DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

An Experimental Study of Constant-sum Centipede Games

Mark Fey

Richard D. McKelvey

Thomas R. Palfrey

SOCIAL SCIENCE WORKING PAPER 877

May 1994

An Experimental Study of Constant-sum Centipede Games

Mark Fey Richard D. McKelvey Thomas R. Palfrey

Abstract

In this paper, we report the results of a series of experiments on a version of the centipede game in which the total payoff to the two players is constant. Standard backward-induction arguments lead to a unique Nash equilibrium outcome prediction, which is the same as the prediction made by theories of "fair" or "focal" outcomes.

We find that subjects frequently fail to select the unique Nash outcome prediction. While this behavior was also observed in McKelvey and Palfrey (1992) in the "growing pie" version of the game they studied, the Nash outcome was not "fair", and there was the possibility of Pareto improvement by deviating from Nash play. Their findings could therefore be explained by small amounts of altruistic behavior. There are no Pareto improvements available in the constant-sum games we examine, hence explanations based on altruism cannot account for these new data.

We examine and compare two classes of models to explain this data. The first class consists of non-equilibrium modifications of the standard "Always Take" model. The other class we investigate, the Quantal Response Equilibrium model, describes an equilibrium in which subjects make mistakes in implementing their best replies and assume other players do so as well. One specification of this model fits the experimental data best, among the models we test, and is able to account for all the main features we observe in the data.

Keywords: Game Theory, Experiments, Bounded Rationality

JEL Classification numbers: 026, 212

An Experimental Study of Constant-sum Centipede Games

Mark Fey Richard D. McKelvey Thomas R. Palfrey

1 Introduction

In this paper, we report on a series of experiments which produce results not easily explained by any of a wide class of game-theoretic or decision theoretic models. The type of experiments we consider are closely related to the "centipede games" studied by McKelvey and Palfrey (1992). Experimental subjects participated in a special form of centipede game in which the total payoff to the two players was constant. We refer to these games as "constant-sum centipedes games." It is easy to see that standard backward-induction arguments lead to a unique Nash equilibrium outcome prediction. The Nash outcome is also the only "fair," "focal," or maximin outcome in the game. However, subjects in our experiments frequently fail to adopt this outcome.

While similar behavior was also observed by McKelvey and Palfrey, the "regular" centipede games they conducted allowed Pareto improvements over the course of the game. Hence, they were able to explain the behavior based on small amounts of altruism. There are no Pareto improvements in the constant-sum centipedes we examine. So we are unable to explain our findings using similar ideas.

The class of games we study was first introduced by Rosenthal (1981). Binmore (1987) considered a version of Rosenthal's game with 100 moves which he dubbed a "centipede" game. Kreps (1990) contains a nice examination of this game, and Megiddo (1986) and Aumann (1988) examine a version with exponentially increasing payoffs, rather than linearly increasing payoffs. This latter game was chosen by McKelvey and Palfrey (1992) for their experimental study. Other, related experiments include ultimatum games from the bargaining literature, studied by Güth et al. (1993) and Thaler (1988). Both McKelvey and Palfrey (1992) and Güth et al. (1993) study experiments in which the players' joint

*This research was supported in part by National Science Foundation grant #SES-9223701 to the California Institute of Technology. We thank Bob Forsythe and Ray Riezman for help in facilitating the use of the experimental laboratory at University of Iowa. We are grateful for comments and suggestions from participants at the 1993 ESA Fall Meetings, an editor, and an anonymous referee.

payoff increases within the game. In the constant-sum version of the centipede game, in contrast to the previously cited studies, there are no efficiency gains available; the game-theoretic, efficiency, and fairness predictions agree.

The centipede game is one of a class of games that call into question the common knowledge of rationality among players that standard game theory requires. Specifically, backward reduction requires that players decide what their "rational" opponent will do after a series of "irrational" moves. Indeed, recent work by Basu (1990) and Reny (1993) argues that in two-person games with perfect information, no fully rational solution concept escapes the paradoxical requirements of players' "rationality." These theoretical results are compelling and lead us to believe that any empirically plausible model of game play in experiments must allow for an element of irrationality in players' actions. We choose to express this by a model that includes "mistakes" by players.

In the next section we describe the design and procedures of the experiments. In the third section we summarize our data and identify its main features. The fourth section examines individual behavior in order to uncover evidence of altruistic behavior in our experimental subjects. Finding none, the following two sections examine and compare several different classes of models to explain the data on our constant-sum centipede games. Some explanations we reject out of hand, such as the Rational, Egalitarian, and Maximin models. We also describe a class of modified "Always Take" models, which we term the Random and Learning models, that admit rigorous statistical testing. As an alternative to these non-equilibrium models, we specify a model, called the Quantal Response Equilibrium model, in which subjects make mistakes in implementing their best replies and assume other players do so as well. Among the models we investigate, this Quantal Response Equilibrium model has the best fit and is able to account for all the main features we observe in the experimental data. The final section presents our conclusions.

2 Experimental Design

The constant-sum centipede game we study is a two player game that can be described as follows. The game involves a fixed amount of money ($3.20) which is initially divided into two equal-size piles ($1.60 each). Player one has the first move, and can choose to take one of the piles or to pass. If the first player takes a pile, the other (equal-size) pile is given to the second player and the game ends. If the first player passes, one fourth of one pile is moved to the other pile, and it is the second player's move. The second player now has the option of taking the big pile and thus leaving the small pile for the first player, or choosing to pass. If the second player passes, one fourth of the small pile is moved to the big pile, and the move returns to the first player. This continues for a predetermined number of moves by each player. Every time a player passes, one fourth of the small pile is moved to the large pile. The game ends as soon as either of the players chooses to take the big pile.

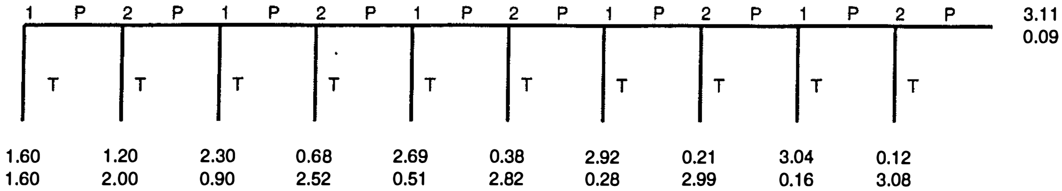

Figure 1: A Ten-Move Constant-Sum Centipede Game.

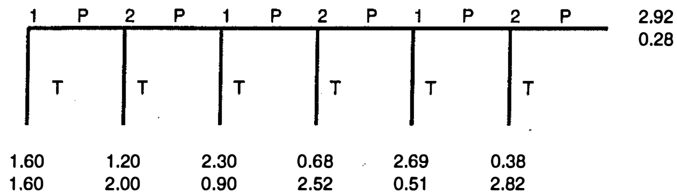

Figure 2: A Six-Move Constant-Sum Centipede Game.

We examined two different game lengths. The first game consists of three "innings", for a total of six moves in the whole game. The second game consists of five innings, for a total of ten moves. The extensive forms of the ten and six move games are shown in Figures 1 and 2, respectively.

We conducted a total of nine experiments – three experiments on each of three different subject pools. The three subject groups used were students from Caltech, Pasadena City College, and the University of Iowa. Experiments involving the first two subject groups were conducted at the Caltech Laboratory for Experimental Economics and Political Science. The third group of experiments were conducted at the University of Iowa. The three experiments in each group consisted of two ten-move centipede experiment and one six-move experiment. The design of all the experiments is summarized in Table 1.[1]

The experimental setup was very similar to that used by McKelvey and Palfrey (1992). Each experiment used either eighteen or twenty subjects, none of whom had previously participated in any form of centipede experiment. The subjects were divided into two equal size groups before the experiment began. We denoted these the Red and Blue groups. The subjects then participated in a series of either 9 or 10 matches. In each match, a Red subject was matched with a Blue subject, and they participated in a constant-sum centipede game. In each game, the Red player moved first, and the Blue player moved second. Each subject played against every player in the other group, using the matching scheme described in McKelvey and Palfrey (1992). This matching scheme was designed to eliminate any possible supergame or cooperative behavior.

In all of the experiments, subjects interacted through computer terminals and were not allowed to communicate in any other way. The subjects in a particular experiment played the same game form (pictured in Figures 1 or 2) throughout the entire experiment.

[1]We conducted two additional experiments that are not included in this analysis. They both used a different payoff structure than the games presented in this paper.

Exp. #	Subject Pool	Game Length	# Subjects	Matches/ Subject	Total # Matches
1	CIT	10	20	10	100
2	CIT	10	20	10	100
3	Iowa	10	20	10	100
4	Iowa	10	20	10	100
5	PCC	10	18	9	81
6	PCC	10	20	10	100
7	CIT	6	20	10	100
8	Iowa	6	20	10	100
9	PCC	6	18	9	81

Table 1: Experimental Design

The subjects were allowed to participate only after listening to the instructions contained in Appendix 1 and completing a quiz designed to test their knowledge of the rules of the game.

3 Data Overview

The complete data for the nine experiments are presented in Appendix 2. The results are summarized in Table 2, which lists the number and proportion of *Takes* played at each node of the matches in each experiment.[2]

Two features of the data are immediately apparent. First, subjects frequently do not play the unique Nash equilibrium prediction of taking at the first move (as little as 22% of the time in one of the experiments). Second, there is some variation in outcomes across experiments that seems to be linked to the differing subject pools. Second, subjects frequently do not play the unique Nash equilibrium prediction of taking at the first move (as little as 22% of the time in one of the experiments). From Table 2 we see that overall, averaging across all experiments and all matches, slightly less than half of the observations correspond to the Nash equilibrium. Even in experiment #8, the most favorable from the standpoint of the game theoretic prediction, in nearly a quarter of the matches there was at least one pass. Considering only the ten-move centipede games, we see that in no session did less than 40% of the red subjects choose *Pass* at the first node.[3]

A different description of the strategic choices being made by the subjects is given by

[2]We number nodes in the game tree from 0 to 9 for the ten-move sessions, and 0 to 5 for the six-move experiments.

[3]The results are even more striking at the individual level, where we find that over 80% of the red subjects passed in at least one of the matches they played. This is discussed in a later section.

Exp. #	Number of Passes* 0	1	2	3	4	5	6	7	8	9	10
1	.57	.31	.11	.00	.01	0	0	0	0	0	0
CIT-10	(57)	(31)	(11)	(0)	(1)	0	0	0	0	0	0
2	.51	.36	.12	.00	.01	0	0	0	0	0	0
CIT-10	(51)	(36)	(12)	(0)	(1)	0	0	0	0	0	0
3	.60	.28	.12	0	0	0	0	0	0	0	0
UI-10	(60)	(28)	(12)	0	0	0	0	0	0	0	0
4	.38	.36	.18	.06	.01	.00	.01	0	0	0	0
UI-10	(38)	(36)	(18)	(6)	(1)	(0)	(1)	0	0	0	0
5	.42	.40	.09	.05	.02	.02	0	0	0	0	0
PCC-10	(34)	(32)	(7)	(4)	(2)	(2)	0	0	0	0	0
6	.22	.23	.26	.13	.08	.06	.01	.00	.00	.01	0
PCC-10	(22)	(23)	(26)	(13)	(8)	(6)	(1)	(0)	(0)	(1)	0
Pooled	.45	.32	.14	.04	.02	.01	.003	.00	.00	.001	0
10 move	(262)	(186)	(86)	(23)	(13)	(8)	(2)	(0)	(0)	(1)	0
7	.62	.31	.07	0	0	0	0				
CIT-6	(62)	(31)	(7)	0	0	0	0				
8	.77	.23	0	0	0	0	0				
UI-6	(77)	(23)	0	0	0	0	0				
9	.33	.48	.15	.02	.01	0	0				
PCC-6	(27)	(39)	(12)	(2)	(1)	0	0				
Pooled	.59	.33	.07	.007	.003	0	0				
6 move	(166)	(93)	(19)	(2)	(1)	0	0				

Table 2: Proportions of matches ending at each outcome

*The number in parentheses is the number of observations at that node in the game tree.

the implied take probabilities at various nodes within a match. These are the aggregate conditional frequencies of players choosing *Take*, given that a particular node has been reached. Thus, the implied take probability q_i is our estimate of the likelihood that a player who faces a choice at node i will choose *Take* at that node. For each experiment, we have shown the implied probability of taking at each node with a significant number of observations.[4] Figure 3 graphs these implied take probabilities for the six ten-move experiments, and Figure 4 presents the results from the three six-move experiments.

It is clear from the bar graphs that with the possible exception of experiment #5, the implied take probabilities are increasing along the game tree. In other words, on average, within a particular match, subjects are more likely to choose *Take* after several passes than after fewer passes. Apparently, subjects use strategies that have systematic differences at different nodes of the game. This is a feature observed by McKelvey and Palfrey (1992), and it should be accounted for by any model of the subjects' behavior.

Insert figures 3 and 4 about here.

In addition to strategic variation *within* matches, we are interested in behavioral differences *across* matches. Specifically, we seek to determine whether the subjects' performance changes with experience. Table 3 presents the implied take probabilities for the first and last five matches in each experiment. In almost every case, no matter how far out in the game, subjects are more likely to choose *Take* in the later matches than in the earlier ones. We also examined the cumulative frequencies of outcomes, derived by summing the outcome proportions in the first and last five matches of each experiment. We observed that, with the exception of experiment #5, the distribution of outcomes in later matches are stochastically dominated by the earlier matches.[5] Thus, subjects choose *Take* at nodes closer to the start of the game as they play more. So this is another feature that any potential model must explain - subjects play "closer" to the Nash prediction as they gain more experience in the game. This was included as a learning parameter in McKelvey and Palfrey (1992).

As we have shown, there are two main features of the data we would like to explain. First, *within* matches, players are increasingly more likely to choose *Take* as moves are passed back and forth. Second, *across* matches, players are increasingly more likely to choose *Take* as they gain experience with the game.

[4]We have eliminated nodes with fewer than ten observations.

[5]It should be noted that the differences in the cumulative frequencies for experiment #5 are very small and do *not* suggest that subjects move farther from the Nash equilibrium over time.

Exp. #	Match Group	Conditional Probabilities* p_0	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	p_9
1	<6	.46 (50)	.70 (27)	.88 (8)	.00 (1)	1.00 (1)					
CIT-10	>5	.68 (50)	.75 (16)	1.00 (4)							
2	<6	.34 (50)	.64 (33)	.92 (12)	.00 (1)	1.00 (1)					
CIT-10	>5	.68 (50)	.94 (16)	1.00 (1)							
3	<6	.48 (50)	.65 (26)	1.00 (9)							
UI-10	>5	.72 (50)	.79 (14)	1.00 (3)							
4	<6	.32 (50)	.38 (34)	.67 (21)	.71 (7)	.50 (2)	.00 (1)	1.00 (1)			
UI-10	>5	.44 (50)	.82 (28)	.80 (5)	1.00 (1)						
5	<6	.40 (45)	.67 (27)	.56 (9)	.50 (4)	.50 (2)	1.00 (1)				
PCC-10	>5	.44 (36)	.70 (20)	.33 (6)	.50 (4)	.50 (2)	1.00 (1)				
6	<6	.22 (50)	.15 (39)	.36 (33)	.42 (21)	.50 (12)	.67 (6)	.50 (2)	.00 (1)	.00 (1)	1.00 (1)
PCC-10	>5	.22 (50)	.44 (39)	.64 (22)	.50 (8)	.50 (4)	1.00 (2)				
7	<6	.48 (50)	.73 (26)	1.00 (7)							
CIT-6	>5	.76 (50)	1.00 (12)								
8	<6	.70 (50)	1.00 (15)								
UI-6	>5	.84 (50)	1.00 (8)								
9	<6	.20 (45)	.67 (36)	.75 (12)	.67 (3)	1.00 (1)					
PCC-6	>5	.50 (36)	.83 (18)	1.00 (3)							

Table 3: Implied Take Probabilities Across Matches

*The number in parentheses is the number of observations at that node in the game tree.

4 Altruism and Individual Behavior

This section examines the individual behavior of the experimental subjects in order to examine the Altruism model of McKelvey and Palfrey (1992) as applied to the data reported here on the constant-sum centipede game. Many of the qualitative features we have just described for the constant-sum centipede game are shared by the data reported by McKelvey and Palfrey. It is natural to investigate whether similar conclusions can be drawn about explanations for the data. In fact, the constant-sum games were designed to test the earlier explanation of passing behavior in centipede games.

A general approach to understanding the actions of experimental subjects postulates that the subject pools we draw from are composed of several different types of players, each with a different utility function. McKelvey and Palfrey (1992) used a version of this idea that they termed the Altruism model.[6] Specifically, the authors supposed that individuals in their experiments were either selfish or altruistic. Altruists attempted to maximize the *social* gains available in the experiment. Selfish subjects simply maximized their own expected utility, given their beliefs about the mix of subject types in the experiment. Thus, in equilibrium selfish subjects mimic altruists by passing with some probability in the early stages of a match. A key finding of McKelvey and Palfrey was that the presence of a small percentage of altruists (approximately 5%) was able to account for most of the systematic patterns in their data. Given the joint benefits from choosing *Pass* in the centipede games they studied, McKelvey and Palfrey defined an altruist as "an individual who always chooses *Pass*" (p. 812). They observed a total of 9 subjects out of 138 who met this definition of altruism. These 9 subjects consisted of 5 Red subjects and 4 Blue subjects. More striking, they observed that "only 1 out of all 138 subjects chose *Take* at every opportunity" (p. 811).

In our experiments, there are no joint benefits from choosing *Pass*, thus altruism carries no prediction for these experiments. However, one possible explanation for the data is that there are subjects who, for some reason (perhaps "reciprocation") play the same role in our experiments as the altruists in McKelvey and Palfrey (1992), and always choose *Pass*. Thus, we investigate whether such subjects exist in our constant-sum games. Our findings for the constant-sum centipede are given in Table 4. It is clear that subjects are behaving much differently in the constant-sum game. Out of 176 subjects, 45 chose *Take* at every opportunity[7] and only 2 chose *Pass* at every opportunity. Both of these findings are in sharp contrast to the data reported by McKelvey and Palfrey (1992).

Table 4 also reveals some differences in the strategies used by the Red and Blue subjects. Because of the sequential nature of the game, a Blue player is allowed the chance to move only if the Red player matched against her chooses *Pass*. Thus, while Red subjects make at least one choice in every match, the Blue subjects can have significantly

[6]Similar approaches appear in Palfrey and Rosenthal (1988) and Cooper et al. (1989).

[7]In other words, over 25% of our subjects are perfectly consistent with the benchmark "Always Take" models. This suggests that a mixed model with multiple "types" might fit our data quite well. See El-Gamal and Grether (1993) and Stahl and Wilson (1993) for two different approaches to estimating mixture models with experimental data.

Exp. #	Subject Pool	All Takes Red	All Takes Blue	All Passes Red	All Passes Blue	Total Players
1	CIT-10	3	4	0	2	20
2	CIT-10	2	3	0	0	20
3	UI-10	3	3	0	0	20
4	UI-10	1	1	0	0	20
5	PCC-10	1	3	0	0	18
6	PCC-10	1	0	0	0	20
7	CIT-6	3	6	0	0	20
8	UI-6	1	9	0	0	20*
9	PCC-6	0	1	0	0	18
ALL		15	30	0	2	176

Table 4: Classification of Strategies Used by Experimental Subjects

*Includes one Blue player with no observed actions.

fewer opportunities. Indeed, in experiment #8, Blue player #10 did not make a single choice! The Red subjects matched against Blue #10 choose *Take* at the first node in every match. Therefore, some of the observed differences between Red and Blue subjects may be due to the lack of observations for the Blue subjects' behavior.

This phenomenon also calls into question our weak findings on subjects who always *Pass*. Significantly, there are no Red subjects who chose *Pass* at every opportunity; both are Blue subjects. But in one case, this classification rests on only three choices and in the other, it depends on only five choices. It seems reasonably likely that if these Blue subjects were given as many opportunities as the Red subjects, they would sometimes choose *Take*.[8]

We conclude that there is no evidence of altruism in our experimental data; a finding in sharp contrast to the centipede experiments of McKelvey and Palfrey (1992). This is not surprising, given the constant-sum nature of the games reported here. It does cast doubt on the explanation offered for the earlier data of McKelvey and Palfrey. It also forces us to consider some alternative explanations for the constant-sum data reported here. We do so in the next section.

[8]Despite the evidence that there do not seem to be players that always pass, we have estimated a version of the model of McKelvey and Palfrey (1992) on the six-move constant-sum experiments. As expected, we find a worse fit to our experimental data, on the whole, than offered by the models estimated in the previous section.

5 Models

In this section, we describe several different models which might be proposed to account for the data. In order to justify our conclusions, our main interest is in models that generate testable predictions. This testing, using standard econometric methods, will be discussed in the next section. It should be noticed that with several of the models outlined in this section some latitude exists regarding how a model might best be applied to the experiments we conducted. These complications will also be dealt with later.

The "Always Take" Models

The first model of behavior we describe is the usual Nash equilibrium prediction, which we term the Rational model. By a standard backward induction argument, the unique subgame-perfect Nash equilibrium[9] to the constant-sum centipede game is for both subjects to take with probability one at every opportunity. This model predicts that, following backward induction, every match will end with the Red player taking at the first node. Sharing this prediction is the Egalitarian model. In our experiments, the principle of equal division is both intuitively "fair" and has the enviable distinction of being a Pareto efficient "equitable allocation."[10] In addition to this "fairness" property, equal division is also an obvious "focal" point of the game since it is the only symmetric outcome in the game that is possible to achieve. As equal payoffs can only be guaranteed at the initial node of the game, the Egalitarian model also predicts that every match will end with Red taking at the first node. This prediction is also common to the Maximin model, in which players maximize their minimum payoff. This can be viewed as an extreme form of risk aversion. It is clear that players acting in this way will choose *Take* at their first opportunity, as passing at any node risks a lower payoff.

The three models we have described so far, which we refer to collectively as the "Always Take" models, share the common prediction that the Red player takes at the first node of the game. This is a point prediction, in that the models rule out any other outcome as impossible. Econometrically, this makes it impossible to specify the likelihood of our data, which contain prohibited outcomes. The point predictions of these models also rule out any kind of errors or experimentation by the players. This is a weakness of these simple models.

In order to give the "Always Take" model a fair shake, we consider next a statistical version of that model. A modified "Always Take" model, which we call the Random model, assumes that every player, at every node reached in every match, chooses *Take* with a fixed probability p and chooses *Pass* with probability $1-p$. Estimates of p close

[9]In fact, any Nash equilibrium to the constant-sum centipede involves the first player taking at the first node. Given this fact, all of the usual refinements of Nash equilibrium make the same prediction. Thus, "traditional" game theory makes a unambiguous prediction about play.

[10]An "equitable allocation" in welfare economics is an allocation such that no agent prefers the bundle of any other agent to his own. See Varian (1984).

to 1 would indicate that the "Always Take" model does fairly well in explaining our data. However, since our data appears to have trends in p within and across matches, we would expect this model to explain the data rather poorly.

A further modified "Always Take" model, the Learning model, is actually a combination of the random and rational models. It assumes that subjects gradually change their behavior as they gain experience, moving from the random model to the rational model. Thus, this model assumes that p varies over time, increasing with experience. However, it maintains the assumption that p is the same at all nodes *within* a match.[11] Recall that Table 3 showed that these implied take probabilities are not constant.

It is important to note that the Random and Learning models do *not* require that subjects act with little or no deterministic intent. The Random model can be interpreted as a particular specification of the error structure in the Rational or Egalitarian model. Thus, if subjects adopt a common (best) strategy that they can only imperfectly implement, they would fall under the Random model. In other words, subjects playing the game with good intentions but "trembling" hands could be described by the Random or Learning models.[12] However, this "noisy play" interpretation assumes a very simplistic model of errors which is not internally consistent. Specifically, these models do not allow players to take into account the errors that other players make. The next model overcomes this deficiency by offering an internally consistent description of play by error-prone agents.

The Quantal Response Equilibrium Model

Building off econometric models of discrete choice, the Quantal Response Equilibrium model parametrically describes the actions (and equilibrium play) of subjects who imperfectly implement their best replies.[13] In this model, sophisticated players play mutually consistent strategies with the knowledge that other players may make mistakes in their choice of action. These mistakes have the feature that "costlier" (in terms of expected payoff) mistakes are less likely.[14] Thus, the Quantal Response Equilibrium model has the advantage that it has a plausible theoretical foundation. It is the only *equilibrium* model we investigate that accommodates mistakes by players. The model is also unique in that it can capture the changing pattern of behaviors both within and across matches. The statistical specification and estimation of this model are presented in the next section.

[11] This assumption is maintained in order to keep the number of free parameters in the model tractable.

[12] These trembles are "real" in the sense that we do not view them in the limit as they vanish to zero, rather we suppose there is a detectable level of error in our subjects' play.

[13] The general theoretical development of this model is presented in McKelvey and Palfrey (1994) and McKelvey and Palfrey (1993). Similar approaches are analyzed in Rosenthal (1989) and Zauner (1993).

[14] This feature is also present in the related work of Beja (1992) on "imperfect equilibrium."

6 Estimation

In this section, we statistically evaluate the non-deterministic models described in the previous section. We begin by selecting a formulation which incorporates all of the "Always Take" models described in the previous section. We define p_t as the probability that a player will choose *Take* at any node of the match she plays in match number t.[15] We attempt to account for learning over time by allowing the probability of a *Take* to change as each subsequent match is played. We select the following, exponential specification for the process by which p_t changes:

$$p_t = 1 - (1 - p_0)e^{-\alpha t}$$

where p_0 and α are parameters to be estimated. This specification has several appealing features. First, it has two structural parameters which have clear interpretations: p_0 is the probability of *Take* used by subjects in the first match of the experiment and α represents the learning rate of the subjects. Second, it allows for a non-linear learning dynamic, instead of requiring a constant rate of learning. Third, this specification encompasses several of the "Always Take" models as special cases. The Rational, Egalitarian and Maximin models are obtained when $p_0 = 1$. The Random model is obtained when $\alpha = 0$, and the Learning model is obtained when there are no constraints on the two parameters, p_0 and α.

In order to estimate these parameters, we define the likelihood function for our data. Let M be the total number of matches in the experiment, indexed by $i = 1, \ldots, M$. Our data consists of n_i, the final node reached in a particular match i, as well as t_i, the match number. With this notation, the log likelihood function for our data is given by

$$\log L = \sum_{i=1}^{M} \log(1 - (1 - p_0)e^{-\alpha t_i}) + \sum_{i=1}^{M} n_i \log((1 - p_0)e^{-\alpha t_i})$$

Table 5 reports the results of maximum likelihood estimation of the unknown parameters p_0 and α. First, we recall that we can reject out of hand the pure versions of the Rational, Egalitarian, and Maximin models, as they all share the common prediction of *Take* at the first node. Each of the experiments has zero likelihood under any of these models, for reasons we have already mentioned.

Regarding the Random model, the table also reports tests for the significance of $\alpha = 0$. These tests are done using the $-2(\log L - \log L_c)$ chi-squared likelihood ratio test, where $\log L_c$ represents the log-likelihood under the constrained model in which $\alpha = 0$. In the overall data as well as in six of the nine experiments, we reject the hypothesis of no learning at conventional significance levels. The Learning model picks up the unraveling of behavior across matches. However, by its very nature, the model is incapable of explaining the increasing probability of *Take* within a match.

[15] We number the matches played in an experiment from 0 to 8 or 9, depending on the number of subjects in the experiment.

Exp. #	Subject Pool	MLE p_0	α	$-\log L$
1	CIT-10	0.526	0.0683	101.19
2	CIT-10	0.416	0.117**	104.13
3	UI-10	0.496	0.1067**	94.15
4	UI-10	0.380	0.0568*	136.04
5	PCC-10	0.554	−0.0198	108.54
6	PCC-10	0.254	0.0335*	184.72
Pooled	10 move	0.410	0.0531**	764.78
7	CIT-6	0.511	0.1292**	85.61
8	UI-6	0.702	0.1225	57.62
9	PCC-6	0.367	0.0917**	102.84
Pooled	6 move	0.493	0.1212**	257.45

Table 5: MLE of Learning model parameters

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

We now turn to the statistical specification and estimation of the Quantal Response Equilibrium model. Recall that one key feature of the model is that costly mistakes are less likely to made by players in the game. We now define a specific quantal response function with this feature to use in our equilibrium analysis. In our setting, suppose that a player at node j in match number t faces a choice between *Take*, with payoff u_T, and *Pass*, with expected payoff u_P. Then the *logistic* response function gives the probability p_t^j that the player will play *Take* as

$$p_t^j = \frac{e^{\lambda u_T}}{e^{\lambda u_T} + e^{\lambda u_P}} = \frac{1}{1 + e^{\lambda(u_P - u_T)}}$$

where $\lambda \geq 0$ is a parameter which is inversely related to the level of error. Specifically, $\lambda = 0$ means that a player's choice is totally random (corresponding to $p = \frac{1}{2}$ in the last estimate), as *Take* and *Pass* are equally likely regardless of the relative expected payoffs. Also, $\lambda = \infty$ means that a player's choice is perfectly rational and exhibits no error; the highest expected payoff choice will be played with certainty. Intermediate values of λ generate varying levels of "noisy play" of the game. An important aspect of this specification is that the probability of implementing a choice is increasing in the equilibrium expected payoff of the choice. Thus, the more costly a mistake would be (in expected payoff), the less likely the player is to make that mistake.

Moreover, the error structure of the Quantal Response Equilibrium model is common knowledge; there is no incomplete information or "mimicking" behavior. However, the possibility of errors, and players' active adjustment to take advantage of this possibility, can account for much of what is usually explained by these techniques. In order to show this by estimating the parameters of the model, we must derive the logistic Quantal

Response Equilibrium (QRE) of our game. We start at the last decision node of the game, where a Blue player would face a choice between *Take* and *Pass* with fixed monetary payoffs. A simple computation gives p_t^N, the probability that Blue will choose *Take* at N, the final node of the game.[16] Given this decision probability, we can compute the expected payoff of a *Pass* by a Red player at the second to last node of the game. This value yields p_t^{N-1}, the take probability of the second to last node. With this, we can calculate the expected payoff of the choices faced by a Blue player at her second to last decision node. Note that these calculations account for mistakes that both she and her opponent might make if she chooses *Pass*. We continue these calculations backwards from the last node of the game until we obtain a complete vector of equilibrium take probabilities, $\boldsymbol{p}_t = \{p_t^1, \ldots, p_t^N\}$.[17]

This process is illustrated in Figure 5, which graphs the equilibrium take probabilities, $\boldsymbol{p}_t$, generated by various values of λ in the six-move constant-sum centipede game. In the figure, λ values are displayed on the horizontal axis on a geometric scale and the various take probabilities are shown as labeled. The figure shows that, as noted above, $\lambda = 0$ corresponds to completely random play at all nodes and $\lambda = \infty$ corresponds to errorless play - the Nash prediction of *Take* with probability 1 at every node. For intermediate values of λ, the model makes a specific prediction about the complete vector of take probabilities, as shown in the figure. The results of the Quantal Response Equilibrium model for the ten-move game is similar.

Insert figure 5 about here.

In order to construct the likelihood function for the logistic specification of the Quantal Response Equilibrium model, we use the equilibrium take probabilities to obtain predicted frequencies of outcomes. Specifically, we let $\hat{f}_t^1 = p_t^1$, $\hat{f}_t^2 = (1 - p_t^1)p_t^2$, $\hat{f}_t^3 = (1 - p_t^1)(1 - p_t^2)p_t^3$, etc. Of course, these values are all derived for a specific value of λ, which is common to all players. Different values of λ will lead to different predicted outcome frequencies. Indeed, we observe such differences across matches in our experiments. The Quantal Response Equilibrium model can easily capture this by supposing that the (common) λ value changes over time. Specifically, we assume that

$$\lambda_t = \lambda_0 + \beta t$$

where λ_0 is the initial value and β represents the rate of change over time. As λ is inversely related to the level of error, the natural interpretation is that as players gain

[16] For example, in the ten-move game, the Blue player can take \$3.08 or play *Pass* and receive \$0.09. We calculate $p_t^N = 1/(1 + e^{\lambda(-2.99)})$. If $\lambda = 1$, then $p_t^N = 0.952$.

[17] This procedure is also used by Zauner (1993) with a different error structure to analyze the centipede experiments of McKelvey and Palfrey (1992).

Exp. #	Subject Pool	MLE λ_0	β	$-\log L$	Vuong Stat.
1	CIT-10	0.745	0.128	98.62	-0.909
2	CIT-10	0.298	0.314**	97.18	-2.679**
3	UI-10	0.684	0.232**	90.56	-1.551
4	UI-10	0.032	0.142*	133.30	-1.169
5	PCC-10	0.769	-0.095	107.09	-1.021
6	PCC-10	-0.479	0.094	198.49	3.173**
Pooled	10 move	-0.028	0.103**	764.08	-0.157
7	CIT-6	0.827	0.324**	79.48	-2.949**
8	UI-6	2.184	0.343*	52.42	-3.415**
9	PCC-6	0.376	0.201*	96.96	-2.596**
Pooled	6 move	0.683	0.292**	243.57	-3.314**

Table 6: MLE for Quantal Response Equilibrium model (logistic specification)

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

more experience they are less prone to make mistakes. Thus, if we again let n_i represent the final node reached in a particular match i, the log likelihood function is given by

$$\log L = \sum_{i=1}^{M} \log[(\hat{f}_t[\lambda_0, \beta])^{n_i}]$$

where M is the total number of matches in the experiment.

Table 6 reports the estimates of λ_0 and β for the Quantal Response Equilibrium model in each of the experiments. We again use the $-2(\log L - \log L_c)$ chi-squared likelihood ratio test to reject the hypothesis that $\beta = 0$ in six of our experiments,[18] at conventional confidence levels. Thus, we conclude that in most cases there is a learning trend in the direction of fewer errors in the later stages of the experiments.

At this stage, we have parameter estimates for both the Learning model and the Quantal Response Equilibrium model. The next question to address is which model better explains the data, the modified "Always Take" model or the Quantal Response Equilibrium model. As both models generate predicted frequencies which are compared to the actual data to produce likelihood scores, we can compare these log-likelihood scores to determine which model fits the data better. The results are striking. In eight of the nine individual experiments, the Quantal Response Equilibrium model does better.

As the Quantal Response Equilibrium model is not nested with the Learning model, we cannot use the standard chi-squared likelihood ratio test to determine if these dif-

[18]Learning is significant in *all* of the six-move games but only half of the ten-move games. Error rates in ten-move games appear to be somewhat higher overall, as well.

ferences are significant. We instead use Vuong's (1989) model selection test for strictly non-nested models. This test is based on the asymptotic distribution of the likelihood ratio statistic under general conditions. In Table 6 we report the Vuong test statistic, the difference in the maximum log-likelihood values for the two models suitably normalized, for each of the experiments and the pooled ten-move game and pooled six-move game data. This test statistic is to be compared to a critical value c from the standard normal distribution for some significance level. If the test statistic is smaller than $-c$, then we can reject the null hypothesis that the models are equivalent in favor of the Quantal Response Equilibrium model being better than the Learning model. The opposite conclusion is reached when the test statistic is larger than c. Finally, if the test statistic is between $-c$ and c, we cannot discriminate between the two competing models given the data, at that significance level.

Of the eight experiments with a lower log likelihood under the Quantal Response Equilibrium model, four have likelihood ratios that select this model over the Learning model, at the .01 level. This includes all of the six-move sessions. In addition, the pooled six-move game data selects the Quantal Response Equilibrium model at the same significance level. The other four experiments individually and the pooled ten-move data do not offer differences significant enough for model selection.[19] Therefore we reject the Learning model in favor of the Quantal Response Equilibrium model in the six-move games. While we cannot reject the Learning model outright in the ten-move centipede games, these data support the Quantal Response Equilibrium model more than the Learning model.

The Quantal Response Equilibrium model not only fits well, it also accounts for the two main features of the data that we listed in a previous section. First, for the equilibrium vector of take probabilities $\boldsymbol{p}_t$, we find that the predicted take probabilities increase with later nodes in the game. Figure 5 illustrates this for the six-move centipede game. For all λ values we estimate, later nodes of the game correspond to larger equilibrium take probabilities. Thus, the model captures the feature that *within* matches, players are more likely to choose *Take* at later nodes. The second main feature of our data is that *across* matches, players are more likely to choose *Take* as they gain experience with the game. The Quantal Response Equilibrium model captures this by allowing the value of λ to change over time. Again this is illustrated for the six-move game by Figure 5. As the curves in the figure are all increasing, higher values of λ (over time) lead to higher equilibrium take probabilities. Therefore, the Quantal Response Equilibrium model accommodates our main qualitative findings, in addition to offering the best fit of the experimental data among our models.

[19]Experiment # 6 has a test statistic that indicates that the data from the experiment is better explained by the Learning model. It also contains a much lower rate of Taking than in the other experiments.

7 Conclusion

We have proposed several different models for our data on constant-sum centipede games. We find that the Rational, Egalitarian and Maximin models do not fare well, even when adapted to allow for statistical behavior. The Random model does not explain unraveling across experiments. The Learning model fits the behavior across matches relatively well, but it does not account for the finding that subjects take with increasing probability within a match. Also, it is a model that is not internally consistent, in that if subjects knew the model that was estimated, best response behavior by the subjects would lead to different data. Finally, variants of the incomplete information model used in McKelvey and Palfrey (1992) are implausible here because altruistic behavior does not imply that a subject should ever (much less *always*) pass. Further, even if some other rationale could be found for subjects always passing (such as reciprocation), there is no evidence in the data that such individuals exist.

Among the models we evaluate, the Quantal Response Equilibrium model best explains the data. It offers a better fit than the Learning model and, as it is an equilibrium model, is internally consistent. It also accounts for the pattern of increasing take probabilities within a match. These facts lend strong support to the Quantal Response Equilibrium model.

The results reported here suggest a natural further test of the Quantal Response Equilibrium model. The Altruism explanation offered by McKelvey and Palfrey for their data seems doubtful given our finding that there is no evidence of altruism in the constant-sum centipede data. This suggests that this earlier data needs to be reexamined, and we believe such an analysis with the Quantal Response Equilibrium model would be very successful.

Appendix 1: Experiment Instructions

DECISION MAKING EXPERIMENT

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.

You have been divided into two groups, with 10 subjects each. The groups are called the RED group and the BLUE group. 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.

[DISPLAY PAYOFF TABLE]

In this experiment, you will be making the following decision, for real money.

First, you are matched with a subject of the other color. There is a sum of money, divided into two equal piles, the Top pile, and the Bottom pile. At the beginning of the match the Top Pile has $1.60 and the Bottom Pile has $1.60.

RED has the first move and can either "Pass" or "Take". If RED chooses "Take", RED gets the Top Pile of $1.60, BLUE gets the Bottom pile of $1.60 cents, and the match is over. If RED chooses "Pass", one fourth of the Bottom Pile is moved into the Top Pile, and it is BLUE's turn.

The Top Pile now contains $2.00 and the Bottom Pile $1.20. BLUE can take or pass. If BLUE takes, BLUE ends up with the Top pile of $2.00 and RED ends up with the Bottom pile of $1.20 and the match is over. If BLUE passes, one fourth of the Bottom Pile is moved into the Top Pile, and it is RED's turn again.

This continues for a total of ten [six] moves, or five [three] moves for each subject. On each move, if a subject takes, he or she gets the Top pile, the other subject gets the Bottom pile, and the match is over. If he or she passes, one fourth of the Bottom Pile is moved to the Top Pile, and it is the other subject's turn.

The last move of the match is move ten [six], and is BLUE's move, if the match gets this far. The Top pile now contains $3.08 [$2.82] and the Bottom pile contains $0.12 [$0.38]. If BLUE takes, BLUE gets the Top pile of $3.08 [$2.82] and RED gets the Bottom pile of $0.12 [$0.38] cents. If BLUE passes, then one fourth of the Bottom pile is moved to the Top Pile. RED then gets the Top Pile, containing $3.11 [$2.92] and BLUE gets the Bottom Pile, containing $0.09 [$0.28].

[GO OVER THE TABLE TO EXPLAIN WHAT IS IN EACH CELL]

[DISPLAY MATCHING SCHEME]

The experiment consists of 10 matches. In each match, you are matched with a different subject of the other color from yours. Thus, if you are a BLUE subject, in each match, you will be matched with a RED subject. If you are a RED subject, in each match you are matched with a BLUE subject. In the first match, you are matched with the subject of the other color with the same number as yours. So Red #1 is matched with Blue #1, Red #2 is matched with Blue #2, etc. In each successive match if you are Red, you are matched with Blue subject with the next higher number. If you are Blue, you are matched with the Red subject with the next lower number. So in match 2, Red #1 is matched with Blue #2, Red #2 with Blue #3, etc. 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.

[PAUSE FOR QUESTIONS]

We will now begin the computer instruction session. During the instruction session, we will teach you how to use the computer by going through a few practice matches. 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 matches.

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 SUBJECT'S TO ENTER COLORS]

You now see the experiment screen. The screen is divided into three sections. The top of the screen tells you the current match number, your subject number and the subject number of the subject with whom you are currently matched. The center of the screen shows the payoff table. This is the same as the table that was just displayed. The bottom of the screen tells you what is currently happening, and prompts you for input. Since the experiment has not begun yet, the bottom part of the screen simply tells you to wait. 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 match.

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

[MASTER HIT KEY]

You now see on the bottom part of the screen that the first match has begun. If you are a RED subject, you are told that it is your move, and are given a description of the choices available to you. If you are a BLUE subject, you are told that it is the other subjects move, and are told the choices available to the other subject. On the payoff table, you see that there are two sets of red arrows pointing to the first column of the table. This indicates that it is the first move of the match and that it is the Red subject's turn.

Will all the RED subjects now choose PASS by typing a P on your terminals now.

[WAIT FOR SUBJECTS TO CHOOSE]

Notice that the two sets of red arrows on the payoff table have now moved to the second column and are Blue. This indicates that Red chose Pass, and that it is now the second move, which is Blue's turn. Note that the Bottom pile has decreased by one fourth to $1.20 and the top pile has increased to $2.00.

On the bottom part of the screen, the BLUE subjects are now told that it is their turn to choose, and are told the choices they can make. The RED subjects are told that it is the other subject's turn to choose, and are told the choices that their other subject can make.

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

[WAIT FOR SUBJECTS TO CHOOSE]

Since BLUE chose T, the first match has ended. This is indicated on the payoff table by the fact that the double set of arrows is now white, and point to the second column. The payoff to you and your opponent in the highlighted column, and are also recorded below the payoff table.

Despite the fact that your match is over (in fact all matches are over), the program requires that you enter a response at each turn. The computer will prompt you to hit "y" and return until each subject has entered a total of five [three] responses. Note that this has no effect on your payoffs. Your match is already over, and your payoffs are those recorded.

[TERMINATE MATCH WITH SEQUENCE OF "y" KEYSTROKES]

On the Bottom part of the screen, you are told that the match is over, and that the next match will begin shortly. You are also prompted to record your payoffs, and the subject with whom you are matched. Please do that now.

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 match, and you would be paid this amount for that match at the end of the experiment. The total you earn over all ten real matches is what you will be paid for your participation in the experiment.

[WAIT FOR SUBJECTS TO RECORD PAYOFFS]

[HIT KEY TO START NEXT MATCH]

You have now been matched with a new subject. The subject number of the subject with whom you are matched is in the upper right hand corner of your screen. We will now proceed to the second practice match.

[MASTER HIT KEY]

The second match has begun. The rules for the second match are exactly like the first. The RED subject gets the first move.

[DO ALL P UNTIL BLUE's LAST MOVE]

Now notice that it is BLUE's move. It is the last move of the match, The Top Pile now contains $3.08 [$2.82], and the Bottom Pile contains $0.12 [$0.38]. If the BLUE subject chooses TAKE, then the match ends. The BLUE subject receives the Top Pile and the RED subject receives the Bottom Pile. If the BLUE subject chooses PASS, one fourth of the Bottom Pile is moved to the Top Pile, and then the match ends. The RED subject receives the Top Pile, which now contains $3.11, [$2.92] and the BLUE subject receives the Bottom Pile, containing $0.09 [$0.28].

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

[WAIT FOR SUBJECTS TO CHOOSE]

Appendix 2: Experimental Data

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

COLUMN 1: Experiment number.

COLUMN 2: Subject number of Red player.

COLUMN $2+j$: Outcome of match j. This is the number of passes before the first *Take*.

The matching scheme used is the same as in McKelvey and Palfrey (1992). In match j, Red subject i is matched with Blue subject $[(i+j-1)\text{mod}m]$, where m is the number of subjects of each color in the experiment. Thus, with twenty total subjects, in the first match Red i is matched with blue i. In the second match, Red i is matched with Blue $1+i$, except for Red 10, who is matched with Blue 1.

1	1	0	0	0	0	0	0	0	0	0	0
1	2	2	1	1	0	1	1	0	1	0	1
1	3	2	1	1	0	0	0	0	0	0	0
1	4	0	0	0	0	0	0	0	0	0	0
1	5	0	0	0	0	0	0	0	0	0	0
1	6	2	1	2	1	2	2	2	1	0	1
1	7	1	2	1	1	0	0	0	0	0	0
1	8	0	0	1	1	0	0	1	0	2	1
1	9	1	1	1	2	1	1	1	1	1	2
1	10	1	1	4	0	1	0	0	0	0	1

Experiment 1
CIT 10 move

2	1	2	1	2	1	1	2	1	1	0	1
2	2	2	1	1	0	0	1	0	0	1	0
2	3	2	1	1	0	0	0	0	0	0	0
2	4	2	1	4	1	1	0	1	1	1	1
2	5	0	0	0	0	0	0	0	0	0	0
2	6	0	2	2	1	0	0	0	0	0	0
2	7	2	2	1	1	1	0	0	0	0	0
2	8	2	1	1	1	1	0	1	0	0	0
2	9	0	0	0	0	0	0	0	0	0	0
2	10	1	0	2	1	1	1	1	1	1	1

Experiment 2
CIT 10 move

3	1	0	0	2	0	0	0	0	1	0	0
3	2	0	0	0	0	0	0	0	0	0	0
3	3	2	1	1	1	0	0	0	0	0	0
3	4	0	0	0	0	0	0	0	0	0	0
3	5	2	1	1	2	2	1	1	1	0	0
3	6	1	2	2	1	1	2	1	2	1	1
3	7	1	1	1	0	0	1	0	0	0	1
3	8	0	0	0	0	0	0	0	0	0	0
3	9	1	1	0	1	0	0	0	0	0	0
3	10	2	2	1	1	1	2	1	1	0	0

Experiment 3
UI 10 move

4	1	2	1	1	1	1	0	1	1	0	2
4	2	0	1	0	0	0	0	0	0	0	0
4	3	0	0	0	0	0	0	0	0	0	0
4	4	1	2	2	1	0	1	1	1	0	0
4	5	2	4	2	3	2	1	2	1	1	1
4	6	0	0	0	0	2	0	1	0	1	0
4	7	6	1	3	3	3	1	1	1	0	3
4	8	2	2	1	1	2	1	1	1	1	2
4	9	3	2	1	1	2	1	1	2	1	0
4	10	2	1	0	2	0	1	0	0	0	1

Experiment 4
UI 10 move

5	1	2	1	1	2	1	2	1	3	1
5	2	2	1	1	0	0	0	0	0	0
5	3	1	1	1	3	3	1	1	4	5
5	4	0	0	0	0	4	1	2	1	1
5	5	1	1	0	1	0	0	1	0	1
5	6	0	0	0	0	0	0	0	0	0
5	7	0	2	1	1	2	1	1	1	3
5	8	0	1	1	0	1	0	0	1	0
5	9	0	0	5	1	1	0	1	0	0

Experiment 5
PCC 10 move

6	1	5	1	2	3	2	2	2	2	2	4
6	2	2	9	1	3	4	5	2	3	1	1
6	3	0	2	4	5	4	3	3	1	1	1
6	4	0	0	0	0	0	5	1	1	1	1
6	5	6	5	3	4	3	2	2	1	1	1
6	6	5	4	3	3	2	1	1	0	0	0
6	7	0	0	0	0	0	0	0	0	0	0
6	8	3	4	1	2	1	1	0	2	0	0
6	9	3	2	3	1	1	2	3	2	4	1
6	10	2	2	2	2	2	2	2	2	1	2

Experiment 6
PCC 10 move

7	1	0	0	0	0	0	0	0	0	0	0
7	2	1	1	1	0	0	0	0	1	1	0
7	3	0	0	0	0	0	0	0	0	0	0
7	4	1	0	1	0	2	1	0	0	1	0
7	5	1	1	2	0	1	0	1	1	1	1
7	6	0	0	0	0	0	0	0	0	0	0
7	7	1	1	1	0	0	0	0	0	0	0
7	8	1	2	2	1	1	1	0	1	0	1
7	9	2	2	1	1	0	0	1	0	0	0
7	10	2	1	1	1	0	0	0	0	0	0

Experiment 7
CIT 6 move

8	1	1	1	1	0	0	0	1	0	0	0
8	2	1	0	1	0	0	0	0	0	0	0
8	3	0	1	0	1	0	0	1	0	0	1
8	4	0	1	0	0	0	0	0	0	0	0
8	5	0	0	0	0	0	0	0	0	0	0
8	6	0	0	0	0	0	0	0	0	1	1
8	7	1	1	0	0	0	0	0	0	0	0
8	8	1	1	0	0	1	0	0	1	0	0
8	9	1	0	0	0	0	0	0	1	0	0
8	10	0	1	0	0	0	0	0	1	0	0

Experiment 8
UI 6 move

9	1	1	0	1	0	0	0	0	0	0
9	2	1	2	1	2	1	1	0	1	0
9	3	4	1	1	1	0	1	0	0	0
9	4	1	0	1	0	0	2	1	0	0
9	5	2	2	2	1	1	1	0	0	0
9	6	1	3	0	1	0	1	1	0	0
9	7	3	1	1	1	1	1	1	0	0
9	8	1	2	2	1	2	1	2	1	1
9	9	1	1	1	1	2	1	1	1	2

Experiment 9
PCC 6 move

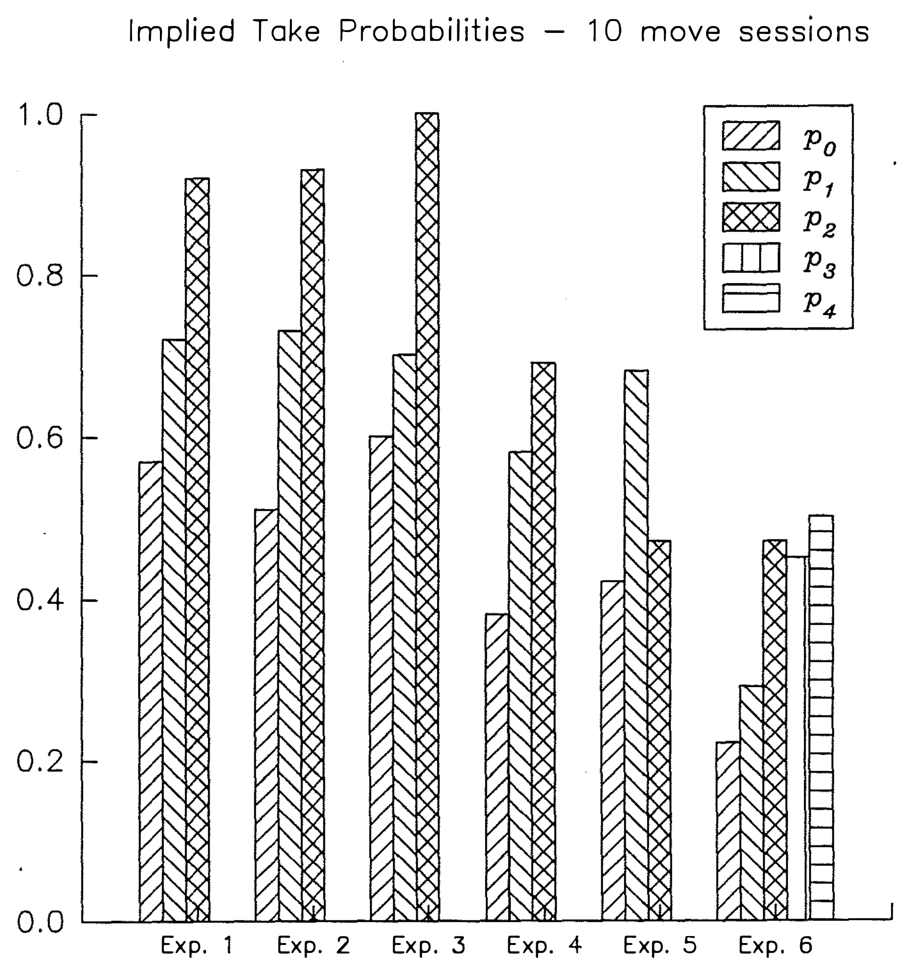
Implied Take Probabilities — 10 move sessions
1.0
0.8
0.6
0.4
0.2
0.0
p_0
p_1
p_2
p_3
p_4
Exp. 1
Exp. 2
Exp. 3
Exp. 4
Exp. 5
Exp. 6

Figure 3

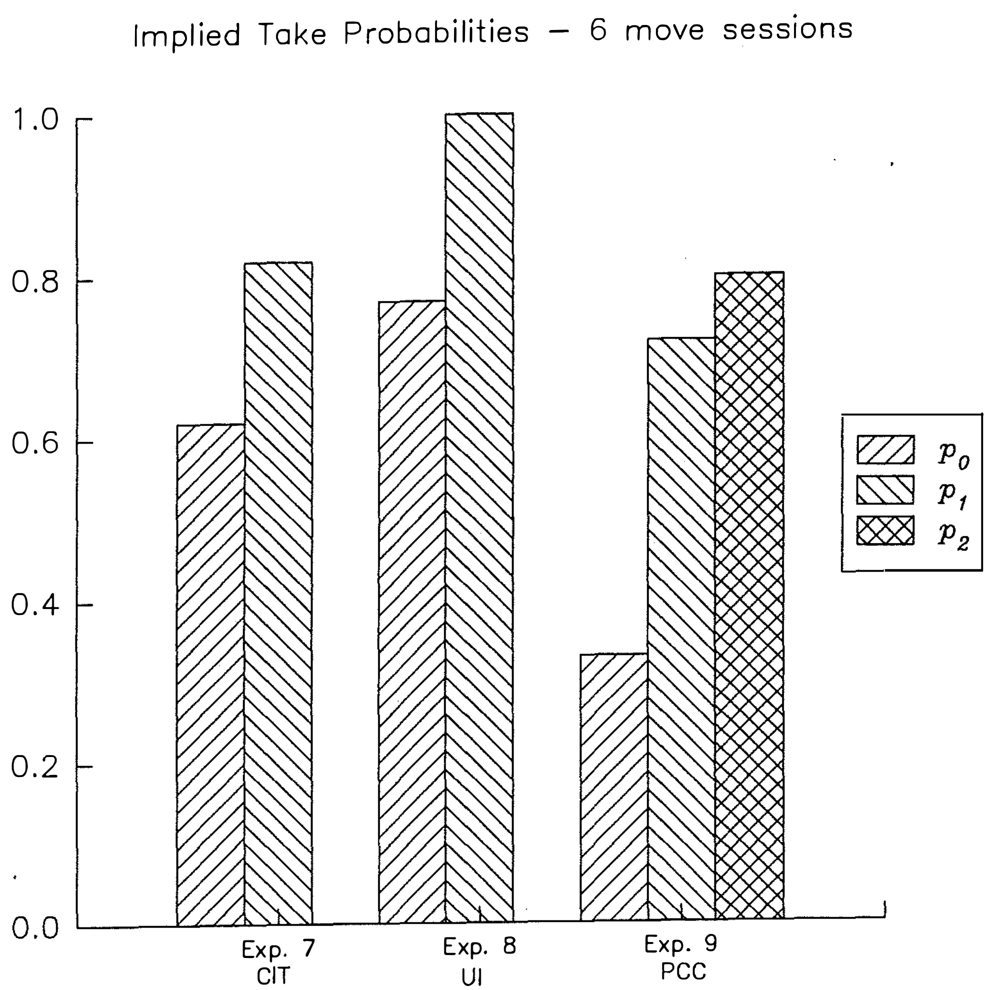

Implied Take Probabilities — 6 move sessions
1.0
0.8
0.6
0.4
0.2
0.0
Exp. 7
CIT
Exp. 8
UI
Exp. 9
PCC
p_0
p_1
p_2

Figure 4

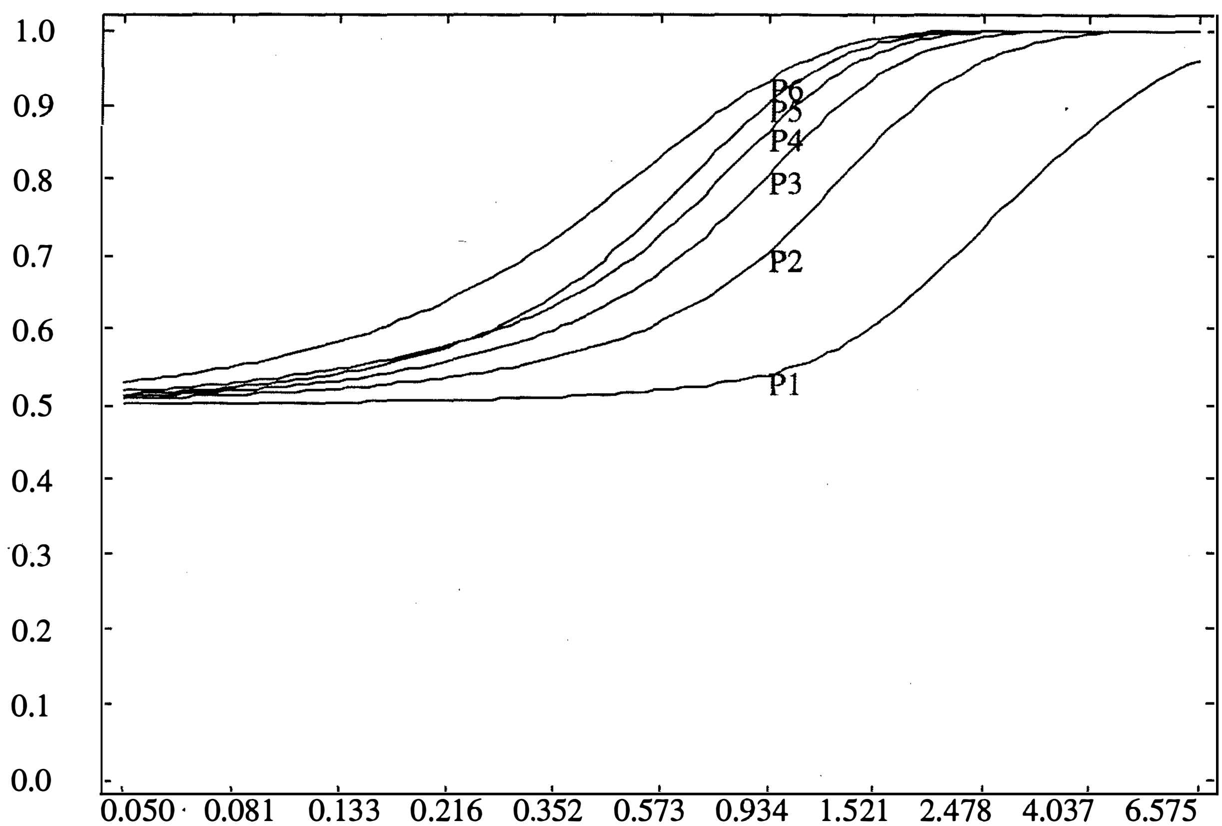

Figure 5: Quantal Response Equilibrium of the Six-Move Constant-Sum Centipede Game.

References

Aumann, R. 1988. Preliminary notes on integrating irrationality into game theory. Mimeo. International Conference on Economic Theories of Politics, Haifa.

Basu, K. 1990. On the non-existence of a rationality definition for extensive games. *International Journal of Game Theory* 19:33–44.

Beja, A. 1992. Imperfect equilibrium. *Games and Economic Behavior* 4:18–36.

Binmore, K. 1987. Modeling rational players. *Economics and Philosophy* 3:179–214.

Cooper, R., D. V. DeJong, R. Forsythe, and T. W. Ross. 1989. Communication in the battle of the sexes game: Some experimental results. *RAND Journal of Economics* 20:568–587.

El-Gamal, M. and D. Grether. 1993. Uncovering behavioral strategies: Likelihood based experimental data mining. Social science working paper # 850, California Institute of Technology.

Güth, W., P. Ockenfels, and M. Wendel. 1993. Efficiency by trust in fairness? multiperiod ultimatum bargaining experiments with an increasing cake. *International Journal of Game Theory* 22:51–73.

Kreps, D. M. 1990. *A course in microeconomic theory.* New Jersey: Princeton University Press.

McKelvey, R. D. and T. R. Palfrey. 1992. An experimental study of the centipede game. *Econometrica* 60(4):803–836.

McKelvey, R. D. and T. R. Palfrey. 1993. Quantal response equilibria for normal form games. Social science working paper # 883, California Institute of Technology.

McKelvey, R. D. and T. R. Palfrey. 1994. Quantal response equilibria for extensive form games. Presented at Conference on Learning Equilibrium Behavior, Texas A & M University.

Megiddo, N. 1986. Remarks on bounded rationality. Technical report. IBM Research Report RJ 54310, Computer Science.

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. J. 1993. Common belief and the theory of games with perfect information. *Journal of Economic Theory* 59:257–274.

Rosenthal, R. 1981. Games of perfect information, predatory pricing, and the chain store paradox. *Journal of Economic Theory* 25:92–100.

Rosenthal, R. 1989. A bounded-rationality approach to the study of noncooperative games. *International Journal of Game Theory* 18:273–292.

Stahl, D. and P. Wilson. 1993. On players' models of other players - a new theory and experimental evidence. Mimeo, University of Texas.

Thaler, R. H. 1988. The ultimatum game. *Journal of Economic Perspectives* 2(4):195–206.

Varian, H. 1984. *Microeconomic analysis*, second ed. New York: Norton & Co.

Vuong, Q. H. 1989. Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica* 57(2):307–333.

Zauner, K. 1993. Bubbles, speculations and a reconsideration of the centipede game experiments. Mimeo, University of California, San Diego.