with a relatively large number of subjects. We could have used a random matching process, but even a low probability of meeting the same partner again dilutes the two-person nature of the game. If this dilution were acceptable, a random rematching procedure could be used for more matchings, thereby giving subjects more time to adjust. Because the deterministic rotation for 12 matchings already required a full two hours, we decided against using random matchings. The unannounced role reversal after 6 matchings does introduce discontinuity in experience, but it doubles the number of data points without introducing obvious repeated game effects.

4. RESULTS

Games 1 and 2

As indicated in section 1, Banks, Camerer and Porter found strong support for the sequential equilibrium concept in a game that we will denote by 1* (their game 3). The results for this game are summarized in the top two rows of table 1. Each entry is a proportion of outcomes, i.e. message/response combinations. For example, 63 percent of the first five matchings for game 1* yielded a sequential outcome: message S followed by response C. Since our game 1 is a simplified version of game 1*, we expected behavior in the two games to be similar. Therefore, we conducted two experimental sessions with game 1, and the results are shown in the third and fourth rows of table 1. As expected, the frequencies for the different types of outcomes in game 1 are qualitatively similar to the ones in game 1*. The increase in experience leads to an even stronger swing from Nash to sequential outcomes, and there are only two nonNash outcomes in the last 6 periods. A breakdown of messages and responses for game 1 is provided in Appendix B.

The pattern of adjustment for game 1 is shown in figure 1, where data for adjacent periods have been combined. In periods 1 and 2, the type A players send message S in 64% of the matchings, and type B players only send S in 38% of the matchings. This type dependence, which was anticipated in the earlier discussion of adjustment, is reversed in periods 3-4, but reappears in later periods as indicated by the location of the dark S/A line above the dashed S/B line in the upper part of figure 1. The lower part of the figure shows that the S message almost always receives the sequential equilibrium response, C, as indicated by the C/S line between .88 and 1.0 in the bottom part of the figure. Moreover, subjects with the player 2 role respond to N by choosing the D response that is appropriate when this message is very likely to have come from a type B player. After the initial period, only 2 of the 36 responses to the N message were different from the D response. The D response to N is a “punishment” for the type B players, giving them a payoff of only 15. As a result the B players tend to switch to the S message, and by the last period all 12 subjects with player 1 roles use the S message. These features of the data from game 1 are consistent with the adjustment pattern suggested above.

The results from two experimental sessions with game 2 are summarized below those for game 1 in table 1. The data for game 2 are quite different. With 1-6 periods of play, 50% of

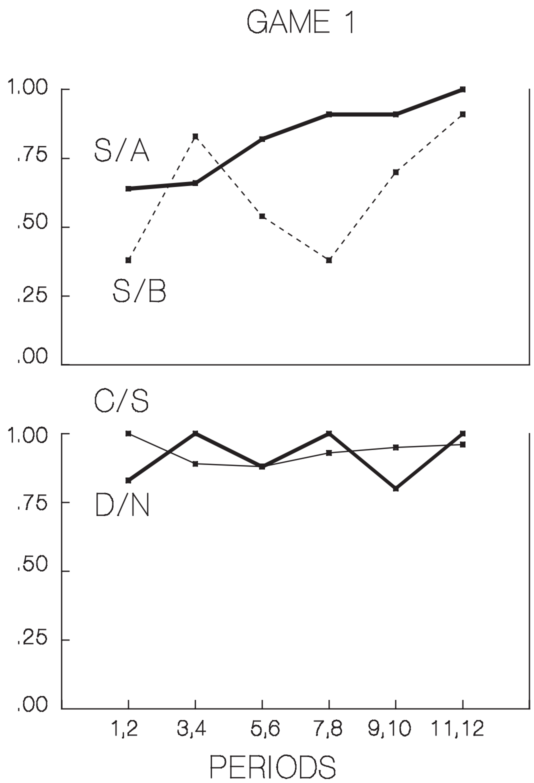

Figure 1. Adjustment for Game 1

the observations correspond to the Nash outcome, while 22% correspond to the sequential equilibrium. With more experience, 7-12 periods of play, the frequency of the Nash outcome does not change, but the frequency of the sequential outcome increases to 38% due to the decrease in nonNash play. In game 2, the sequential equilibrium fails to be the dominant outcome for the two levels of experience that we explored.[12] Even though half of the outcomes are nonsequential Nash outcomes, the dominated response D to message S is almost never observed, despite the fact that this is the response off the equilibrium path that supports the nonsequential Nash equilibrium. This can be seen from the location of the dark D/S line in the bottom part of figure 2.

Recall that game 2 was designed to induce both player 1 types to choose the Nash message N from the beginning. The intention was that subjects with player 1 roles would not acquire experience on the S message side of the game and, hence, they would not learn about the actual use, as a reply to message S, of the dominated D response that supports the Nash equilibrium. The data sustain our intuition and show that both A and B types choose message N overwhelmingly in the first period. In this initial period, all 6 type A players chose N, an extreme outcome that would occur with probability $(1/2)^6$ if the two messages were equally likely. Similarly, 5 out of 6 type B players chose N in the first period,

[12] We conducted one additional session with a slight variation of game 2. The row sums of player 1's payoffs were less sensitive to the message. In this session, the payoffs used in the first period were printed incorrectly in the instructions, but the intended payoffs were used in periods 2-12. In general, the pattern of data is similar to that for game 2, but with far fewer sequential messages and with somewhat more disequilibrium (C and E) responses to Nash messages. The percentage of Nash signals was 97 % in periods 1-6 and 83 % for periods 7-12. In periods 7-12, 53 % of the message/response combinations were Nash and 11 % were sequential.

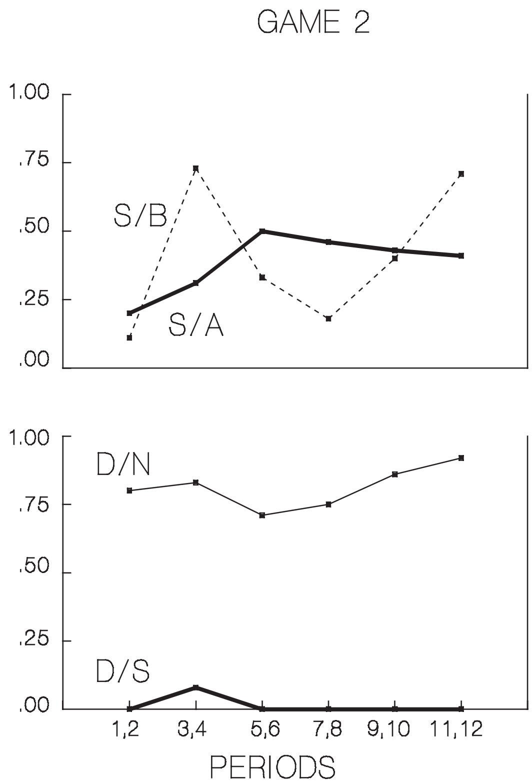

Figure 2. Adjustment for Game 2

which is also rather unlikely ($p < .08$) under a null hypothesis that the two decisions are selected with equal probability. The individual data for this game also show that 6 out of 24 subjects never used the S message, and therefore did not get any experience on that side of the game. The aggregate data for game 2 in table 1 and in Appendix B show that this preponderance of the N signal diminishes in later periods, but the N message remains the more frequent one, probably as a consequence of the initial behavior we induced. In any event, a comparison of the top parts of figures 1 and 2 indicates that behavior in game 2 is not converging to the sequential equilibrium as in game 1.

Games 3-5

Next compare the data for our game 3 with the Banks, Camerer, and Porter data for game 3* (their game 4). As explained before, the payoffs for games 3 and 3* are identical, except for the omission of a third message in our game. This omission makes no difference, both for the structure of the equilibrium outcomes and, as can be seen from table 1, for the behavior observed. The data that we obtain are virtually the same as those of Banks, Camerer, and Porter. The intuitive equilibrium outcomes predominate, even in early periods, and the incidence of nonintuitive, sequential outcomes is low and diminishes with experience.

The message and response data for this game provide further support for the intuitive equilibrium in game 3. But, at the same time, the adjustment patterns are quite consistent with those that we have described in the earlier section on design. In particular, there is strong initial type dependence. For period 1, type dependence is significant at the 10% level on the basis of a Fisher exact probability test. The stronger propensity for type A players to send the I message can be seen by the location of the I/A line over the dashed I/B line in the upper part of figure 3. The prediction for respondents was that they would naively choose the C response to either message, since the C response provides the greatest expected payoff if a message is naively expected to be equally likely to have come from either type. After respondents come to anticipate the relationship between messages and player 1 types, they should switch to a D response for the S signal, which is appropriate when the S signal comes from a type B. In fact, respondents were somewhat less naive than this story suggests, as the D response to the S signal was fairly common even in the first matching. Moreover, almost all of the disequilibrium outcomes for this game result from type B players choosing the S signal, which has the highest row payoff sum for them, and getting the D response. The location of the D/S line above .5 in the lower part of the figure indicates that the respondents understand that it is likely to have come from a type B player. On the other hand, the high level of the C/I line is due to the fact that the I signal almost never receives the D response that is ruled out by the intuitive criterion.

The next two rows in table 1 show the results for game 4, which has the same configuration of intuitive and sequential equilibria as game 3. In spite of this, the drawing power of the two equilibria is reversed. As can be seen from a comparison of the top parts of figures 3 and 4, the pattern of type dependence has been reversed with type B players now choosing message I more often and type A players now choosing message S more frequently.[13] Although the C response to an I message is common in the initial matchings, the proportion of D responses to this message rises steadily toward 1.0, indicating that the respondent believes that the message came from a type B player. These are precisely the beliefs that are ruled out by equilibrium dominance. The D response to the I message, together with the high incidence of C responses to the S message, provides an incentive to switch to the S message, and this almost doubles the frequency with which type B players choose the S message in the final 6 periods.

A comparison of figures 3 and 4 indicates that the convergence of behavior to the sequential equilibrium in game 4 is about as strong as the convergence of behavior in game 3 to the sequential equilibrium. In the final 2 periods, 79% of the messages are intuitive in game

[13] The type dependence in the first matching is significant at the 5% level using a Fisher exact probability test.

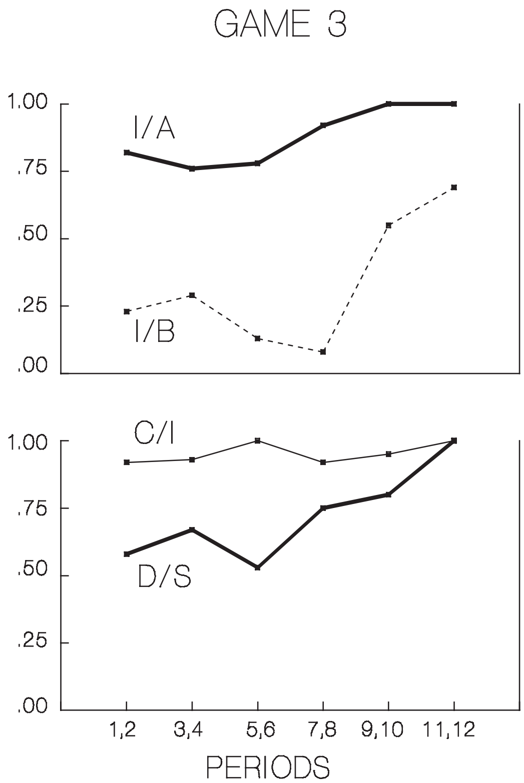

Figure 3. Adjustment for Game 3

3, and 79% of the messages are sequential in game 4. Also, the 100% of the responses in the final 2 periods of game 3 are those specified by the intuitive equilibrium (C for I and D for S), and 92% of the responses in the final 2 periods of game 4 are those specified by the sequential equilibrium (C for S and D for I). To the extent that the data for game 3 are interpreted as being supportive of the intuitive refinement, the data for game 4 have to be interpreted as evidence against the refinement.

The bottom two rows of table 1 show the summary results for game 5, which has the same equilibrium configuration as games 3 and 4. Although there are considerably more nonNash outcomes in game 5 than in the other games, there are more than five times as many sequential outcomes as intuitive outcomes for the last 6 periods of game 5.[14]

Type dependence was perfect in the first matching for game 5, with all 5 type A players choosing message S and all 7 type B players choosing message I, a difference that is significant at the 1% level using a Fisher exact probability test. As can be seen from figures 4 and 5, the adjustment pattern in game 5 is similar to that for game 4, with the same reverse type

[14] We conducted an additional session with a slight variation of game 5. Only two payoff numbers were changed, and the effect was to make the E a dominated response to message S. The data for this session were quite similar to the data for the two sessions with game 5. In part B, 58% of the message/response combinations were sequential and only 17% were intuitive.

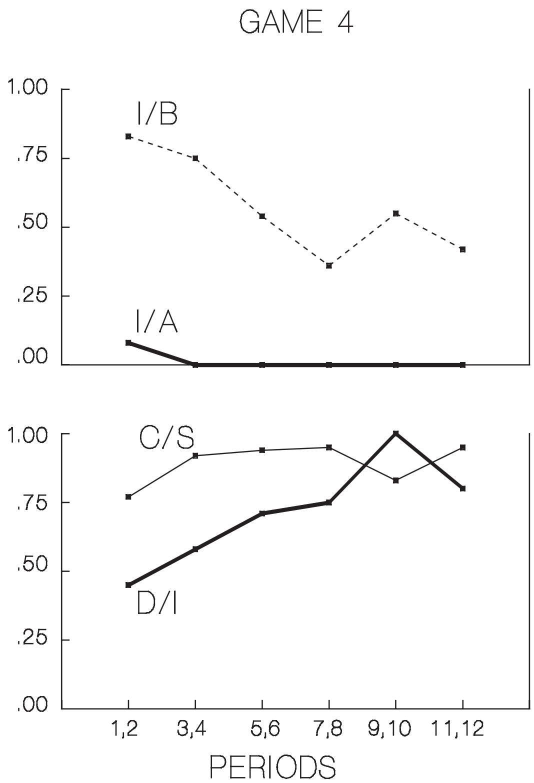

Figure 4. Adjustment for Game 4

dependence that diminishes over time as type B players switch to the S decision as a consequence of the likelihood of obtaining the D response to the I message. As indicated in section 2, the initial reverse type dependence would tend to vanish if the respondents used a naive response of C to each message, which would complicate the adjustment process. The fact that the type dependence persists may be due to the tendency for respondents to anticipate the type dependence by following message I with response D even in the early matchings, as can be seen from the bottom part of figure 5 and as was also the case in games 3 and 4. Unlike the situation in games 3 and 4, this less naive behavior for respondents is crucial here since it prevents a change in the initial type dependence, and thereby allows for convergence.

Statistical Comparisons

Since the behavior of individuals in early matchings of a session can influence other's behavior in later matchings, it is not appropriate to use the outcomes of individual matchings (after the first matching) as independent observations in statistical tests. What then can be concluded if each session is only counted as a single data point? In a strict sense, a refinement implies that the more refined equilibrium will always be observed, and a single session with a high proportion of less refined equilibrium outcomes could be considered a contradiction of the theory. A weaker interpretation, which we consider reasonable, is that a refinement implies that one equilibrium is somewhat more likely to be observed that another. We have observed a higher proportion of less refined, sequential outcomes in all four sessions with the reverse type dependence designs (games 4 and 5). If we interpret a refinement as only predicting that the

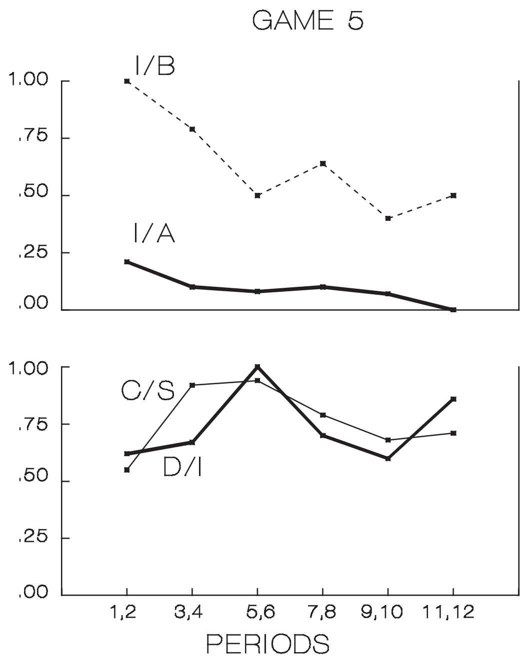

Figure 5. Adjustment for Game 5

more refined equilibrium outcomes will predominate with some probability $p > .5$, then the chance of observing a predominance of less refined outcomes in all 4 sessions is $(1-p)^4$. For example, the probability of a predominance of less refined outcomes in all 4 sessions is only .05 if $p = .53$, which would correspond to a refinement that provides only a very slight advantage over using a coin flip to predict which of the two equilibria will predominate. Similarly, if $p = .68$ (a prediction that the more refined outcomes predominate about twice as often as less refined outcomes), then the chance of seeing the less refined outcomes predominate in all 4 sessions is only .01.

An alternative way to proceed is to test whether the change from treatments with normal type dependence to treatments with reverse type dependence affects the proportion of intuitive equilibrium outcomes. If we only compare the two sessions for game 3 with the two sessions for game 4, we cannot reject the null hypothesis (at any reasonable level of significance) that the proportions of intuitive outcomes for sessions with games 3 and 4 are drawn from the same distribution. But if we treat the Banks, Camerer, and Porter data for game 3* as a separate observation, and if we pool the four observations from games 4 and 5, then there are three "control" observations and four "treatment" observations. All four of the treatment observations have lower proportions of intuitive outcomes than each of the control observations. The number of ways in which the treatment and control observations could be ranked is 7!/(3!4!) = 35, and of these, we observe the most extreme outcome in which the four treatment observations have the highest ranks. Under the null hypothesis, the chance of this outcome is only 1/35, which is significant at about the 3% level in a one-tailed (Mann-Whitney) test.

The Nash-versus-sequential results can be analyzed in the same manner by treating the Banks, Camerer, and Porter game 1* as a single observation, which taken with the two sessions for game 1 yields three control sessions. The two treatment sessions are ones that were conducted with game 2. These five observations can be ranked on the basis of the proportion of sequential outcomes. The number of possible rankings of this measure for treatment and control sessions is 5!/(3!2!), or 10. Under a null hypothesis that the proportions of sequential equilibrium outcomes for treatment and control sessions are drawn from the same distribution, each of these rankings is equally likely. The most extreme of these 10 outcomes was observed, since each of the treatment sessions yielded lower proportions of sequential outcomes than the control sessions, so we reject the null hypothesis at the 10% level for a one-tailed test.[15]

The weakness of these arguments is that we are pooling observations from different games, and therefore, our conclusions are tentative. On the other hand, the fact that we both replicate the results of others and that we see divergence from the preponderance of sequential outcomes in different games gives us more confidence in the generality of our results.

6. CONCLUSIONS

The experimental evidence presented here shows how specific features of a game affect the pattern of adjustment and, consequently, how these features lead to more or less refined equilibria. We consider two types of contexts: (1) games with two Nash equilibria, of which only one is sequential, and (2) games with two sequential equilibria, of which only one is intuitive. The evidence contained in this paper shows how the subjects' experience in the process of adjustment may lead to behavior that is inconsistent with the most popular theoretical refinements of the Nash concept, the sequential and intuitive equilibrium concepts.

As with the previous experiment reported by Banks, Camerer, and Porter (1990), we find that the sequential equilibrium outcome is at least three times as likely as the nonsequential Nash outcome in the last half of the matchings for game 1. In a second game, which has the same equilibrium configuration as game 1, the nonsequential Nash outcome is more common than the

[15] The same null hypothesis can be rejected at the 5% level if the variation of game 2 described in footnote 11 is included as a treatment observation.

sequential equilibrium outcome, which we believe is due to a payoff structure that keeps a significant fraction of subjects from experimenting with the sequential-message side of the game tree. Although aggregate behavior does not converge to the sequential equilibrium in this second game, there is a significant fraction of sequential outcomes in later matchings. Nevertheless, the refinement does not provide a good qualitative prediction for this game.

We also replicate the Banks, Camerer, and Porter results for a third game in which a high proportion of the outcomes correspond to an intuitive equilibrium, and few nonintuitive, sequential equilibrium outcomes are observed. In a fourth game with the same equilibrium configuration, the relative drawing powers of the intuitive and (nonintuitive) sequential equilibria are almost exactly reversed. In the latter game, the responses to deviations are consistent with beliefs “off of the equilibrium path” that are ruled out by the intuitive criterion. These “unreasonable” beliefs are induced by an adjustment process that begins with relatively naive behavior and develops as subjects respond to patterns observed during the adjustment. Once a behavioral equilibrium is reached, a subject’s beliefs about the type of person who would deviate are influenced by the type of person who tended to send the deviant message in earlier matchings. We conclude that the process of adjustment provides a better explanation of actual belief formation than abstract deductive arguments that assume a specific equilibrium and consider what beliefs are reasonable for deviations from that equilibrium.

To conclude, one of the most unsettling aspects of economic theory is the absence of a widely accepted theory of adjustment to equilibrium. Laboratory experiments can provide data that shed light on the appropriate way to proceed. We do not present a formal model of adjustment in this paper, but rather, we use observed regularities in behavior to generate different adjustment paths. The data generated by these experiments may be useful to theorists in the development of more formal models of out-of-equilibrium learning and adjustment.

REFERENCES

Banks, Jeffrey S., Camerer, Colin F. and Porter, David (1990) "Experimental Tests of Nash Refinements in Signaling Games," unpublished manuscript, forthcoming in *Games and Economic Behavior*.

Brandts, Jordi and Holt, Charles A. (1992) "An Experimental Test of Equilibrium Dominance in Signaling Games," *American Economic Review*, *82*, *1350-1365*.

Cooper, Russell W., DeJong, Douglas V., Forsythe, Robert, and Ross, Thomas W. (1991) "Cooperation without Reputation", working paper, University of Iowa.

Cho, In-Koo and Kreps, David M. (1987) "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, *102*, 179-221.

Crawford, Vincent P. (1991) "An 'Evolutionary' Interpretation of Van Huyck, Battalio, and Beil's Experimental Results on Coordination," *Games and Economic Behavior*, 3, 25-29.

Friedman, Dan (1991) "Evolutionary Games in Economics," *Econometrica*, 59, 637-666.

Friedman, Dan (1992) "Evolutionary Games: Some Experimental Results," working paper, University of California at Santa Cruz.

Harrison, Glenn W. and McCabe, Kevin (1992) "Testing Bargaining Theory in Experiments," in M. Isaac, ed., *Research in Experimental Economics*, vol. 5, Greenwich, Conn., JAI Press.

Jordan, James S. (1991) "Bayesian Learning in Normal Form Games," *Games and Economic Behavior*, 3, 60-81.

Kalai, Ehud and Ehud Lehrer (1991) "Rational Learning Leads to Nash Equilibrium," Center for Mathematical Studies in Economics and Management Science Discussion Paper No. 925, Northwestern University.

Kohlberg, Elon and Mertens, Jean-Francois (1986) "On the Strategic Stability of Equilibria," *Econometrica*, *54*, 1003-37.

Mailath, George (1992) "Introduction: Symposium on Evolutionary Game Theory," *Journal of Economic Theory*, 57, 259-277.

Marimon, Ramon and Ellen McGrattan (1992) "On Adaptive Learning in Strategic Games," forthcoming in A. Kirman and M. Salmon, eds., *Learning and Rationality in Economics*, Basil Blackwell.

Selten, Reinhard and Rolf Stoecker (1986) "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames," *Journal of Economic Behavior and Organization*, 7, 47-70.

Appendix A: Instructions For Game 1

Introduction

You are going to take part in an experimental study of decision making. The funding for this study has been provided by several foundations. The instructions are simple, and by following them carefully, you may earn a considerable amount of money. At this time, you will be given $10.00 for coming on time. All the money that you earn subsequently will be yours to keep, and your earnings will be paid to you in cash today at the end of this experiment. We will start by reading the instructions, and then you will have the opportunity to ask questions about the procedures described.

Before we proceed, we will choose a "supervisor". The supervisor will observe the way in which we perform the experiment and will help us at several moments. The supervisor will earn the maximum of the earnings that the remainder of you have earned at the end of the experiment. Now, we will assign a number to each of you, and then the supervisor will be chosen by the throw of two dice.

Earnings

The experiment consists of a series of separate periods in which each of you must choose among several alternatives. Each period consists of two parts. In part I you will be matched with another participant, and the decision that you and the other participant make will determine the "points" earned by each of you. In part II, you will have a chance of winning an amount of money in accordance with the points you earned in Part I. We will start by describing Part II so that you will understand how the points that you earn determine your earnings in dollars and cents. We will subsequently describe Part I in detail so that you can understand how the points are earned.

Instructions for Part II

At the end of Part I, you will have earned between 0 and 100 points, according to the rules to be explained later. The amount of money that you earn in Part II will depend partly upon the number of points you won in Part I and partly on chance. Here we have two dice, each with ten sides that are equally likely. In Part II, these two dice will be thrown to determine a number between 0 and 99. The white die will determine the "tens" digit and the red die will determine the "units" digit. If the result of the dice throw is a number smaller than the number of points you won in Part I, you will earn $2.00. If the number resulting from the throw is greater than or equal to the number of points you won in Part I, you will earn $0.50. In this way, if you have earned 100 points in a period, you will be assured of cash earnings of $2.00. If you have earned 0 points in a period, you will be assured of cash earnings of $0.50. Note that the more points you earn, the higher your chances will be to earn $2.00 for the period.

Summary

The period begins.
Part I: Decisions determine the points that you earn.
Part II: The result of throwing the dice determines your earnings in cash:
    $2.00 if the result is < than the points you earned.
    $0.50 if the result is ≥ than the points you earned.
The period ends, and the following period begins.

Instructions for Part I

In each period, you will be matched with another participant; one of you will be referred to as the "proponent", and the other will referred to as the "respondent". At the beginning of each period the proponent will choose between two possibilities, N or S. Then this decision will be communicated to the respondent. After learning about the decision made by the proponent, the respondent will make a decision by choosing among different possibilities: C, D, and E. Next, the decision of the respondent will be communicated to the proponent. Given these decisions, the earnings in points will be determined by one of the following two tables:

TABLE A (selected with probability 1/2)

Proponent's Decision	Respondent's Decision	Proponent's Earnings (in points)	Respondent's Earnings (in points)
N	C	15	30
S	C	45	15
N	D	30	30
S	D	0	0
N	E	0	45
S	E	30	15

TABLE B (selected with probability 1/2)

Proponent's Decision	Respondent's Decision	Proponent's Earnings (in points)	Respondent's Earnings (in points)
N	C	30	30
S	C	30	30
N	D	15	60
S	D	0	0
N	E	45	30
S	E	30	15

Note that the earnings will depend on two decisions made and on the table of earnings being used. The table of earnings, A or B, will be chosen by a throw of a 6-sided die; a 1, 2, or 3 yields table B and a 4, 5, or 6 yields A. The proponent will observe the die throw, and therefore will know the table employed before making his/her decision. On the other hand, the respondent will not know which table is being employed until after his/her decision is made.

Supervisor

At the beginning of each period, the supervisor throws the die that determines the table of earnings in points to be employed in this period. The supervisor may communicate with the organizers of the experiment in order to comment on the procedures, but he/she will not be allowed to communicate with the participants in the experiment, unless there is a question about procedures being followed.

Record of Results

Now, each of you will examine the record sheet for part A. This sheet is the last one attached to these instructions. Your identification number is written in the top part of this sheet. The next line indicates your designation: proponent or respondent. All of the proponents for this experiment will be in the same room, and an equal number of the respondents will be in an adjacent room.

Now examine the column on the left side of your record sheet for Part A. The number in this column shows the period (the first row corresponds to period 1 and so on). Going from left to right, you will see the column titled "other participant" and the number in this column shows the identification number of the participant that you are matched with in each period. Note that in each period you are matched with a different participant. To the right, you can see columns for your decision, the other participant's decision, and the earnings table. You must record your decision as soon as you make it.

Once you know the decision of the other participant and the relevant earnings table, you will be asked to record this information in the suitable column, and your earnings in points for each period in the "your earnings" column. Then the throw of dice will determine your earnings in dollars and cents. Both the result of the throw of the dice and the amount earned ($2.00 or $0.50) will be recorded in the appropriate columns of your record sheet.

More on the Pairing of Proponents and Respondents

The box below shows the initial matching of subject numbers for the first period.

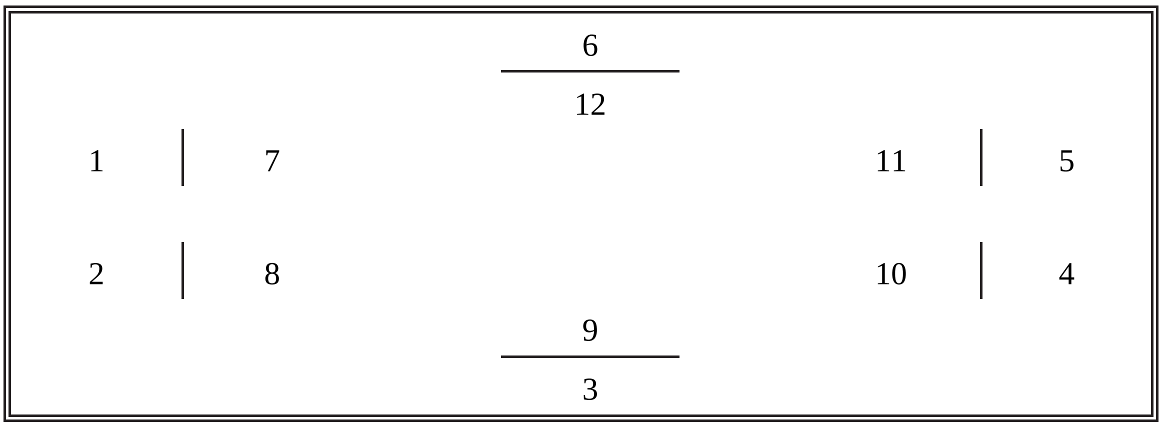

The proponents' numbers are 1, 2, ... 6, so you can think of the proponents as being located on the outside of a circle, with the respondents (7, 8, ...12) located on the inside. After each period the proponents move to the next respondent in a clockwise direction. For example, proponent 1 is matched with respondent 7 in the first period, with respondent 12 in the second period, etc. Respondent 7 is matched with proponent 1 in the 1st period and with proponent 2 in the second period, etc. The 6 periods in this experiment allow each proponent to go around the circle. Notice that no proponent is ever matched with the same respondent twice, and no respondent is matched with the same proponent twice. Moreover, no person is ever matched with anyone who has ever been matched with someone who has been matched with them previously. If you were able to tell a story to every person that you meet, and they told it to everyone they met, etc., then you would never meet anyone who had heard the story. The circular movement in the pairing process causes all participants who have had contact with you, either direct or indirect, to be located "behind" you in one direction around the circle.

Summary

Each period begins with the throw of a 6-sided die that determines the table of earnings that will be employed during this period. Once the proponent has observed the die throw, he/she will make his/her decision and record it. Afterwards, the respondent will be informed of the decision made by the proponent, and the respondent in turn will make a decision, without knowing the result of the die throw. Finally, the decision of the respondent will be shown to the proponent, and the result of the die throw that determines the earnings table will be communicated to the respondent. With this information, each participant will be able to use the appropriate earnings table to determine his/her earnings in points for this period. Then one of us will come to the table of each participant to throw the 10-sided dice separately for each person, and the resulting number will determine whether that person's earnings are $2.00 or $0.50. Finally, each participant will register this amount in the square of the right side corresponding to the period that has just finished. Later periods will be performed in the same way until the end of period 6, at which time we will read to you the instructions for another experiment on decision making.

Final Remarks

At the end of today's session, we will pay to you in private the amount that you have earned. You have already received the $10.00 participation payment. Therefore, if you earn an amount X during the following experiments, your total earnings for today's session will be $10.00 + X. Your earnings are your own business, and you do not have to discuss them with anyone. During the experiment, you are not permitted to speak or communicate with the other participants. If you have a question while the experiment is going on, please raise your hand and one of us will come to your desk to answer it. At this time, do you have any questions about the instructions or procedures?

The experiment is about to begin. You are not permitted to speak with other participants during the experiment. Please, don't ask public questions during the experiment. If you need to ask anything during the experiment, please raise your hand and one of us will come to your desk to answer it. Now please change to the seat that will be indicated to you.

INSTRUCTIONS (for part B)

Please look now at the new record sheet for Part B which pertains to this final part of the experiment, and which is attached. At the top part of the sheet, you will find your new identification number. Your role has changed: those who were proponents of Part A will now be respondents, and those who were respondents in Part A will now be proponents. Although your identification number has changed, the circular rotation scheme is the same as used as before, so you will be matched with a different person in every period, as can be seen in the two left columns of the new record sheet. This part will be just like Part A, with the exception that the roles have changed. The earnings tables are the same as before: (The same earnings tables are presented again.)

APPENDIX B: DATA FOR GAMES 1-5

Data for Game 1

periods 1-6

N	C	D	E
A	-	9	1
B	-	14	2

S	C	D	E
A	20	-	3
B	22	-	1

periods 7-12

N	C	D	E
A	-	1	1
B	-	13	-

S	C	D	E
A	34	-	-
B	22	-	1

Data for Game 2

periods 1-6

N	C	D	E
A	2	20	5
B	1	16	2

S	C	D	E
A	7	1	5
B	9	-	4

periods 7-12

N	C	D	E
A	-	24	2
B	-	12	5

S	C	D	E
A	16	-	3
B	11	-	-

Data for Game 3

periods 1-6

I	C	D	E
A	28	1	-
B	6	-	1

S	C	D	E
A	3	5	-
B	11	16	1

periods 7-12

I	C	D	E
A	30	2	-
B	18	-	-

S	C	D	E
A	-	1	-
B	4	17	-

Data for Game 4

periods 1-6

I	C	D	E
A	-	1	-
B	10	16	3

S	C	D	E
A	27	-	3
B	10	-	2

periods 7-12

I	C	D	E
A	-	-	-
B	1	13	1

S	C	D	E
A	34	1	3
B	18	1	-

Data for Game 5

periods 1-6

I	C	D	E
A	2	3	-
B	7	20	1

S	C	D	E
A	25	1	5
B	8	0	1

periods 7-12

I	C	D	E
A	2	-	-
B	4	16	-

S	C	D	E
A	24	1	7
B	12	1	5