

An Experimental Study of the Centipede Game

Author(s): Richard D. McKelvey and Thomas R. Palfrey

Source: *Econometrica*, Jul., 1992, Vol. 60, No. 4 (Jul., 1992), pp. 803-836

Published by: The Econometric Society

Stable URL: <https://www.jstor.org/stable/2951567>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



The Econometric Society is collaborating with JSTOR to digitize, preserve and extend access to *Econometrica*

JSTOR

AN EXPERIMENTAL STUDY OF THE CENTIPEDE GAME

BY RICHARD D. MCKELVEY AND THOMAS R. PALFREY¹

We report on an experiment in which individuals play a version of the centipede game. In this game, two players alternately get a chance to take the larger portion of a continually escalating pile of money. As soon as one person takes, the game ends with that player getting the larger portion of the pile, and the other player getting the smaller portion. If one views the experiment as a complete information game, all standard game theoretic equilibrium concepts predict the first mover should take the large pile on the first round. The experimental results show that this does not occur.

An alternative explanation for the data can be given if we reconsider the game as a game of incomplete information in which there is some uncertainty over the payoff functions of the players. In particular, if the subjects believe there is some small likelihood that the opponent is an altruist, then in the equilibrium of this incomplete information game, players adopt mixed strategies in the early rounds of the experiment, with the probability of taking increasing as the pile gets larger. We investigate how well a version of this model explains the data observed in the centipede experiments.

KEYWORDS: Game theory, experiments, rationality, altruism.

1. OVERVIEW OF THE EXPERIMENT AND THE RESULTS

THIS PAPER REPORTS THE RESULTS of several experimental games for which the predictions of Nash equilibrium are widely acknowledged to be intuitively unsatisfactory. We explain the deviations from the standard predictions using an approach that combines recent developments in game theory with a parametric specification of the errors individuals might make. We construct a structural econometric model and estimate the extent to which the behavior is explainable by game-theoretic considerations.

In the games we investigate, the use of backward induction and/or the elimination of dominated strategies leads to a unique Nash prediction, but there are clear benefits to the players if, for some reason, some players fail to behave in this fashion. Thus, we have intentionally chosen an environment in which we expect Nash equilibrium to perform at its worst. The best known example of a game in this class is the finitely repeated prisoners' dilemma. We focus on an even simpler and, we believe more compelling, example of such a game, the closely related alternating-move game that has come to be known as the "centipede game" (see Binmore (1987)).

The centipede game is a finite move extensive form two person game in which each player alternately gets a turn to either terminate the game with a favorable payoff to itself, or continue the game, resulting in social gains for the pair. As

¹Support for this research was provided in part by NSF Grants #IST-8513679 and #SES-878650 to the California Institute of Technology. We thank Mahmoud El-Gamal for valuable discussions concerning the econometric estimation, and we thank Richard Boylan, Mark Fey, Arthur Lupia, and David Schmidt for able research assistance. We thank the JPL-Caltech joint computing project for granting us time on the CRAY X-MP at the Jet Propulsion Laboratory. We also are grateful for comments and suggestions from many seminar participants, from an editor, and from two very thorough referees.

far as we are aware, the centipede game was first introduced by Rosenthal (1982), and has subsequently been studied by Binmore (1987), Kreps (1990), and Reny (1988). The original versions of the game consisted of a sequence of a hundred moves (hence the name “centipede”) with linearly increasing payoffs. A concise version of the centipede game with exponentially increasing payoffs, called the “Share or Quit” game, is studied by Megiddo (1986), and a slightly modified version of this game is analyzed by Aumann (1988). It is this exponential version that we study here.

In Aumann’s version of the centipede game, two piles of money are on the table. One pile is larger than the other. There are two players, each of whom alternately gets a turn in which it can choose either to take the larger of the two piles of money or to pass. When one player takes, the game ends, with the player whose turn it is getting the large pile and the other player getting the small pile. On the other hand, whenever a player passes, both piles are multiplied by some fixed amount, and the play proceeds to the next player. There are a finite number of moves to the game, and the number is known in advance to both players. In Aumann’s version of the game, the pot starts at \$10.50, which is divided into a large pile of \$10.00 and a small pile of \$.50. Each time a player passes, both piles are multiplied by 10. The game proceeds a total of six moves, i.e., three moves for each player.

It is easy to show that any Nash equilibrium to the centipede game involves the first player taking the large pile on the first move—in spite of the fact that in an eight move version of the game, both players could be multi-millionaires if they were to pass every round. Since all Nash equilibria make the same outcome prediction, clearly any of the usual refinements of Nash equilibrium also make the same prediction. We thus have a situation where there is an unambiguous prediction made by game theory.

Despite the unambiguous prediction, game theorists have not seemed too comfortable with the above analysis of the game, wondering whether it really reflects the way in which anyone would play such a game (Binmore (1987), Aumann (1988)). Yet, there has been no previous experimental study of this game.²

In the simple versions of the centipede game we study, the experimental outcomes are quite different from the Nash predictions. To give an idea how badly the Nash equilibrium (or iterated elimination of dominated strategies) predicts outcomes, we find only 37 of 662 games end with the first player taking the large pile on the first move, while 23 of the games end with both players passing at every move. The rest of the outcomes are scattered in between.

One class of explanations for how such apparently irrational behavior could arise is based on reputation effects and incomplete information.³ This is the approach we adopt. The idea is that players believe there is some possibility

²There is related experimental work on the prisoner’s dilemma game by Selten and Stoecker (1986) and on an ultimatum bargaining game with an increasing cake by Guth et al. (1991).

³See Kreps and Wilson (1982a), Kreps et al. (1982), Fudenberg and Maskin (1986), and Kreps (1990, pp. 536–543).

that their opponent has payoffs different from the ones we tried to induce in the laboratory. In our game, if a player places sufficient weight in its utility function on the payoff to the opponent, the rational strategy is to always pass. Such a player is labeled an *altruist*.⁴ If it is believed that there is some likelihood that each player may be an altruist, then it can pay a selfish player to try to mimic the behavior of an altruist in an attempt to develop a reputation for passing. These incentives to mimic are very powerful, in the sense that a very small belief that altruists are in the subject pool can generate a lot of mimicking, even with a very short horizon.

The structure of the centipede game we run is sufficiently simple that we can solve for the equilibrium of a parameterized version of this reputational model. Using standard maximum likelihood techniques we can then fit this model. Using this assumption of only a single kind of deviation from the “selfish” payoffs normally assumed in induced-value theory⁵ we are able to fit the data well, and obtain an estimate of the proportion of altruistic players on the order of 5 percent of the subject pool. In addition to estimating the proportion of altruists in the subject pool, we also estimate the beliefs of the players about this proportion. We find that subjects’ beliefs are, on average, equal to the estimated “true” proportion of altruists, thus providing evidence in favor of a version of rational expectations. We also estimate a decision error rate to be on the order of 5%–10% for inexperienced subjects and roughly two-thirds that for experienced subjects, indicating two things: (i) a significant amount of learning is taking place, and (ii) even with inexperienced subjects, only a small fraction of their behavior is unaccounted for by a simple game-theoretic equilibrium model in which beliefs are accurate.

Our experiment can be compared to that of Camerer and Weigelt (1988) (see also Neral and Ochs (1989) and Jung et al. (1989)). In our experiment, we find that many features of the data can be explained if we assume that there is a belief that a certain percentage of the subjects in the population are altruists. This is equivalent to asserting that subjects did not believe that the utility functions we attempted to induce are the same as the utility functions that all subjects really use for making their decisions. I.e., subjects have their own personal beliefs about parameters of the experimental design that are at odds with those of the experimental design. This is similar to the point in Camerer and Weigelt, that one way to account for some features of behavior in their experiments was to introduce “homemade priors”—i.e., beliefs that there were more subjects who always act cooperatively (similar to our altruists) than were actually induced to be so in their experimental design. (They used a rule-of-thumb procedure to obtain a homemade prior point estimate of 17%.) Our analysis differs from Camerer and Weigelt partly in that we integrate it into a structural econometric model, which we then estimate using classical tech-

⁴We called them “irrationals” in an earlier version of the paper. The equilibrium implications of this kind of incomplete information and altruism has been explored in a different kind of experimental game by Palfrey and Rosenthal (1988). See also Cooper et al. (1990).

⁵See Smith (1976).

niques. This enables us to estimate the number of subjects that actually behave in such a fashion, and to address the question as to whether the beliefs of subjects are on average correct.

Our experiment can also be compared to the literature on repeated prisoner’s dilemmas. This literature (see e.g., Selten and Stoecker (1986) for a review) finds that experienced subjects exhibit a pattern of “tacit cooperation” until shortly before the end of the game, when they start to adopt noncooperative behavior. Such behavior would be predicted by incomplete information models like that of Kreps et al. (1982). However, Selten and Stoecker also find that inexperienced subjects do not immediately adopt this pattern of play, but that it takes them some time to “learn to cooperate.” Selten and Stoecker develop a learning theory model that is not based on optimizing behavior to account for such a learning phase. One could alternatively develop a model similar to the one used here, where in addition to incomplete information about the payoffs of others, all subjects have some chance of making errors, which decreases over time. If some other subjects might be making errors, then it could be in the interest of all subjects to take some time to learn to cooperate, since they can masquerade as slow learners. Thus, a natural analog of the model used here might offer an alternative explanation for the data in Selten and Stoecker.

2. EXPERIMENTAL DESIGN

Our budget is too constrained to use the payoffs proposed by Aumann. So we run a rather more modest version of the centipede game. In our laboratory games, we start with a total pot of \$.50 divided into a large pile of \$.40 and a small pile of \$.10. Each time a player chooses to pass, both piles are multiplied by two. We consider both a two round (four move) and a three round (six move) version of the game. This leads to the extensive forms illustrated in Figures 1 and 2. In addition, we consider a version of the four move game in which all payoffs are quadrupled. This “high payoff” condition therefore produced a payoff structure equivalent to the last four moves of the six move game.




FIGURE 1.—The four move centipede game.




FIGURE 2.—The six move centipede game.

TABLE I
EXPERIMENTAL DESIGN

Session #	Subject pool	# subjects	Games/subject	Total # games	# moves	High Payoffs
1	PCC	20	10	100	4	No
2	PCC	18	9	81	4	No
3	CIT	20	10	100	4	No
4	CIT	20	10	100	4	Yes
5	CIT	20	10	100	6	No
6	PCC	18	9	81	6	No
7	PCC	20	10	100	6	No

In each experimental session we used a total of twenty subjects, none of whom had previously played a centipede game. The subjects were divided into two groups at the beginning of the session, which we called the Red and the Blue groups. In each game, the Red player was the first mover, and the Blue player was the second mover. Each subject then participated in ten games, one with each of the subjects in the other group.⁶ The sessions were all conducted through computer terminals at the Caltech Laboratory for Experimental Economics and Political Science. Subjects did not communicate with other subjects except through the strategy choices they made. Before each game, each subject was matched with another subject, of the opposite color, with whom they had not been previously matched, and then the subjects who were matched with each other played the game in either Figure 1 and Figure 2 depending on the session.

All details described above were made common knowledge to the players, at least as much as is possible in a laboratory setting. In other words, the instructions were read to the subjects with everyone in the same room (see Appendix B for the exact instructions read to the subjects). Thus it was common knowledge that no subject was ever matched with any other subject more than once. In fact we used a rotating matching scheme which insures that no player i ever plays against a player who has previously played someone who has played someone that i has already played. (Further, for any positive integer n , the sentence which replaces the phrase “who has previously played someone who has played someone” in the previous sentence with n copies of the same phrase is also true.) In principle, this matching scheme should eliminate potential supergame or cooperative behavior, yet at the same time allow us to obtain multiple observations on each individual’s behavior.

We conducted a total of seven sessions (see Table I). Our subjects were students from Pasadena Community College (PCC) and from the California Institute of Technology (CIT). No subject was used in more than one session. Sessions 1–3 involved the regular four move version of the game, session 4

⁶Only one of the three versions of the game was played in a given session. In sessions 2 and 6, not all subjects showed up, so there were only 18 subjects, with 9 in each group, and consequently each subject played only 9 games.

TABLE IIA
PROPORTION OF OBSERVATIONS AT EACH TERMINAL NODE

		Session	<i>N</i>	<i>f</i> ₁	<i>f</i> ₂	<i>f</i> ₃	<i>f</i> ₄	<i>f</i> ₅	<i>f</i> ₆	<i>f</i> ₇
Four Move	1	(PCC)	100	.06	.26	.44	.20	.04		
	2	(PCC)	81	.10	.38	.40	.11	.01		
	3	(CIT)	100	.06	.43	.28	.14	.09		
	Total	1-3	281	.071	.356	.370	.153	.049		
High Payoff	4	(High-CIT)	100	.150	.370	.320	.110	.050		
Six Move	5	(CIT)	100	.02	.09	.39	.28	.20	.01	.01
	6	(PCC)	81	.00	.02	.04	.46	.35	.11	.02
	7	(PCC)	100	.00	.07	.14	.43	.23	.12	.01
	Total	5-7	281	.007	.064	.199	.384	.253	.078	.014

involved the high payoff four move game, and sessions 5-7 involved the six move version of the game. This gives us a total of 58 subjects and 281 plays of the four move game, and 58 subjects with 281 plays of the six move game, and 20 subjects with 100 plays of the high payoff game. Subjects were paid in cash the cumulative amount that they earned in the session plus a fixed amount for showing up (\$3.00 for CIT students and \$5.00 for PCC students).⁷

3. DESCRIPTIVE SUMMARY OF DATA

The complete data from the experiment is given in Appendix C. In Table II, we present some simple descriptive statistics summarizing the behavior of the subjects in our experiment. Table IIA gives the frequencies of each of the terminal outcomes. Thus f_i is the proportion of games ending at the i th terminal node. Table IIB gives the implied probabilities, p_i of taking at the i th decision node of the game. In other words, p_i is the proportion of games among those that reached decision node i , in which the subject who moves at node i chose TAKE. Thus, in a game with n decision nodes, $p_i = f_i / \sum_{j=i}^{n+1} f_j$.

All standard game theoretic solutions (Nash equilibrium, iterated elimination of dominated strategies, maximin, rationalizability, etc.) would predict $f_i = 1$ if $i = 1$, $f_i = 0$ otherwise. The requirement of rationality that subjects not adopt dominated strategies would predict that $f_{n+1} = 0$ and $p_n = 1$. As is evident from Table II, we can reject out of hand either of these hypotheses of rationality. In only 7% of the four move games, 1% of the six move games, and 15% of the high payoff games does the first mover choose TAKE on the first round. So the subjects clearly do not iteratively eliminate dominated strategies. Further, when a game reaches the last move, Table IIB shows that the player with the last move adopts the dominated strategy of choosing PASS roughly 25% of the time

⁷The stakes in these games were large by usual standards. Students earned from a low of \$7.00 to a high of \$75.00, in sessions that averaged less than 1 hour—average earnings were \$20.50 (\$13.40 in the four move, \$30.77 in the six move, and \$41.50 in the high payoff four move version).

TABLE IIB^a
IMPLIED TAKE PROBABILITIES FOR THE CENTIPEDE GAME

	Session	p_1	p_2	p_3	p_4	p_5	p_6
Four Move	1 (PCC)	.06 (100)	.28 (94)	.65 (68)	.83 (24)		
	2 (PCC)	.10 (81)	.42 (73)	.76 (42)	.90 (10)		
	3 (CIT)	.06 (100)	.46 (94)	.55 (51)	.61 (23)		
	Total 1-3	.07 (281)	.38 (261)	.65 (161)	.75 (57)		
High Payoff	4 (CIT)	.15 (100)	.44 (85)	.67 (48)	.69 (16)		
	5 (CIT)	.02 (100)	.09 (98)	.44 (89)	.56 (50)	.91 (22)	.50 (2)
	6 (PCC)	.00 (81)	.02 (81)	.04 (79)	.49 (76)	.72 (39)	.82 (11)
	7 (PCC)	.00 (100)	.07 (100)	.15 (93)	.54 (79)	.64 (36)	.92 (13)
	Total 5-7	.01 (281)	.06 (279)	.21 (261)	.53 (205)	.73 (97)	.85 (26)

^aThe number in parentheses is the number of observations in the game at that node.

in the four move games, 15% in the six move games, and 31% in the high payoff games.⁸

The most obvious and consistent pattern in the data is that in all of the sessions, the probability of TAKE increases as we get closer to the last move (see Table IIB). The only exception to this pattern is in session 5 (CIT) in the last two moves, where the probabilities drop from .91 to .50. But the figure at the last move (.50) is based on only two observations. Thus any model to explain the data should capture this basic feature. In addition to this dominant feature, there are some less obvious patterns of the data, which we now discuss.

Table III indicates that there are some differences between the earlier and later plays of the game in a given treatment which are supportive of the proposition that as subjects gain more experience with the game, their behavior appears "more rational." Recall that with the matching scheme we use, there is no game-theoretic reason to expect players to play any differently in earlier games than in later games. Table IIIA shows the cumulative probabilities, $F_j = \sum_{i=1}^j f_i$ of stopping by the j th node. We see that the cumulative distribution in the first five games stochastically dominates the distribution in the last five games both in the four and six move experiments. This indicates that the games end earlier in later matches. Table IIIB shows that in both the four and six move sessions, in later games subjects chose TAKE with higher probability at all

⁸For sessions 1-3, 7 of the 14 cases in this category are attributable to 2 of the 29 subjects. In the high payoff condition, 4 of the 5 events are attributable to 1 subject.

TABLE IIIA
 CUMULATIVE OUTCOME FREQUENCIES
 ($F_j = \sum_{i=1}^j f_i$)

Treatment	Game	<i>N</i>	F_1	F_2	F_3	F_4	F_5	F_6	F_7
Four Move	1-5	145	.062	.365	.724	.924	1.00		
	6-10	136	.081	.493	.875	.978	1.00		
Six Move	1-5	145	.000	.055	.227	.558	.889	.979	1.000
	6-10	136	.015	.089	.317	.758	.927	.993	1.000

TABLE IIIB
 IMPLIED TAKE PROBABILITIES
 COMPARISON OF EARLY VERSUS LATE PLAYS IN THE LOW PAYOFF CENTIPEDE GAMES

Treatment	Game	p_1	p_2	p_3	p_4	p_5	p_6
Four Move	1-5	.06 (145)	.32 (136)	.57 (92)	.75 (40)		
	6-10	.08 (136)	.49 (125)	.75 (69)	.82 (17)		
Four Move	1-5	.00 (145)	.06 (145)	.18 (137)	.43 (112)	.75 (64)	.81 (16)
	6-10	.01 (136)	.07 (134)	.25 (124)	.65 (93)	.70 (33)	.90 (10)

stages of the game (with the exception of node 5 of the six move games). Further, the number of subjects that adopt the dominated strategy of passing on the last move drops from 14 of 56, or 25%, to 4 of 27, or 15%.

A third pattern emerges in comparing the four move games to the six move games (in Table IIB). We see that at every move, there is a higher probability of taking in the four move game than in the corresponding move of the six move game (.07 vs .01 in the first move; .38 vs .06 in the second move, etc.). The same relation holds between the high payoff games and the six move games. However, if we compare the four move games to the *last* four moves of the six move games, there is more taking in the six move games (.75 vs .85 in the last move; .65 vs .73 in the next to last move, etc.). This same relationship holds between the high payoff games and the six move games even though the payoffs in the high payoff games are identical to the payoffs in the last four moves of the six move games.

There is at least one other interesting pattern in the data. Specifically, if we look at individual level data, there are several subjects who PASS at every opportunity they have.⁹ We call such subjects *altruists*, because an obvious way to rationalize their behavior is to assume that they have a utility function that is

⁹Some of these subjects had as many as 24 opportunities to TAKE in the 10 games they played. See Appendix C.

monotonically increasing in the *sum* of the red and blue payoffs, rather than a selfish utility function that only depends on that players' own payoff. Overall, there were a total of 9 players who chose PASS at every opportunity. Roughly half (5) of these were red players and half (5) were in four move games. At the other extreme (i.e., the Nash equilibrium prediction), only 1 out of all 138 subjects chose TAKE at every opportunity. This indicates the strong possibility that players who will always choose PASS do exist in our subject pool, and also suggests that a theory which successfully accounts for the data will almost certainly have to admit the existence of at least a small fraction of such subjects.

Finally, there are interesting *non-patterns* in the data. Specifically, unlike the ten cases cited above, the preponderance of the subject behavior is inconsistent with the use of a single pure strategy throughout all games they played. For example, subject #8 in session #1 (a red player) chooses TAKE at the first chance in the second game it participates in, then PASS at both opportunities in the next game, PASS at both opportunities in the fourth game, TAKE at the first chance in the fifth game, and PASS at the first chance in the sixth game. Fairly common irregularities of this sort, which appear rather haphazard from a casual glance, would seem to require some degree of randomness to explain. While some of this behavior may indicate evidence of the use of mixed strategies, some such behavior is impossible to rationalize, even by resorting to the possibility of altruistic individuals or Bayesian updating across games. For example, subject #6 in session #1 (a blue player), chooses PASS at the last node of the first game, but takes at the first opportunity a few games later. Rationalization of this subject's behavior as altruistic in the first game is contradicted by the subject's behavior in the later game. Rational play cannot account for some sequences of plays we observe in the data, even with a model that admits the possibility of altruistic players.

4. THE MODEL

In what follows, we construct a structural econometric model based on the theory of games of incomplete information that is simultaneously consistent with the experimental design and the underlying theory. Standard maximum likelihood techniques can then be applied to estimate the underlying structural parameters.

The model we construct consists of an incomplete information game together with a specification of two sources of errors—errors in actions and errors in beliefs. The model is constructed to account for both the time-series nature of our data and for the dependence across observations, features of the data set that derive from a design in which every subject plays a sequence of games against different opponents. The model is able to account for the broad descriptive findings summarized in the previous section. By parameterizing the structure of the errors, we can also address issues of whether there is learning going on over time, whether there is heterogeneity in beliefs, and whether individuals' beliefs are on average correct.

We first describe the basic model, and then describe the two sources of errors.

4.1. *The Basic Model*

If, as appears to be the case, there are a substantial number of altruists in our subject pool, it seems reasonable to assume that the possible existence of such individuals is commonly known by all subjects. Our basic model is thus a game of two sided incomplete information where each individual can be one of two types (selfish or altruistic), and there is incomplete information about whether one's opponent is selfish or altruistic.

In our model, a *selfish* individual is defined as an individual who derives utility only from its own payoff, and acts to maximize this utility. In analogy to our definition of a selfish individual, a natural definition of an altruist would be as an individual who derives utility not only from its own payoff, but also from the payoff of the other player. For our purposes, to avoid having to make parametric assumptions about the form of the utility functions, it is more convenient to define an altruist in terms of the strategy choice rather than in terms of the utility function. Thus, we define an *altruist* as an individual who always chooses PASS. However, it is important to note that we could obtain an equivalent model by making parametric assumptions on the form of the utility functions. For example, if we were to assume that the utility to player i is a convex combination of its own payoff and that of its opponent, then any individual who places a weight of at least $\frac{2}{5}$ on the payoff of the opponent has a dominant strategy to choose PASS in every round of the experiment. Thus, defining altruists to be individuals who satisfy this condition would lead to equivalent behavior for the altruists.

The extensive form of the basic model for the case when the probability of a selfish individual equals q is shown in Figure 8 in Appendix A. Hence, the probability of an altruist is $1 - q$. This is a standard game of incomplete information. There is an initial move by nature in which the types of both players are drawn. If a player is altruistic, then the player has a trivial strategy choice (namely, it can PASS). If a player is selfish, then it can choose either PASS or TAKE.

4.1.1. *Equilibrium of the Basic Model*

The formal analysis of the equilibrium appears in Appendix A, but it is instructive to provide a brief overview of the equilibrium strategies, and to summarize how equilibrium strategies vary over the family of games indexed by q . We analytically derive in Appendix A the solution to the n -move game, for arbitrary values of q (the common knowledge belief that a randomly selected player is selfish) ranging from 0 to 1.

For any given q , a strategy for the first player in a six move game is a vector, (p_1, p_3, p_5) , where p_i specifies the probability that Red chooses TAKE on move i conditional on Red being a selfish player. Similarly, a strategy for Blue is a

vector (p_2, p_4, p_6) giving the probability that Blue chooses TAKE on the corresponding move conditional that Blue is selfish. Thus a strategy pair is a vector $p = (p_1, p_2, \dots, p_6)$, where the odd components are moves by Red and the even components are moves by Blue. Similar notation is used for the four move game.

In Appendix A, we prove that for generic q there is a unique sequential equilibrium to the game, and we solve for this equilibrium as a function of q . Let us write $p(q)$ for the solution as a function of q . It is easy to verify the following properties of $p(q)$, which are true in both the four move and six move games.

Property 1: For any q , Blue chooses TAKE with probability 1 on its last move.

Property 2: If $1 - q > \frac{1}{7}$, both Red and Blue always choose PASS, except on the last move, when Blue chooses TAKE.

Property 3: If $1 - q \in (0, \frac{1}{7})$ the equilibrium involves mixed strategies.

Property 4: If $q = 1$, then both Red and Blue always choose TAKE.

From the solution, we can compute the implied probabilities of choosing TAKE at each move, accounting for the altruists as well as the selfish players. We can also compute the probability $s(q) = (s_1(q), \dots, s_7(q))$ of observing each of the possible outcomes, T, PT, PPT, PPPT, PPPPT, PPPPPT, PPPPPP. Thus, $s_1(q) = qp_1(q)$, $s_2(q) = q^2(1 - p_1(q))p_2(q) + q(1 - q)p_2(q)$, etc. Figures 3 and 4 graph the probabilities of the outcomes as a function of the level of altruism, $1 - q$.

It is evident from the properties of the solution and from the outcome probabilities in Figures 3 and 4 that the equilibrium predictions are extremely sensitive to the beliefs that players have about the proportion of altruists in the population. The intuition of why this is so is well-summarized in the literature on signalling and reputation building (for example, Kreps and Wilson (1982a), Kreps et al. (1982)) and is expounded very nicely for a one-sided incomplete information version of the centipede game more recently in Kreps (1990). The guiding principle is easy to understand, even if one cannot follow the technical details of the Appendix. Because of the uncertainty in the game when it is not common knowledge that everyone is self-interested, it will generally be worthwhile for a selfish player to mimic altruistic behavior. This is not very different from the fact that in poker it may be a good idea to bluff some of the time in order to confuse your opponent about whether or not you have a good hand. In our games, for *any* amount of uncertainty of this sort, equilibrium will involve some degree of imitation. The form of the imitation in our setting is obvious: selfish players sometimes pass, to mimic an altruist. By imitating an altruist one might lure an opponent into passing at the next move, thereby raising one's final payoff in the game. The amount of imitation in equilibrium depends directly on the beliefs about the likelihood $(1 - q)$ of a randomly selected player being an altruist. The more likely players believe there are altruists in the population, the more imitation there is. In fact, if these beliefs are sufficiently high (at least $\frac{1}{7}$, in our versions of the centipede game), then selfish players will always imitate altruists, thereby completely reversing the predictions of game theory when it is




FIGURE 3.—Equilibrium outcome probabilities for basic four move game.




FIGURE 4.—Equilibrium outcome probabilities for basic six move game.

common knowledge that there are no altruists. Between 0 and $\frac{1}{7}$, the theory predicts the use of mixed strategies by selfish players.

4.1.2. *Limitations of the Basic Model*

The primary observation to make from the solution to the basic model is that this model can account for the main feature of the data noted in the previous section—namely that probabilities of taking increase as the game progresses.

For any level of altruism above $1/7^2$ in the four move game, and for any value above $1/7^3$ in the six move game, the solution satisfies the property that $p_i \geq p_j$ whenever $i > j$.

Despite the fact that the basic model accounts for the main pattern in the data, it is just as obvious that the basic model cannot account for the remaining features of the data. It is apparent from Figures 3 and 4 that for any value of q , there is at least one outcome with a 0 or close to 0 probability of occurrence. So the model will fit poorly data in which all of the possible outcomes occur. Nor can it account for any consistent patterns of learning in the data, for substantial variations across individuals, or for some of the irregularities described earlier.

To account for these features of the data, we introduce two additional elements to the model—the possibility of errors in actions, and the possibility of errors in beliefs.

4.2. *Errors in Actions—Noisy Play*

One explanation of the apparently bizarre irregularities that we noted in the previous section is that players may “experiment” with different strategies in order to see what happens. This may reflect the fact that, early on, a subject may not have settled down on any particular approach about how to play the game. Alternatively, subjects may simply “goof,” either by pressing the wrong key, or by accidentally confusing which color player they are, or by failing to notice that is the last round, or some other random event. Lacking a good theory for how and why this experimentation or goofing takes place, a natural way to model it is simply as noise. So we refer to it as *noisy play*.

We model noisy play in the following way. In game t , at node s , if p^* is the equilibrium probability of TAKE that the player at that node attempts to implement, we assume that the player actually chooses TAKE with probability $(1 - \varepsilon_t)p^*$, and makes a random move (i.e. TAKE or PASS with probability .5) with probability ε_t . Therefore, we can view $\varepsilon_t/2$ as the probability that a player experiments, or, alternatively, goofs in the t th game played. We call ε_t the error rate in game t . We assume that both types (selfish and altruistic) of players make errors at this rate, independently at all nodes of game t , and that this is common knowledge among the players.

4.2.1. *Learning*

If the reasons for noisy play are along the lines just suggested, then it is natural to believe that the incidence of such noisy play will decline with experience. For one thing, as experience accumulates, the informational value of experimenting with alternative strategies declines, as subjects gather information about how other subjects are likely to behave. Perhaps more to the point, the informational value will decline over the course of the 10 games a subject plays simply because, as the horizon becomes nearer, there are fewer and fewer games where the information accumulated by experimentation can be capital-

ized on. For different, but perhaps more obvious reasons, the likelihood that a subject will goof is likely to decline with experience. Such a decline is indicated in a wide range of experimental data in economics and psychology, spanning many different kinds of tasks and environments. We call this decline *learning*.

We assume a particular parametric form for the error rate as a function of t . Specifically, we assume that individuals follow an exponential learning curve. The initial error rate is denoted by ε and the learning parameter is δ . Therefore,

$$\varepsilon_t = \varepsilon e^{-\delta(t-1)}.$$

Notice that, while according to this specification the error rate may be different for different t , it is assumed to be the same for all individuals, and the same at all nodes of the game. More complicated specifications are possible, such as estimating different ε 's for altruistic and selfish players, but we suspect that such parameter proliferation would be unlikely to shed much more light on the data. When solving for the equilibrium of the game, we assume that players are aware that they make errors and learn, and are aware that other players make errors and learn too.¹⁰ Formally, when solving for the Bayesian equilibrium TAKE probabilities, we assume that ε and δ are common knowledge.

4.2.2. *Equilibrium with Errors in Actions*

For $\varepsilon > 0$, we do not have an analytical solution for the equilibrium. The solutions were numerically calculated using GAMBIT, a computer algorithm for calculating equilibrium strategies to incomplete information games, developed by McKelvey (1990). For comparison, the equilibrium outcome probabilities as a function of q , for $\varepsilon_t = .2$, are illustrated graphically in Figures 5 and 6.

4.3. *Errors in Beliefs—Heterogeneous Beliefs*

In addition to assuming that individuals can make errors in their strategies, we also assume that there can be errors in their beliefs. Thus, we assume that there is a *true probability* Q that individuals are selfish (yielding probability $1 - Q$ of altruists), but that each individual has a *belief*, q_i , of the likelihood of selfish players, which may be different from the true Q .¹¹ In particular, individuals' beliefs can differ from each other, giving rise to heterogeneous beliefs.

¹⁰An alternative specification would have it common knowledge that subjects believe others make errors, but believe they do not commit these "errors" themselves. Such a model is analytically more tractable and leads to similar conclusions, but seems less appealing on theoretical grounds.

¹¹This is related to the idea proposed independently by Camerer and Weigelt (1988) and Palfrey and Rosenthal (1988), where they posit that subjects' beliefs about the distribution of types may differ from the induced-value distribution of types announced in the instructions.




FIGURE 5.—Equilibrium outcome probabilities for four move game ($\epsilon_i = .2$).




FIGURE 6.—Equilibrium outcome probabilities for six move game ($\epsilon_i = .2$).

For individual i , denote by q_i the belief individual i holds that a randomly selected opponent is selfish. (We assume that each individual maintains its belief throughout all 10 games that it plays.) Because this converts the complete information centipede game into a Bayesian game, it is necessary to make some kind of assumption about the beliefs a player has about its opponent's beliefs, etc. etc. If there were no heterogeneity in beliefs, so that $q_i = q$ for all i , then one possibility is that a player's beliefs are correct—that is, q is common knowledge, and $q = Q$. We call this *rational expectations*. One can then solve for

the Bayesian equilibrium of the game played in a session (which is unique), as a function of ε , δ , t , and q . An analytical solution is derived in Appendix A for the case of $\varepsilon = 0$.

To allow for heterogeneity, we make a parametric assumption that the beliefs of the individuals are independently drawn from a Beta distribution with parameters (α, β) , where the mean of the distribution, q , is simply equal to $\alpha/(\alpha + \beta)$. There are several ways to specify higher order beliefs. One possibility is to assume it is common knowledge among the players that beliefs are independently drawn from a Beta distribution with parameters (α, β) and that the pair (α, β) is also common knowledge among the players. This version of higher order beliefs leads to serious computational problems when numerically solving for the equilibrium strategies. Instead, we use a simpler¹² version of the higher order beliefs, which might be called an *egocentric model*. Each player plays the game as if it were common knowledge that the opponent had the same belief. In other words, while we, the econometricians, assume there is heterogeneity in beliefs, we solve the game in which the players do have heterogeneous beliefs, but believe that everyone's beliefs are alike. This enables us to use the same basic techniques in solving for the Bayesian equilibrium strategies for players with different beliefs as one would use if there were homogeneous beliefs. We can then investigate a weaker form of rational expectations: is the average belief $(\alpha/(\alpha + \beta))$ equal to the true proportion (Q) of selfish players?

Given the assumptions made regarding the form of the heterogeneity in beliefs, the introduction of errors in beliefs does not change the computation of the equilibrium for a given individual. It only changes the aggregate behavior we will expect to see over a group of individuals. For example, at an error rate of $\varepsilon_t = .2$, and parameters α, β for the Beta distribution, we will expect to see aggregate behavior in period t of the six move games which is the average of the behavior generated by the solutions in Figure 6, when we integrate out q with respect to the Beta distribution $B(\alpha, \beta)$.

5. MAXIMUM LIKELIHOOD ESTIMATION

5.1. Derivation of the Likelihood Function

Consider the version of the game where a player draws belief q . For every t , and for every ε_t , and for each of that player's decision nodes, ν , the equilibrium solution derived in the previous section yields a probability that the decision at that node will be TAKE, conditional on the player at that decision node being selfish, and conditional on that player *not* making an error. Denote that probability $p_s(\varepsilon_t, q, \nu)$. Therefore, the probability that a selfish type of that player would TAKE at ν is equal to $P_s(\varepsilon_t, q, \nu) = (\varepsilon_t/2) + (1 - \varepsilon_t)p_s(\varepsilon_t, q, \nu)$, and the probability that an altruistic type of this player would take is $P_a(\varepsilon_t, q, \nu) = \varepsilon_t/2$. For each individual, we observe a collection of decisions that

¹²While it is simpler, it is no less arbitrary. It is a version of beliefs that does not assume "common priors" (Aumann (1987)), but is consistent with the standard formulation of games of incomplete information (Harsanyi (1967–68)).

are made at all nodes reached in all games played by that player. Let N_{it} denote the set of decision nodes visited by player i in the t th game played by i , and let D_{it} denote the corresponding set of decisions made by i at each of those nodes. Then, for any given $(\varepsilon, \delta, q, \nu)$ with $\nu \in N_{it}$ we can compute $P_s(\varepsilon, q, \nu)$ from above, by setting $\varepsilon_t = \varepsilon e^{-\delta(t-1)}$. From this we compute $\pi_{it}(D_{it}; \varepsilon, \delta, q)$, the probability that a selfish i would have made decisions D_{it} in game t , with beliefs q , and noise/learning parameters (ε, δ) , and it equals the product of $P_s(\varepsilon, q, \nu)$ over all ν reached by in game t . Letting D_i denote the set of all decisions by player i , we define $\pi_i^s(D_i; \varepsilon, \delta, q)$ to be the product of the $\pi_{it}(D_{it}; \varepsilon, \delta, q)$ taken over all t . One can similarly derive $\pi_i^a(D_i; \varepsilon, \delta, q)$, the probability that an altruistic i would have made that same collection of decisions. Therefore, if Q is the true population parameter for the fraction of selfish players, then the likelihood of observing D_i , without conditioning on i 's type is given by:

$$\pi_i(D_i; Q, \varepsilon, \delta, q) = Q\pi_i^s(D_i; \varepsilon, \delta, q) + (1 - Q)\pi_i^a(D_i; \varepsilon, \delta, q).$$

Finally, if q is drawn from the Beta distribution with parameters (α, β) , and density $B(q; \alpha, \beta)$, then the likelihood of observing D_i without conditioning on q is given by:

$$s_i(D_i; Q, \varepsilon, \delta, \alpha, \beta) = \int_0^1 \pi_i(D_i; Q, \varepsilon, \delta, q) B(q; \alpha, \beta) dq.$$

Therefore, the log of the likelihood function for a sample of observations, $D = (D_1, \dots, D_I)$, is just

$$L(D; Q, \varepsilon, \delta, \alpha, \beta) = \sum_{i=1}^I \log [s_i(D_i; Q, \varepsilon, \delta, \alpha, \beta)].$$

For any sample of observations, D , we then find the set of parameter values that maximize L . This was done by a global grid search using the Cray X-MP at the Jet Propulsion Laboratory.

5.2. Treatments and Hypotheses

We are interested in testing four hypotheses about the parameters of the theoretical model.

- (1) Errors in action: Is ε significantly different from 0?
- (2) Heterogeneity (errors in beliefs): Is the variance of the estimated distribution of priors significantly different from 0?¹³
- (3) Rational Expectations: Is the estimated value of Q equal to the mean of the estimated distribution of priors, $\alpha/(\alpha + \beta)$?
- (4) Learning: Is the estimated value of δ positive and significantly different from 0?

¹³While homogeneity of beliefs is not strictly nested in the Beta distribution model (since the Beta family does not include degenerate distributions), the homogeneous model can be approximated by the Beta distribution model by constraining $(\alpha + \beta)$ to be greater than or equal to some large number.

The first two hypotheses address the question of whether the two components of error in our model (namely errors in action and errors in belief) are significantly different from 0. The second two hypotheses address issues of what properties the errors have, if they exist.

In our experimental design, there were three treatment variables:

- (1) The length of the game (either four move or six move).
- (2) The size of the two piles at the beginning of the game (either high payoff, (\$1.60, \$.40), or low payoff, (\$.40, \$.10)).
- (3) The subject pool (either Caltech undergraduates (CIT) or Pasadena City College students (PCC)).

We also test whether any of these treatment variables were significant.

5.3. Estimation Results

Table IV reports the results from the estimations. Before reporting any statistical tests, we summarize several key features of the parameter estimates.

First, the mean of the distribution of beliefs about the proportion of altruists in the population in all the estimations was in the neighborhood of 5%. Figure 7 graphs the density function of the estimated Beta distribution for the pooled sample of all experimental treatments for the four and six move experiments, respectively. Second, if one looks at the rational expectations estimates (which constrain $\hat{\mu} = \hat{\alpha}/(\hat{\alpha} + \hat{\beta})$ to equal \hat{Q}), the constrained estimate of the Beta

TABLE IV
RESULTS FROM MAXIMUM LIKELIHOOD ESTIMATION^a

	Treatment	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\mu}$	\hat{Q}	$\hat{\varepsilon}$	$\hat{\delta}$	$-\ln L$
Four Move	unconstrained	42	2.75	.939	.956	.18	.045	327.35
	$\mu = Q$	44	2.75	.941	.941	.18	.045	327.41
	$\delta = 0$	68	2.50	.965	.950	.21	.000	345.08
	$\sigma = 0$	—	—	.972	.850	.23	.020	371.04
Six Move	unconstrained	40	2.00	.952	.904	.06	.030	352.07
	$\mu = Q$	38	2.00	.950	.950	.06	.030	352.76
	$\delta = 0$	34	1.75	.951	.908	.05	.000	371.01
	$\sigma = 0$	—	—	.976	.850	.22	.030	442.96
PCC	unconstrained	42	2.75	.939	.974	.14	.030	464.14
	$\mu = Q$	40	2.75	.936	.936	.11	.040	464.57
	$\sigma = 0$	—	—	.952	.882	.18	.050	508.60
CIT	unconstrained	42	1.25	.971	.880	.22	.040	340.27
	$\mu = Q$	28	1.00	.966	.966	.22	.040	342.57
	$\sigma = 0$	—	—	.994	.750	.27	.010	424.83
	High Payoff	64	2.25	.966	.900	.22	.050	107.11
	All Four Move	48	2.25	.955	.938	.22	.050	435.73
	All Low	28	1.75	.941	.938	.14	.050	702.80
	All Sessions	40	2.00	.952	.930	.18	.050	813.38

^aRows marked $\mu = Q$ report parameter estimates under the rational expectations restriction that $\hat{\alpha}/(\hat{\alpha} + \hat{\beta}) = \hat{Q}$. Rows marked $\delta = 0$ are parameter estimates under the hypothesis of no learning. Rows marked $\sigma = 0$ are parameter estimates under the assumption of no heterogeneity.

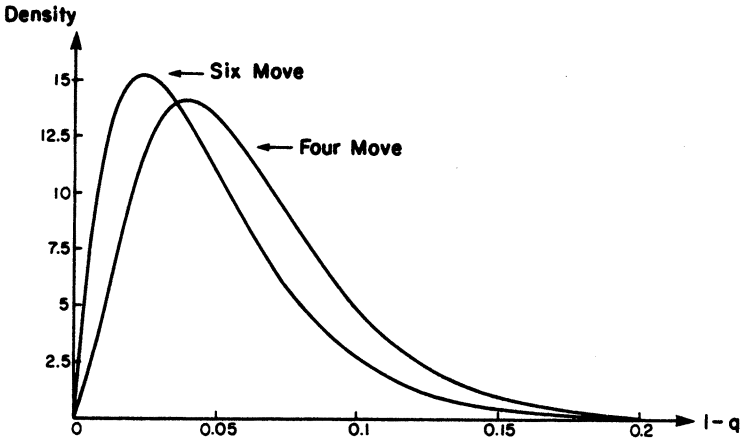


FIGURE 7.—Estimated distribution of beliefs for four and six move games.

distribution is nearly identical to the unconstrained estimate of the Beta distribution.

Furthermore, the rational expectations estimates of μ are nearly identical across all treatments. Therefore, if we are unable to reject rational expectations, then it would seem that these beliefs, as well as the true distribution of altruists are to a large extent independent of the treatments. The greatest difference across the treatments is in the estimates of the amount of noisy play. While the estimates of δ are quite stable across treatments, the estimates of ε are not. This is most apparent in the comparison of the four move and the six move estimates of ε . We discuss this below after reporting statistical tests. Finally, observe that in the $\sigma = 0$ estimates (no heterogeneity of beliefs), the estimate of μ is consistently much larger than the estimate of Q , and the recovered error rates are higher.

It might seem paradoxical that we estimate a level of altruism on the order of 5%, while the proportion of subjects who choose pass on the last move is on the order of 25% for the four move experiments and 15% for the six move experiments. One reason for the low estimate of the level of altruism is that in the theoretical model, part of this behavior is attributed to errors in action. But less obviously, it should be noted that because our equilibrium is a mixed strategy, there is a sample selection bias in the set of subjects who get to the last move: altruists are more likely to make it to this stage, since they pass on earlier moves whereas selfish subjects mix. Thus, even with no errors, we would expect to see a higher proportion of passing on the last move than the proportion of altruists in the subject pool.

5.4. Statistical Tests

Table V reports likelihood ratio χ^2 tests for comparisons of the various treatments, and for testing the theoretical hypotheses of rational expectations, learning, and heterogeneity.

TABLE V
LIKELIHOOD RATIO TESTS

	Hypothesis	Treatment	d.f.	-2 log likelihood ratio
Theoretical Hypotheses	Heterogeneity ($\sigma = 0$)	4-move	1	87.38*
		6-move	1	181.78*
	Rational expectations ($\mu = Q$)	4-move	1	.12
		6-move	1	1.38
	Learning ($\delta = 0$)	4-move	1	35.46*
		6-move	1	37.88*
Treatment Effects	4-move vs 6-move		5	51.16*
	4-high vs 4-low (Payoff Treatment)		5	2.54
	PCC vs CIT	All	5	17.94*
		4-move	5	9.63
		6-move	5	8.42

*Significant at 1% level.

Our first hypothesis—that $\varepsilon = 0$, can be dispensed with immediately. In both the four and six move experiments, setting $\varepsilon = 0$ yields a likelihood of zero. So this hypothesis can be rejected at any level of significance, and is not included in Table V.

To test for heterogeneity, we estimate a model in which a single belief parameter, q , is estimated instead of estimating the two parameter model, α and β (see Table IV, rows marked $\sigma = 0$). As noted earlier, homogeneity is approximately nested in our heterogeneity model. Therefore, we treat the homogeneous model as if it is nested in the heterogeneous model, and report a standard χ^2 test based on likelihood ratios. All statistical tests were highly significant (see Table V). Note that by setting Q and all q_i to one, then one obtains a pure random model where individuals PASS with probability $\varepsilon/2$. Hence for any probability of passing less than or equal to $\frac{1}{2}$, the pure random model is a special case of the homogeneous model. Thus the above findings mean we also reject the pure random model ($\varepsilon = 1$).

In the test for learning, the null hypothesis that $\delta = 0$ is clearly rejected for all treatments, at essentially any level of significance. We conclude that learning effects are clearly identified in the data. Subjects are experimenting less and/or making fewer errors as they gain experience. The magnitude of the parameter estimates of δ indicate that subjects make roughly two-thirds as many errors in the tenth game, compared to the first game.

The test for rational expectations is unrejectable in most cases. The one exception is for the CIT subject pool, where the difference is significant at the 5% level, but not the 1% level.

The payoff level treatment variable is not significant. This is reassuring, as it indicates that the results are relatively robust.

The other treatment effects, CIT/PCC and four move/six move, are both significant. One source of the statistical difference between the PCC and CIT

TABLE VI
CIT-PCC
ESTIMATES, BROKEN DOWN INTO FOUR MOVE AND SIX MOVE TREATMENTS

	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\mu}$	\hat{Q}	$\hat{\varepsilon}$	$\hat{\delta}$	$-\ln L$
PCC 4	74.0	3.75	.95	.996	.210	.04	216.47
CIT 4	36.0	1.50	.96	.866	.220	.05	214.35
PCC 6	40.0	2.50	.94	.906	.060	.04	231.28
CIT 6	80.0	2.25	.97	.902	.060	.03	116.58

TABLE VII
CHI-SQUARED TESTS FOR DIFFERENCES IN ε AND δ ACROSS TREATMENTS
(UNDER ASSUMPTION THAT $\mu = Q$)

Parameter	Treatment	d.f.	$-2 \log L$ ratio
ε	4-move vs. 6-move	1	39.28 ^a
	CIT vs. PCC	1	3.90
δ	4-move vs. 6-move	1	5.76
	CIT vs. PCC	1	.02

^aSignificant at $p = .01$ level.

estimates apparently derives from the fact that two-thirds of the CIT data were for four move games, while only half of the PCC data were for four move games. Consequently, we break down the statistical comparison between PCC and CIT into the four and six move game treatments (see Table VI). The subject pool effect is not significant in the six move treatment and is barely significant at the 10% level in the four move treatment (see Table V).

In order to pin down the source of the treatment effects we performed several tests. The first one was to test for differences in learning effects across treatments. This test is done by simultaneously reestimating the parameters for each of the different treatments, subject to the constraint that δ is the same for each treatment, and then conducting a likelihood ratio test. Second, we tested for differences in the noise parameter, ε . The χ^2 tests are reported in Table VII. The results reflect the estimates in Tables IV and V. The only significant (1% level) difference is the estimated initial error rate ε in the four move versus six move games. The CIT/PCC difference in ε is significant at the 5% level, but this is due to reasons given in the previous paragraph. The difference between δ in the six move game and four move game is significant at the 5% level.

5.5. Fit

In order to get a rough measure of how well our model fits the data, we use the unconstrained parameter estimates from Table IV to obtain predicted aggregate outcomes. Table VIII displays the predicted frequencies of each of the five possible outcomes of the four move game, and compares these frequencies to the observed frequencies. This comparison is done for each of the ten periods, t , $t = 1, \dots, 10$, and for the time-aggregated data. Table IX displays

TABLE VIII
 COMPARISON OF PREDICTED (FIRST) vs.
 ACTUAL (SECOND) FREQUENCIES, FOUR MOVE
 (CELLS 4 AND 5 COMBINED FOR COMPUTATION OF χ^2)

Period	n	ϵ	Predicted					Actual					χ^2
			\hat{f}_1	\hat{f}_2	\hat{f}_3	\hat{f}_4	\hat{f}_5	f_1	f_2	f_3	f_4	f_5	
1	58	.18	.106	.269	.291	.291	.042	.000	.276	.379	.241	.103	7.72
2	58	.17	.101	.292	.298	.271	.037	.103	.276	.345	.241	.034	.69
3	58	.16	.097	.314	.316	.241	.032	.034	.310	.379	.172	.103	3.11
4	58	.16	.096	.314	.316	.241	.032	.103	.276	.379	.172	.069	1.24
5	58	.15	.100	.329	.320	.223	.028	.069	.379	.310	.172	.069	1.04
6	58	.14	.095	.349	.335	.197	.024	.103	.310	.414	.172	.000	2.00
7	58	.14	.095	.349	.335	.197	.024	.069	.414	.448	.069	.000	9.39*
8	58	.13	.090	.369	.338	.181	.021	.069	.414	.345	.138	.034	.87
9	58	.13	.090	.369	.338	.182	.021	.103	.483	.310	.069	.034	5.10
10	40	.12	.095	.380	.348	.159	.017	.050	.450	.400	.050	.050	2.99
Total	562	—	.097	.333	.324	.218	.028	.071	.356	.370	.153	.050	

*Significant at .05 level.

similar numbers for the six move games. The last column of the table displays a χ^2 for each period, comparing the predicted to the actual values.¹⁴

It is evident from Tables VIII and IX that the fit to the four move games is better than that of the six move games. For the four move games, in only one of the periods (period 7), are the predicted and actual frequencies significantly different at the .05 level. In the six move games, there are six periods in which there are differences significant at the .05 level, and four of these differences are significant at the .01 level. The predicted frequencies for the four move games also pick up reasonably well the trends in the data between periods. Thus \hat{f}_2 and \hat{f}_3 increase over time, while \hat{f}_4 decreases, all in accordance with the actual data. In the six move data, on the other hand, the predicted frequencies show very small trends in comparison with those in the actual data. The tables can also be used to help identify where the model seems to be doing badly. In the four move games, the model overestimates f_1 and f_4 and underestimates f_2 and f_3 . In the six move games, the model underestimates f_3 .

The differences in fit between the four and six move games can be accounted for by the limitations of the model we fit. As noted earlier, the main difference between the estimates for the four and six move games is in the estimate of ϵ . In the four move games, we find a high value of $\epsilon = .18$, while in the six move game we find a lower value of $\epsilon = .06$. The reason for the differences in the estimates of ϵ between the four and six move games is fairly clear. In order to obtain a high value of the likelihood function in the six move games, the model is forced to estimate a low error rate, simply because there are almost no observations at the first terminal node.¹⁵

¹⁴The χ^2 statistic is not reported for the time-aggregated data, since the assumption of independence is violated.

¹⁵In fact, there were *no* such observations in the first 9 periods and only 2 in the last period.

TABLE IX
COMPARISON OF PREDICTED (FIRST) VS. ACTUAL (SECOND) FREQUENCIES, SIX MOVE
(CELLS 1 AND 2 COMBINED AND CELLS 6 AND 7 COMBINED FOR COMPUTATION OF χ^2)

Period	n	e	Predicted							Actual							χ^2
			\hat{f}_1	\hat{f}_2	\hat{f}_3	\hat{f}_4	\hat{f}_5	\hat{f}_6	\hat{f}_7	f_1	f_2	f_3	f_4	f_5	f_6	f_7	
1	58	.06	.035	.077	.096	.435	.275	.072	.010	.000	.103	.310	.345	.103	.034	5.33	
2	58	.06	.035	.077	.096	.435	.275	.072	.010	.000	.069	.276	.379	.103	.000	10.41*	
3	58	.06	.035	.077	.096	.435	.275	.072	.010	.000	.069	.345	.345	.069	.069	5.32	
4	58	.05	.030	.076	.113	.432	.273	.068	.009	.000	.034	.414	.207	.103	.000	12.72*	
5	58	.05	.030	.076	.113	.432	.273	.068	.009	.000	.000	.310	.379	.069	.000	18.99**	
6	58	.05	.030	.076	.113	.432	.273	.068	.009	.000	.069	.552	.138	.069	.000	8.39	
7	58	.05	.030	.076	.113	.432	.273	.068	.009	.000	.069	.448	.172	.034	.000	17.98**	
8	58	.05	.030	.076	.113	.432	.273	.068	.009	.000	.069	.345	.241	.034	.034	15.68**	
9	58	.05	.030	.076	.113	.432	.273	.068	.009	.000	.069	.483	.207	.069	.000	3.86	
10	40	.05	.030	.076	.113	.432	.273	.068	.009	.100	.100	.350	.050	.150	.000	20.66**	
Total	562	—	.031	.076	.108	.433	.274	.069	.009	.007	.064	.384	.253	.078	.014		

*Significant at .05 level.
**Significant at .01 level.

The difference in the estimates of ε also accounts for the failure of the model to explain the trends in the six move data. The way that our model explains time trends is through the learning parameter, δ , which represents the rate at which the error rate, ε , declines over time. As ε declines, the model implies that the game will end sooner, by predicting more frequent outcomes at the earlier terminal nodes. In the four move game, the initial error rate is $\varepsilon = .18$. With this high initial error rate, a decay of $\delta = .045$ is able to lead to substantial changes in predicted behavior. However, in the six move games, we estimate a significantly lower $\varepsilon = .06$. With this lower initial error rate, a similar rate of learning will lead to less dramatic changes in behavior.

Our model of errors assumes that in any given game, the error rate is a constant. In other words the likelihood of making an error is the same at each node regardless of whether the equilibrium recommends a pure or mixed strategy at that node. It might be reasonable to assume instead that individuals make more errors when they are indifferent between alternatives (i.e., when the equilibrium recommends a mixed strategy) than when they have preferences over alternatives (i.e., when the equilibrium recommends a pure strategy). In other words, the error rate may be a function of the utility differences between the choices. Such a model might be able to explain the behavior of the six move games with error rates on the same order of magnitude as those of the four move games. This in turn might lead to a better fit, and a better explanation of the trends in the six move games. We have not investigated such a model here because of computational difficulties.

One final point regards the comparisons between the take probabilities in the four move games and the six move games. As noted in Section 3, the aggregate data from the four move games exhibit higher take probabilities than the six move games at each of the decision nodes 1, 2, 3, and 4. But in contrast to this, when one compares the take probabilities of the last four moves of the six move game with the four move game (where the terminal payoffs are exactly comparable), this relationship is reversed; the six move data exhibit the higher take probabilities. Both of these features of the data are picked up in the estimates from our model. From Tables VIII and IX, the predicted take probabilities for the four move and six move games are $\hat{p}_4 = (.097, .368, .568, .886)$, and $\hat{p}_6 = (.031, .078, .120, .552, .778, .885)$, respectively.

6. CONCLUSIONS

We conclude that an incomplete information game that assumes the existence of a small proportion of altruists in the population can account for many of the salient features of our data. We estimate a level of altruism on the order of 5%. In the version of the model we estimate, we allow for errors in actions and errors in beliefs. Both sources of errors are found to be significant. Regarding errors in action, we find that there are significant levels of learning in the data, in the sense that subjects learn to make fewer errors over time. Subjects make roughly two thirds as many mistakes at the end of a session as they make in

early play. Regarding errors in beliefs, we reject homogeneity of beliefs. However, we find that rational expectations, or on-average correct beliefs, cannot be rejected.

While we observe some subject pool differences, they are small in magnitude and barely significant. The payoff treatment had no significant effect. The only significant difference between the parameter estimates of the four and the six move games was that the estimated initial error rate was lower in the six move game. A model in which the error rate is a function of the expected utility difference between the available action choices might well account for the observed behavior in both the four and the six move games with similar estimates of the error rate. This might be a promising direction for future research.

Division of the Humanities and Social Sciences, California Institute of Technology, Pasadena, CA 91125, U.S.A.

Manuscript received May, 1990; final revision received February, 1992.

APPENDIX A

In this appendix, we prove that there is a unique sequential equilibrium to the n move centipede game with two sided incomplete information over the level of altruism (see Figure 8).

There are n nodes, numbered $i = 1, 2, \dots, n$. Player 1 moves at the odd nodes, and Player 2 moves at the even nodes. We use the terminology "player i " to refer to the player who moves at node i . Let the payoff if player i takes at node i be (a_i, b_i) where a_i is payoff to player i and b_i is

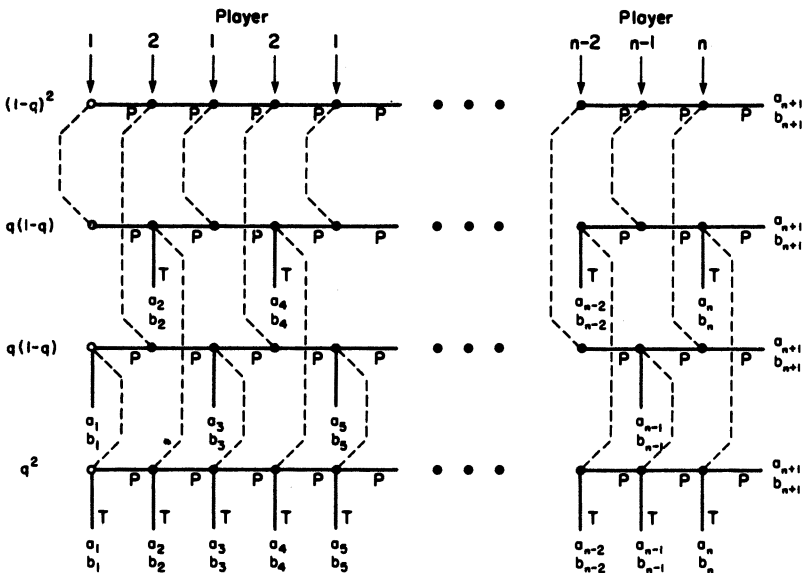


FIGURE 8.—Centipede game with incomplete information (dashed lines represent information sets and open circles are starting nodes, with probabilities indicated).

payoff to player $i - 1$ (or $i + 1$). Also, if $i = n + 1$, then (a_i, b_i) refers to the payoff if player n passes at the last node. Define

$$\eta_i = \frac{a_i - b_{i+1}}{a_{i+2} - b_{i+1}}.$$

We assume that $a_{i+2} > a_i > b_{i+1}$, and that η_i is the same for all i . We write $\eta = \eta_i$. (A similar solution can be derived when the η_i are different.)

Now a strategy for the game can be characterized by a pair of vectors $p = (p_1, \dots, p_n)$ and $r = (r_1, \dots, r_n)$, where for any node i , p_i is the probability that a selfish type takes at that node, and r_i is the conditional probability, as assessed by player i at node i , of the other player being selfish. Let q be the initial probability of selfishness. So $r_1 = q$.

LEMMA 1: *If $p \in \mathcal{P}^n$ and $r \in \mathcal{R}^n$ are a sequential equilibrium, then:*

- (a) for all i , $p_i = 0 \Rightarrow p_j = 0$ for all $j \leq i$;
- (b) $p_n = 1$. Also, $p_i = 1 \Rightarrow i = n$ or $i = n - 1$.

PROOF: (a) Assume $p_i = 0$. Clearly $p_{i-1} = 0$, because at node $i - 1$, the value to player $i - 1$ of passing is at least a_{i+1} , which is by assumption greater than a_{i-1} , the payoff from taking.

(b) By dominance, $p_n = 1$. Suppose $p_i = 1$ for $i < n - 1$. Then by Bayes rule $r_{i+1} = 0$, so $p_{i+1} = 0$ is the best response by player $i + 1$, since i 's type has been revealed. But $p_{i+1} = 0 \Rightarrow p_i = 0$ by part (a), which is a contradiction. *Q.E.D.*

Define \underline{k} to be the first node at which $0 < p_i$, and \bar{k} to be the first node for which $p_i = 1$. Clearly $\underline{k} \leq \bar{k}$. Then \underline{k} and \bar{k} partition the nodes into at most three sets, which we refer to as the *passing stage* ($i < \underline{k}$), the *mixing stage* ($\underline{k} \leq i < \bar{k}$), and the *taking stage* ($\bar{k} \leq i$). From Lemma 1 it follows that there are no pure strategies in the mixing stage. From Lemma 1(b), it follows that the taking stage is at least one and at most two moves.

LEMMA 2: *In any sequential equilibrium (p, r) , for $1 \leq i \leq n - 1$,*

- (a) $0 < p_i < 1 \Rightarrow r_i p_{i+1} = 1 - \eta$;
- (b) $r_i p_{i+1} < 1 - \eta \Rightarrow p_i = 0$;
- (c) $r_i p_{i+1} > 1 - \eta \Rightarrow p_i = 1$.

PROOF: (a) Assume $0 < p_i < 1$. Let v_i be the value to player i at node i given equilibrium play from node i on. Write $v_{n+1} = a_{n+1}$. Now if $i = n - 1$, then $v_{i+2} = v_{n+1} = a_{n+1} = a_{i+2}$. If $i < n - 1$, then by Lemma 1, $0 < p_i \Rightarrow 0 < p_{i+2}$, which implies $v_{i+2} = a_{i+2}$. Now in both cases $0 < p_i$ implies $v_i = a_i = r_i p_{i+1} b_{i+1} + (1 - r_i p_{i+1}) a_{i+2}$. Solving for $r_i p_{i+1}$, we get

$$r_i p_{i+1} = \frac{a_{i+2} - a_i}{a_{i+2} - b_{i+1}} = 1 - \eta.$$

(b) If $r_i p_{i+1} < 1 - \eta$, then at node i , $v_i \geq r_i p_{i+1} b_{i+1} + (1 - r_i p_{i+1}) a_{i+2} = a_{i+2} - r_i p_{i+1} (a_{i+2} - b_{i+1}) > a_{i+2} - (a_{i+2} - a_i) = a_i$. So $v_i > a_i \Rightarrow p_i = 0$.

(c) If $r_i p_{i+1} > 1 - \eta$, then $p_{i+1} > 0 \Rightarrow p_{i+2} > 0$ (by Lemma 1). Hence, $v_{i+2} = a_{i+2}$. By similar argument to (b), $v_i < a_i \Rightarrow p_i = 1$. *Q.E.D.*

LEMMA 3: *For generic q , for any sequential equilibrium there are an even number of nodes in the mixing stage. I.e., $\bar{k} = \underline{k} + 2K$ for some integer $0 \leq K < n/2$. For any $k < K$,*

- (a) $r_{\underline{k}+2k} = 1 - (1 - q)/\eta^k$;
- (b) $r_{\bar{k}-1-2k} = 1 - \eta^{k+1}$.

PROOF: We first show (a) and (b), and then show $\bar{k} = \underline{k} + 2K$.

(a) For any node i , Bayes rule implies

$$(1) \quad r_{i+2} = \frac{(1 - p_{i+1})r_i}{(1 - p_{i+1})r_i + (1 - r_i)} = \frac{r_i - p_{i+1}r_i}{1 - p_{i+1}r_i}.$$

By assumption $r_1 = q$. And since in the passing stage $p_i = 0$, it follows that $r_{\underline{k}} = q$. Now if both i and $i + 2$ are in the mixing stage, it follows from Lemma 1 that $i + 1$ is also, implying $0 < p_{i+1} < 1$. So by Lemma 3, $r_i p_{i+1} = 1 - \eta$. Hence, (1) becomes

$$(2) \quad r_{i+2} = \frac{r_i - (1 - \eta)}{1 - (1 - \eta)} = 1 - \frac{(1 - r_i)}{\eta}.$$

By induction, it follows that as long as $k < \frac{1}{2}(\bar{k} - \underline{k})$, then $\underline{k} + 2k < \bar{k}$ is in the mixing stage. So

$$r_{\underline{k}+2k} = 1 - \frac{1 - r_{\underline{k}}}{\eta^k} = 1 - \frac{1 - q}{\eta^k}.$$

(b) As above, as long as both i and $i - 2$ are in the mixing stage, we get

$$r_i = 1 - \frac{(1 - r_{i-2})}{\eta}.$$

Solving for r_{i-2} , we get

$$r_{i-2} = 1 - \eta(1 - r_i).$$

Now from Lemma 2, it follows, since $p_{\bar{k}} = 1$,

$$r_{\bar{k}-1} = \frac{1 - \eta}{p_{\bar{k}}} = 1 - \eta.$$

Hence, by induction, as long as $k < \frac{1}{2}(\bar{k} - \underline{k})$, we have

$$r_{\bar{k}-1-2k} = 1 - \eta^k(1 - r_{\bar{k}-1}) = 1 - \eta^{k+1}.$$

Finally, to show that there are an even number of nodes in the mixing stage, assume, to the contrary that there are an odd number. Then we can write $\bar{k} = \underline{k} + 2k + 1$ for some $k \geq 0$. Thus $\underline{k} = \bar{k} - 1 - 2k$. So by part (b) we have $r_{\bar{k}} = 1 - \eta^{k+1}$. But by (a) we have $r_{\bar{k}} = 1 - (1 - q) = q$, implying that $q = 1 - \eta^{k+1}$. For generic q , this is a contradiction, implying that $\bar{k} = \underline{k} + 2K$ for some $K \geq 0$. If $k \geq n/2$, then $\bar{k} \geq \underline{k} + n > n$, which contradicts Lemma 1(b). Hence $\bar{k} = \underline{k} + 2K$ for some $0 \leq K < n/2$. Q.E.D.

THEOREM 4: For generic q , there is a unique sequential equilibrium (p, r) which is characterized as follows: Let I be the smallest integer greater than or equal to $n/2$. If $1 - q < \eta^I$, set $K = I - 1$, $\underline{k} = 1$, and $\bar{k} = 2I - 1$. If $1 - q > \eta^I$, let K be the largest integer with $1 - q < \eta^K$, $\bar{k} = n$, and $\underline{k} = \bar{k} - 2K$. The solution then satisfies:

- (a) if $i < \underline{k}$, then $r_{i+1} = q$ and $p_i = 0$;
- (b) if $i \geq \bar{k}$, then $r_{i+1} = 0$ and $p_i = 1$;
- (c) if $\underline{k} \leq i < \bar{k}$: (i) if $i = \underline{k}$, then $r_{i+1} = 1 - \eta^K$, and $p_i = (q + \eta^K - 1)/q\eta^K$; (ii) if $i = \underline{k} + 2k$, with $1 \leq k < K$, then $r_{i+1} = 1 - \eta^{K-k}$, and $p_i = (1 - \eta)/(1 - \eta^{K+1-k})$; (iii) if $i = \underline{k} + 2k + 1$, with $0 \leq k < K$, then $r_{i+1} = 1 - (1 - q)/\eta^{k+1}$, and $p_i = \eta^k(1 - \eta)/(\eta^k - (1 - q))$.

PROOF: The formulae for r_i and p_i in parts (a), (b), and (c) follow by application of the previous lemmas together with Bayes rule. In particular in (a), $p_i = 0$ follows from the definition of \underline{k} , and $r_{i+1} = q$ follows from $p_i = 0$ for $j \leq i$ together with Bayes rule. In (b), $p_i = 1$ follows from the definition of \bar{k} , and $r_{i+1} = 0$ then follows by Bayes rule. In (c), all the formulae for r_{i+1} follow from Lemma 3. In (c) part (i), we set $\underline{k} = i + 1$ in (1) and solve for $p_{\underline{k}}$ to get

$$p_{\underline{k}} = \frac{r_{\underline{k}-1} - r_{\underline{k}+1}}{r_{\underline{k}-1} - r_{\underline{k}-1}r_{\underline{k}+1}}.$$

But $r_{\underline{k}-1} = q$ and $r_{\underline{k}+1} = 1 - \eta^K$. So

$$p_{\underline{k}} = \frac{q - 1 + \eta^K}{q\eta^K}.$$

In parts (ii) and (iii) of (c), we apply Lemma 2a to get that $p_i = (1 - \eta)/(r_{i-1})$. Substituting in for the values of r_{i-1} gives the required formulae.

Thus, it only remains to prove the assertions about \bar{k} and \underline{k} . We first prove two preliminary inequalities. First, note, $\underline{k} > 1$ implies, by Lemma 2,

$$\begin{aligned} p_{\underline{k}-1} = 0 &\Rightarrow r_{\underline{k}-1} p_{\underline{k}} \leq 1 - \eta \\ &\Rightarrow q \left(\frac{q + \eta^K - 1}{q\eta^K} \right) \leq 1 - \eta \\ &\Rightarrow q + \eta^K - 1 \leq \eta^K - \eta^{K+1} \\ &\Rightarrow 1 - q \geq \eta^{K+1}. \end{aligned}$$

Hence,

(3) $1 - q < \eta^{K+1} \Rightarrow \underline{k} = 1.$

Second, note $\bar{k} = n - 1$ implies, by Lemma 2,

$$\begin{aligned} p_{\bar{k}} = 1 &\Rightarrow r_{\bar{k}} p_{\bar{k}+1} \geq 1 - \eta \\ &\Rightarrow \frac{q + n^K - 1}{\eta^K} \geq 1 - \eta \\ &\Rightarrow 1 - q \leq \eta^{K+1}. \end{aligned}$$

Hence,

(4) $1 - q > \eta^{K+1} \Rightarrow \bar{k} = n.$

Let

$$I = \left\lfloor \frac{n}{2} \right\rfloor.$$

There are two cases.

Case I: $1 - q < \eta^I$. From Lemma 3, we have

$$K < \frac{n}{2} \Rightarrow K \leq \left\lfloor \frac{n}{2} \right\rfloor - 1 \Rightarrow I \geq K + 1.$$

Thus we have $1 - q < \eta^I \leq \eta^{K+1}$. But from (3), this implies $\underline{k} = 1$. Now since $\bar{k} \geq n - 1$, it follows that $K = I - 1$, and $\bar{k} = \underline{k} + 2K = 2I - 1$.

Case II: $1 - q > \eta^I$. Now $p_k > 0 \Rightarrow (q + \eta^k - 1)/q\eta^k > 0 \Rightarrow 1 - q < \eta^k$. Suppose $1 - q < \eta^{K+1}$. Then, from (1), we have $\underline{k} = 1$, and by the same argument as Case I, $K = I - 1 \Rightarrow 1 - q < \eta^{K+1} = \eta^I$, a contradiction. Hence we must have

$$\eta^{K+1} < 1 - q < \eta^K.$$

So K is the largest integer with $1 - q < \eta^K$. But now, from (4), it follows that $\bar{k} = n$. Q.E.D.

In the centipede games described in the text, the piles grow at an exponential rate: There are real numbers $c > d > 1$ with $a_i = cb_i$ and $a_{i+1} = da_i$ for all i . So $\eta = (c - d)/(cd^2 - d)$. In our experiments $c = 4$, and $d = 2$, so $\eta = \frac{1}{7}$. The figures in the text correspond to the solution for the two and three round games ($n = 4$ and $n = 6$) for these parameters.

It is interesting to note that since the solution depends only on η , the above solution also applies if there are linearly increasing payoffs of the form $a_{i+1} = a_i + c$, and $b_{i+1} = b_i + c$ (with $c > 0$), as long as $a_i > b_{i+1} = b_i + c$. Hence picking a_i , b_i , and c so that

$$\frac{a_1 - b_2}{a_3 - b_2} = \frac{a_1 - b_1 - c}{a_i - b_i + c} = \frac{1}{7},$$

(e.g., $a_1 = 60$, $b_1 = 20$, $c = 30$) one can obtain a game with linearly increasing payoffs whose solution is exactly the same as the solution of the game with exponentially increasing payoffs treated in this paper.

APPENDIX B

Experiment Instructions

This is an experiment in group decision making, and you will be paid for your participation in cash, at the end of the experiment. Different subjects may earn different amounts. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance.

The entire experiment will take place through computer terminals, and all interaction between you will take place through the computers. It is important that you not talk or in any way try to communicate with other subjects during the experiments. If you disobey the rules, we will have to ask you to leave the experiment.

We will start with a brief instruction period. During the instruction period, you will be given a complete description of the experiment and will be shown how to use the computers. You must take a quiz after the instruction period. So it is important that you listen carefully. If you have any questions during the instruction period, raise your hand and your question will be answered so everyone can hear. If any difficulties arise after the experiment has begun, raise your hand, and an experimenter will come and assist you.

The subjects will be divided into two groups, containing 10 subjects each. The groups will be labeled the RED group and the BLUE group. To determine which color you are, will you each please select an envelope as the experimenter passes by you.

[EXPERIMENTER PASS OUT ENVELOPES]

If you chose BLUE, you will be BLUE for the entire experiment. If you chose RED, you will be RED for the entire experiment. Please remember your color, because the instructions are slightly different for the BLUE and the RED subjects.

In this experiment, you will be playing the following game, for real money.

First, you are matched with an opponent of the opposite color. There are two piles of money: a Large Pile and a Small Pile. At the beginning of the game the Large Pile has 40 cents and the Small Pile has 10 cents.

RED has the first move and can either "Pass" or "Take." If RED chooses "Take," RED gets the Large Pile of 40 cents, BLUE gets the small pile of 10 cents, and the game is over. If RED chooses "Pass," both piles double and it is BLUE's turn.

The Large Pile now contains 80 cents and the Small Pile 20 cents. BLUE can take or pass. If BLUE takes, BLUE ends up with the Large Pile of 80 cents and RED ends up with the Small Pile of 20 cents and the game is over. If BLUE passes, both piles double and it is RED's turn again.

This continues for a total of six turns, or three turns for each player. On each move, if a player takes, he or she gets the Large Pile, his or her opponent gets the Small Pile, and the game is over. If the player passes, both piles double again and it is the other player's turn.

The last move of the game is move six, and is BLUE's move (if the game even gets this far). The Large Pile now contains \$12.80 and the Small Pile contains \$3.20. If BLUE takes, BLUE gets the Large Pile of \$12.80 and RED gets the Small Pile of \$3.20. If BLUE passes, then the piles double again. RED then gets the Large Pile, containing \$25.60, and BLUE gets the Small Pile, containing \$6.40. This is summarized in the following table.

PAYOFF CHART FOR DECISION EXPERIMENT

		Move #				Large Pile	Small Pile	RED's Payoff	BLUE's Payoff
1	2	3	4	5	6				
T						.40	.10	.40	.10
P	T					.80	.20	.20	.80
P	P	T				1.60	.40	1.60	.40
P	P	P	T			3.20	.80	.80	3.20
P	P	P	P	T		6.40	1.60	6.40	1.60
P	P	P	P	P	T	12.80	3.20	3.20	12.80
P	P	P	P	P	P	25.60	6.40	25.60	6.40

[EXPERIMENTER HAND OUT PAYOFF TABLE]

Go over table to explain what is in each column and row.

The experiment consists of 10 games. In each game, you are matched with a different player of the opposite color from yours. Thus, if you are a BLUE player, in each game, you will be matched with a RED player. If you are a RED player, in each game you are matched with a BLUE player.

Since there are ten subjects of each color, this means that you will be matched with each of the subjects of the other color *exactly* once. So if your label is RED, you will be matched with each of the BLUE subjects *exactly* once. If you are BLUE, you will be matched with each of the RED subjects *exactly* once.

We will now begin the computer instruction session. Will all the BLUE subjects please move to the terminals on the left side of the room, and all the RED subjects move to the terminals on the right side of the room.

[SUBJECTS MOVE TO CORRECT TERMINALS]

During the instruction session, we will teach you how to use the computer by going through a few practice games. During the instruction session, *do not hit any keys until you are told to do so*, and when you are told to enter information, *type exactly what you are told to type*. You are not paid for these practice games.

Please turn on your computer now by pushing the button labeled "MASTER" on the right hand side of the panel underneath the screen.

[WAIT FOR SUBJECTS TO TURN ON COMPUTERS]

When the computer prompts you for your name, type your full name. Then hit the ENTER key.

[WAIT FOR SUBJECTS TO ENTER NAMES]

When you are asked to enter your color, type R if your color is RED, and B if your color is BLUE. Then hit ENTER.

[WAIT FOR SUBJECTS TO ENTER COLORS]

You now see the experiment screen. Throughout the experiment, the bottom of the screen will tell you what is currently happening, and the top will tell you the history of what happened in the previous games. Since the experiment has not begun yet, the top part of the screen is currently empty. The bottom part of the screen tells you your subject number and your color. It also tells you the subject number of the player you are matched against in the first game. Is there anyone whose color is not correct?

[WAIT FOR RESPONSE]

Please record your color and subject number on the top left hand corner of your record sheet. Also record the number of the subject you are matched against in the first game.

Each game is represented by a row in the upper screen, and the player you will be matched with in each of the ten games appears in the column labeled "OPP" (which stands for "opponent") on the right side of the screen. It is important to note that you will never be paired with the same player twice.

We will now start the first practice game. Remember, do not hit any keys until you are told to do so.

[MASTER HIT KEY TO START FIRST GAME]

You now see on the bottom part of the screen that the first game has begun, and you are told who you are matched against. If you are a RED player, you are told that it is your move, and are given a description of the choices available to you. If you are a BLUE player, you are told that it is your opponent's move, and are told the choices available to your opponent.

Will all the RED players now choose PASS by typing in P on your terminals now.

[WAIT FOR SUBJECTS TO CHOOSE]

Since RED chose P, this is recorded on the top part of the screen with a P in the first RED column, and the cursor has moved on to the second column, which is BLUE, indicating that it is BLUE's move.

On the bottom part of the screen, the BLUE players are now told that it is their turn to choose, and are told the choices they can make. The RED players are told that it is their opponent's turn to choose, and are told the choices that their opponent can make. Notice, that there is now a Large Pile of \$.80 and a Small Pile of \$.20.

Will all the BLUE players now please choose TAKE by typing T at your terminal now.

[WAIT FOR SUBJECTS TO CHOOSE]

Since BLUE chose T, the first game has ended. On the bottom part of the screen, you are told that the game is over, and that the next game will begin shortly. On the top part of the screen, BLUE's move is recorded with a T in the second column. The payoffs from the first game for both yourself and your opponent are recorded on the right hand side of the screen in the columns labeled "Payoff." Your own payoff is in your color. That of your opponent is in the opponent's color.

Please record your own payoff on the record sheet that is provided.

[WAIT FOR SUBJECTS TO RECORD PAYOFFS]

You are not being paid for the practice session, but if this were the real experiment, then the payoff you have recorded would be money you have earned from the first game, and you would be paid this amount for that game at the end of the experiment. The total you earn over all ten real games is what you will be paid for your participation in the experiment.

We will now proceed to the second practice game.

[MASTER HIT KEY TO START SECOND GAME]

You now see that you have been matched with a new player of the opposite color, and that the second game has begun. Does everyone see this?

[WAIT FOR RESPONSE]

The rules for the second game are exactly like the first. The RED player gets the first move.

[DO RED-P, BLUE-P, RED-P]

Now notice that it is BLUE's move. It is the last move of the game. The Large Pile now contains \$3.20, and the Small Pile contains \$.80. If the BLUE player chooses TAKE, then the game ends. The BLUE player receives the Large Pile and the RED player receives the Small Pile. If the BLUE player chooses PASS, both piles double, and *then* the game ends. The RED player receives the Large Pile, which now contains \$6.40, and the BLUE player receives the Small Pile, containing \$1.60.

Will the BLUE player please choose PASS by typing P at your terminal now.

[WAIT FOR SUBJECTS TO CHOOSE]

The second practice game is now over. Please record your payoff on the second line of your record sheet.

[WAIT FOR PLAYERS TO RECORD PAYOFFS]

[MASTER HIT KEY TO START THIRD GAME]

We now go to the third practice game. Notice again that you have a new opponent. Will all the RED players please choose TAKE by typing T at your terminal now.

[WAIT FOR PLAYERS TO CHOOSE]

Since the RED player chose TAKE on the first move, the game is over, and we proceed on to the next game. Since RED chose TAKE on the first move, BLUE did not get any chance to move.

Please record your payoff for the third game on the third line of your record sheet.

[WAIT FOR PLAYERS TO RECORD PAYOFFS]

This concludes the practice session. In the actual experiment there will be ten games instead of three, and, of course, it will be up to you to make your own decisions. At the end of game ten, the experiment ends and we will pay each of you privately, in cash, the TOTAL amount you have accumulated during all ten games, plus your guaranteed five dollar participation fee. No other person will be told how much cash you earned in the experiment. You need not tell any other participants how much you earned.

Are there any questions before we pass out the quiz?

[EXPERIMENTER TAKE QUESTIONS]

O.K., then we will now have you take the quiz.

[PASS OUT QUIZ]

[COLLECT AND MARK QUIZZES]

[HAND QUIZZES BACK AND GO THRU CORRECT ANSWERS]

We will now begin with the actual experiment. If there are any problems from this point on, raise your hand and an experimenter will come and assist you. When the computer asks for your name, please start as before by typing in your name. Wait for the computer to ask for your color, then respond with the correct color.

[START EXPERIMENT]
 [CHECK THAT COLORS ARE OK BEFORE BEGINNING EXPERIMENT]

APPENDIX C

Experimental Data

The following tables give the data for our experiment. Each row represents a subject. The columns are

- Col 1: Session number
- Col 2: Subject number of Red Player

Col 2 + *j*: Outcome of game *j*. Letting *k* be the entry in this column, and *n* be the number of moves in the game (*n* = 4 for Exp. 1-4, *n* = 6 for Exp. 5-7), then

$$k = \begin{cases} \leq n \Rightarrow & \text{game ended with } T \text{ on move } k, \\ = n + 1 \Rightarrow & \text{game ended with } P \text{ on move } n. \end{cases}$$

The matching scheme was: In game *j*, Red subject *i* is matched with Blue subject [(*i* + *j* - 1) mod *m*], where *m* is the number of subjects of each color in the session. Thus, with ten subjects of each color, in the first game, Red *i* is matched with Blue *i*. In the second game, Red *i* is matched with Blue 1 + *i*, except Red 10, who is matched with Blue 1.

1	1	3	3	3	3	2	3	2	2	2	3
1	2	4	2	4	4	4	4	2	4	4	2
1	3	3	3	2	2	3	3	3	4	3	2
1	4	4	3	3	4	3	2	3	3	3	2
1	5	2	1	3	1	3	3	3	3	2	2
1	6	5	3	4	3	5	4	4	3	3	3
1	7	4	2	2	3	4	3	2	3	2	3
1	8	2	1	5	4	1	3	3	4	3	3
1	9	3	3	4	3	2	2	3	1	3	2
1	10	4	2	4	2	3	1	3	3	2	5

SESSION 1
 (Four move, PCC)

2	1	4	4	3	3	2	4	2	2	2
2	2	2	1	4	1	3	2	2	2	1
2	3	3	3	2	2	1	1	1	1	1
2	4	3	3	3	3	3	2	3	2	3
2	5	3	4	2	3	2	3	3	2	3
2	6	3	3	3	2	3	3	2	3	2
2	7	2	4	3	4	5	2	3	2	2
2	8	3	2	3	3	2	3	2	3	2
2	9	2	4	2	3	4	2	3	2	2

SESSION 2
 (Four move, PCC)

3	1	3	4	2	3	2	3	2	2	1	3
3	2	4	2	2	2	4	2	2	5	4	3
3	3	2	2	1	1	2	1	1	3	3	2
3	4	3	2	3	2	2	3	3	3	3	2
3	5	3	5	2	2	3	4	3	3	2	3
3	6	5	2	2	5	4	4	4	2	2	2
3	7	2	3	5	4	2	3	2	2	2	1
3	8	2	3	3	3	3	2	2	2	5	2
3	9	5	4	5	5	2	2	2	4	2	4
3	10	4	4	3	2	2	3	3	2	2	3

SESSION 3
(Four move, CIT)

4	1	2	4	5	4	2	4	4	2	4	4
4	2	3	3	2	2	2	1	1	1	3	2
4	3	4	2	3	2	3	2	3	3	2	3
4	4	2	3	2	3	2	3	3	2	2	2
4	5	3	3	3	2	5	4	2	2	2	1
4	6	3	3	2	5	4	2	2	2	2	2
4	7	3	3	5	4	2	1	3	2	1	1
4	8	3	5	4	2	3	3	2	2	2	3
4	9	3	3	3	3	3	2	2	3	3	2
4	10	3	2	1	1	1	1	1	1	1	1

SESSION 4
(Four move, High payoff, CIT)

5	1	5	5	7	4	4	2	4	2	4	3
5	2	3	3	3	3	3	4	2	3	2	2
5	3	5	4	4	3	4	3	3	3	3	5
5	4	3	3	3	3	3	3	3	3	3	3
5	5	5	4	4	3	5	4	4	5	5	1
5	6	4	4	4	3	4	3	3	3	3	3
5	7	3	3	3	3	3	3	5	4	4	2
5	8	5	5	4	4	3	3	3	3	2	1
5	9	5	5	5	5	5	4	4	2	4	4
5	10	5	4	5	5	4	5	2	4	4	6

SESSION 5
(Six move, CIT)

6	1	5	6	4	5	5	5	6	5	4
6	2	6	5	7	4	6	4	4	4	4
6	3	4	5	5	4	5	4	4	4	5
6	4	5	5	5	4	5	4	4	5	5
6	5	4	5	6	6	4	4	5	4	5
6	6	4	2	5	4	4	4	4	7	6
6	7	4	6	4	4	3	4	3	5	4
6	8	2	5	4	5	4	5	4	6	5
6	9	5	4	5	4	5	4	4	3	4

SESSION 6
(Six move, PCC)

7	1	4	4	4	3	3	4	3	3	3	3
7	2	2	5	4	2	5	4	3	5	3	4
7	3	5	3	2	5	5	4	5	5	4	4
7	4	4	2	5	4	6	6	5	4	4	4
7	5	2	5	5	4	5	6	4	4	4	4
7	6	7	6	6	6	5	4	4	4	4	6
7	7	6	4	4	6	4	4	4	4	6	6
7	8	4	5	5	4	4	5	5	4	5	4
7	9	4	3	5	5	5	4	4	5	4	4
7	10	6	4	2	4	3	2	3	3	4	3

SESSION 7
(Six move, PCC)

REFERENCES

- AUMANN, R. (1988): "Preliminary Notes on Integrating Irrationality into Game Theory," Mimeo, International Conference on Economic Theories of Politics, Haifa.
- (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, 1–18.
- BINMORE, K. (1987): "Modeling Rational Players," *Economics and Philosophy*, 3, 179–214.
- CAMERER, C., AND K. WEIGELT (1988): "Experimental Tests of a Sequential Equilibrium Reputation Model," *Econometrica*, 56, 1–36.
- COOPER, R., D. DEJONG, R. FORSYTHE, AND T. ROSS (1990): "Selection Criteria in Coordination Games," *American Economic Review*, 80, 218–233.
- FUDENBERG, D., AND E. MASKIN (1986): "The Folk Theorem in Repeated Games with Discounting and Incomplete Information," *Econometrica*, 54, 533–554.
- GUTH, W., P. OCKENFELS, AND M. WENDEL (1991): "Efficiency by Trust in Fairness?—Multiperiod Ultimatum Bargaining Experiments with an Increasing Cake," Technical Report, J. W. Goethe Universitat, Frankfurt.
- HARSANYI, J. C. (1967–68): "Games of Incomplete Information Played by Bayesian Players," Parts (I, II, and III). *Management Science*, 14, 159–182, 320–334, 486–502.
- JUNG, Y., J. KAGEL, AND D. LEVIN (1989): "On the Existence of Predatory Pricing in the Laboratory: An Experimental Study of Reputation and Entry Deterrence in the Chain-Store Game," Mimeo.
- KREPS, D. M. (1990): *A Course in Microeconomic Theory*. New Jersey: Princeton University Press.
- KREPS, D., P. MILGROM, J. ROBERTS, AND R. WILSON (1982): "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma," *Journal of Economic Theory*, 27, 245–252.
- KREPS, D. M., AND R. WILSON (1982a): "Reputation and Imperfect Information," *Journal of Economic Theory*, 27, 253–279.
- (1982b): "Sequential Equilibria," *Econometrica*, 50, 863–894.
- MCKELVEY, R. D. (1990): "GAMBIT: Interactive Extensive Form Game Program," Mimeo, California Institute of Technology.
- MEGIDDO, N. (1986): "Remarks on Bounded Rationality," Technical Report, IBM Research Report RJ 54310, Computer Science.
- NERAL, JOHN, AND JACK OCHS (1989): "The Sequential Equilibrium Theory of Reputation Building: A Further Test," mimeo, University of Pittsburgh.
- PALFREY, T. R., AND H. ROSENTHAL (1988): "Private Incentives in Social Dilemmas: The Effects of Incomplete Information and Altruism," *Journal of Public Economics*, 35, 309–332.
- RENY, P. (1988): "Rationality, Common Knowledge, and the Theory of Games," Technical Report, University of Western Ontario.
- ROSENTHAL, R. (1982): "Games of Perfect Information, Predatory Pricing, and the Chain Store Paradox," *Journal of Economic Theory*, 25, 92–100.
- SELTEN, R., AND R. STOECKER (1986): "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames," *Journal of Economic Behavior and Organization*, 7, 47–70.
- SMITH, V. (1976): "Experimental Economics: Induced Value Theory," *American Economic Review*, 66, 274–279.