DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES
CALIFORNIA INSTITUTE OF TECHNOLOGY
PASADENA, CALIFORNIA 91125

Repeated Play, Cooperation and Coordination: An Experimental Study

Thomas R. Palfrey
California Institute of Technology

Howard Rosenthall
Carnegie-Mellon University

SOCIAL SCIENCE WORKING PAPER 785

January 1992

Abstract

An experiment was conducted to test whether discounted repeated play leads to greater cooperation and coordination than one-shot play, in a public good environment with incomplete information. The experiment was designed so that, theoretically, repeated play can sustain equilibria with higher group earnings than result in the one-shot Bayesian Nash equilibrium. The design varied a number of environmental parameters, including the size of the group, the marginal rate of transformation between the public and private good, and the statistical distribution of marginal rates of substitution between the public and private good. Marginal rates of substitution were private information but the statistical distribution was common knowledge. The results indicate that repetition leads to greater cooperation, and that the magnitude of these gains depends both on the ability of players to monitor each other's strategy and on the underlying environmental parameters.

Keywords: Repeated Games, Experiments, Public Goods, Bayesian Games

JEL Classificatioin Numbers: 026 215

1. INTRODUCTION

The theory of repeated games[1] has developed rapidly in the last 10 years. A common theme of the large body of work in this area is that repetition enables players to reach outcomes that would otherwise not be possible. Of particular interest is the fact that cooperative outcomes can be sustained as noncooperative equilibria of repeated symmetric games even though they fail to be equilibrium outcomes in the stage-game. Competitive players can do better with repetition than without repetition. On the other hand, one difficulty of most of this theory is that nearly everything can be supported as an equilibrium in the repeated game, not only the cooperative outcome. Thus, to say that cooperation is a clear prediction of the theory is too strong. Nevertheless, the suggestion that repetition of an asymmetric stage game whose Nash equilibria are not Pareto optimal will usually lead to better outcomes is widely viewed as an implication of the theory.

As an empirical matter, this proposition invites testing. The literature in experimental economics and experimental psychology contains many studies of finitely repeated games.[2] In contrast, to our surprise few laboratory experiments have been conducted which attempted to induce an infinite horizon discounted supergame in a laboratory. In addition, there is a paucity of research on the corresponding one-shot games.

What evidence there is about one-shot games comes either from the last round of play of a finitely repeated game with a known terminal period, or from experiments where subjects played the game only once.[3] The former evidence (which generally indicates less cooperation in the terminal period[4]) is suspect because, in the context of a repeated game it is difficult to distinguish between the use of trigger strategies, as opposed to simply myopic one-shot play. The latter evidence (virtually all of it in the context of public goods provision) Marwell and Ames [1979,1980,1981] and Schneider and Pommerehne [1981] is open to the criticism that subjects lacked sufficient task experience. In fact, inexperience has been demonstrated to be an important factor in explaining cooperative play in these games (Isaac, Walker and Thomas [1984] and Andreoni [1988]).

An alternative, and we believe better, approach to studying one-shot games involves matching schemes so that each player's opponents change from game to game. This has been used in auction games[5] (Palfrey [1985]), bargaining games (Roth and Murnighan [1982], and elsewhere in situations where collusion or supergame effects are thought to be potential problems (McKelvey and Palfrey [1989]). We have elsewhere used a version of this design technique to study one-shot public goods games (Palfrey and Rosenthal [1991a,1991b]), and it is becoming the standard methodology for studying one-shot games.

The experiment reported in this paper was designed explicitly to permit a careful comparison of one-shot games with the corresponding discounted, infinitely repeated games. The experiment involved over 200 hundred subjects and thousands of stage games, which varied in their payoff matrices, information conditions, number of players, and monitoring technologies.

The basic stage games played were all variants of a well-known N-person voluntary contribution public goods game where the production decision is binary: a public good is produced if contributions exceed a threshold. The stage game also has an element of private information, in that players have different, privately known marginal rates of substitution between the private good (that they can use to contribute) and the public good. In nearly all of the games, the symmetric equilibria of the stage game are Pareto suboptimal because of both a free-rider problem and a coordination problem, but Pareto improving outcomes are supportable in the repeated game.

Roughly half of our data is generated by a repeated design, in which a group of individuals play the same game repeatedly with a random stopping rule. The other half of the data is a one-shot design, where individuals play the same game many times, but group assignments are randomly and anonymously reassigned after each game is played.

We use our data to investigate to what extent, and in what ways, repeated games can improve the degree to which the players coordinate and produce the public good. The stage game in fact presents two impediments to collective action as a result of the "threshold" nature of the public good technology. First, a player who believes the threshold is being met by others finds it in his interest to "free ride" and not contribute. Second, a player who believes the threshold will not be met, even with his contribution, finds it in his interest not to contribute.

Players' strategic calculations weigh these two disincentives against the potential benefit from obtaining the public good. Since opportunities to explicitly coordinate are absent, it is reasonable to assume that players will only use symmetric strategies represented by "cutpoints": if the value of the private good is below the cutpoint, contribute; otherwise, do not. The stable equilibria for such strategies, for the parameters we have chosen, either imply no contribution or relatively low rates of contribution.

When the games are repeated, higher levels of contribution can be sustained, even when the players continue to use symmetric cutpoint strategies. The expectation of future reciprocation diminishes the effect of the current incentive to free ride. An even higher degree of coordination could be achieved if the players could get together and sign a binding agreement that treated everyone symmetrically. They could agree to always meet the contribution threshold exactly. A random device would determine who would contribute for each repetition of the stage game. Such an agreement would lead to ex ante efficient behavior. Decentralized versions of such agreements can always be sustained, for our parameters, in the repeated game (without contracting). For example, we might observe rotation. In, say, a 2 of 3 game, players "B" and "C" could contribute on the first round, "A" and "C" on the second, "A" and "B" on the third, "B" and "C" on the fourth, and so on. Alternatively, the players might separate, asymmetrically, into sets of "contributors" and "non-contributors."

The above discussion suggests three distinct patterns of behavior that one might expect to observe under repeated play:

1. Symmetric cutpoints higher than the one-shot cutpoints.
2. Separation into roles of contributor and free-rider.
3. Rotation schemes.

The expected patterns of behavior that would be exhibited under these three different solutions to the inefficiency problem are very different[6], with the exception of unanimity games (where all of the private good is required for production of the public good). In this case, the three varieties collapse into a single repeated play equilibrium where all contribute.

What we observe in the repeated laboratory games is closest to the first two of these three patterns. On average, individuals contribute more under the repeated-game treatment than in the one-shot treatment. Although there is greater variance across individuals in their contribution rates in repeated games than in one shot games, there is only very weak evidence indicating the second pattern. Indeed, only a few subjects were pure or almost pure free riders and no subject was a pure contributor. We do not observe any group rotation schemes that would support the third pattern. Only one individual rotated throughout an entire repeated game session. His behavior was not reciprocated by the other members of his group; they used simple cutpoint strategies. While a few individuals at times may have attempted to encourage other players to rotate by systematically alternating their own contribution choices, these attempts invariably fail. Our results contrast sharply with results for analogous symmetric two-person games with complete information, where efficient alternation often emerges after only a few repetitions (Prisbrey 1991).

We suspect that, without communication, the coordination problem is too complicated to overcome simply by repeated play. Even with communication, the coordination problem is by no means easy to overcome. In a one-shot game with preplay communication, Palfrey and Rosenthal (1991a) find no efficiency gain compared with no communication. In one-shot battle of the sexes games, Cooper et al. (1989) find increased efficiency with one-way communication, where the communicator effectively chooses the asymmetric equilibrium that favors her. Neither Palfrey and Rosenthal nor Cooper et al. have a repeated-game design, so alternation schemes are not possible.

Indeed, we find that the extent to which repeated play improves over one-shot play appears to be small. Not only do earnings fail to reach the ex ante efficient level that could be produced by rotation, they generally fail to reach the level that could be achieved by symmetric cutpoint strategies. In contrast, aggregate earnings in the one-shot play setting are almost perfectly predicted by the noncooperative solution (Palfrey and Rosenthal, 1991a). This contrast between our one-shot and repeated play results is not encouraging news for those who might wish to interpret as gospel the oft-spoken suggestion that repeated play with discount rates close to 1 leads to more cooperative behavior. True enough, it does--but not by much.

Perhaps this is not surprising, given both the general difficulty of free rider problems and the difficulty of the coordination problem in the experiment we conducted. After all, our experimental games had more than two players, incomplete information, and, in some cases, multiple symmetric equilibria in the one-shot game. The power of repeated play might be more forceful in simpler experimental environments. But we suspect that the power of repeated play is, if anything, even less forceful in the more complex natural settings that we are ultimately trying to understand. While repeated play makes possible cooperative gains, other ingredients are also needed for the lion's share of these gains to be reaped.

2. THE EXPERIMENTAL DESIGN

The subjects used in the experiment were 228 students recruited from the campuses of Carnegie-Mellon University, California Institute of Technology, and Pasadena Community College. Except where noted in Table 1, each session was run on networked personal computers that. Subjects participated in a sequence of three sessions, one immediately following the other.[7] Component sessions in a triple session typically employed different treatments (for example, a change in the parameters). Table 1 provides the details. Each subject was paid privately in cash immediately following the final session. A triple session lasted between 45 minutes and 2 hours, including the preliminaries (instructions, questions and answers, quiz, and practice rounds). No subject participated in more than one triple session. Some Caltech subjects had prior experience with computerized double auction market experiments.

[Table 1 about here.]

In a typical session, 12 subjects[8] were seated at terminals, separated by partitions, and assigned identification numbers. Instructions were read aloud to everyone at the same time.[9] Each session was divided into a sequence of periods (or rounds). In each period, the group of 12 subjects was either divided into four three-person groups or three four-person groups. Each group then independently played the following game.

THE PUBLIC GOODS GAME

Each subject was endowed with a "token", with value of c_i Francs (an artificial laboratory currency with a publicly known dollar-exchange rate). The computer screen displayed the individual's token value but not the token values of other group members. Thus, prior to play of a stage game, individual token values were always private information. Subjects were told that they could either "spend" their own token or "keep" it. If at least W subjects in a group chose "spend", then each member in the group earned V

Francs. In addition, nonspenders ("keeper" subjects) also earned their token values. If there were fewer than W spenders in the group, spenders earned 0 and nonspenders earned their token values. Token values were drawn independently from a uniform distribution using a random device on the computer in one-franc increments between 1 and a known maximum token value, C. This distribution was publicly announced and explained to all subjects during the instruction period.

To insure that, insofar as possible, all aspects of the game except for the exact draws of token values and the personal identities of the other members in the group were common knowledge, we used the following procedure. First, the rules were publicly announced in great detail. Second, two practice games were played, to help the subjects familiarize themselves with the keyboard and the computer screen. Third, a quiz was given after the practice rounds. Any incorrect answers by a subject were corrected in private, and then the correct answers were read aloud and explained to all subjects, together.

After the quiz, the first (stage) game began. After every subject had made a spending decision, all subjects were told what the other members in their group did, their payoffs were calculated for them, and the session then proceeded to the next game (period). In the new game, subjects' new token values were again drawn randomly by the same procedure as the last game, independently from past draws. During the course of the experiment, subjects could press "H" on their terminals and obtain a history of the play of the last 25 games they had participated in under the current treatment.

EXPERIMENTAL TREATMENTS

In terms of treatments, the design is complicated and multifactorial, and nonsquare. There were two primary treatments variables: the parameters of the game, and whether the groups were randomly reassigned after each round or were fixed for the entire session. The sessions where groups were randomly reassigned after each play of the game approximate the case of "one-shot" play of the game. The sessions where groups were fixed approximate the case of "repeated" play.

A. PARAMETERS

The parameter treatment was three dimensional, since we varied the threshold, W, the group size, N, and the maximal marginal rate of substitution of the private good for the public good, C. A total of six different parameter conditions were used.

B. RANDOM GROUPS VS. FIXED GROUPS

In the "repeated" condition, each subject was assigned to a group whose membership was the same in every game of that treatment. If more than one fixed-group treatment was used in a "triple session", then group membership changed between sessions so that individuals were assigned to a new group with a completely different membership. Except for this piece of information, subjects were told nothing about the identities of the other members of their group.

The repeated game used a random stopping rule to determine when the session would end. Each session began with 20 games. After the twentieth game, a 10-sided die was rolled, and the session ended if a 4 was rolled. Otherwise, the session continued on to the 21st game. The rolling of the die followed every game thereafter, until a 4 was finally rolled.[10]

In the random-group condition (or "one-shot" condition), subjects were randomly reassigned to a new group after every game. They were never told the identity of current, future, or past group members. Each random treatment session lasted a fixed number of games, either 20, 25, or 30, and subjects were always informed of this fixed number at the beginning of the session.

C. THE "REVEAL" TREATMENT: PERFECT VS. IMPERFECT MONITORING

In addition, there was a secondary treatment variable, called "reveal". In some of the sessions, after a game was played, subjects were informed not only what everyone in their group chose but also what their exact token values had been. Sessions where we did this are referred to as "reveal" sessions, and other sessions are called "no reveal" sessions. The motivation for the reveal/no-reveal manipulation is that theory suggests cooperation will be easier to sustain when the token values are revealed after each game, allowing subjects to more accurately monitor the other subjects' strategies. This monitoring permits a richer class of trigger strategies, which may be used to sustain higher contribution rates. In contrast, optimal rotation schemes can be supported even if token values are not revealed after each play of the stage game. This is discussed in more depth in the next section.

Table 1 summarizes relevant design information for all sessions, with one-shot sessions shown in Table 1A and repeated-game sessions in 1B.

3. EQUILIBRIUM

In what follows, without loss of generality, we assume that the value of the public good, V, is the same for every player, and normalize it at 1. Therefore, the marginal rate of substitution between the private good and the public good for individual i is simply equal to i's token value.[11]

A "cutpoint rule" for player i is a strategy with the property that there exists a critical cost, $\hat{c}_i$, with the property that i spends[12] if and only if $c_i \leq \hat{c}_i$.

BAYES-NASH EQUILIBRIUM IN THE ONE-SHOT GAME

In the one-shot game, it is easy to show that a cutpoint strategy is an optimal response for player i, given any strategy profile of the other players in the group. Therefore, we restrict attention to such strategies in our analysis of the one-shot game. The equilibrium to this game has been derived in Palfrey and Rosenthal (1988,1991a,1991b); the reader is referred to those papers for details.[13]

The basic features of the equilibrium are as follows:

--If $W > 1$ then there always exists an equilibrium where all players adopt a cutpoint $c^* = 0$. That is, no one ever spends. This is never true if $W = 1$.

--If $W = N$, then there exists another equilibrium at $c^* = C$, if and only if $C \leq 1$.

--The other equilibria are characterized as the set of all c^* between 0 and 1 that satisfy the equation:

$$c^* = Q(c^*, W, N, C) \tag{1}$$

where Q is the probability that a player is pivotal; this is the probability that exactly W-1 out of the N-1 other players contribute, given that they are using the cutpoint rule c*, and given that their token values are uniformly distributed between 0 and C. The exact formula for Q is:

$$Q(c^*,W,N,C) = \frac{(N-1)!}{(W-1)!(N-W)!}\left(\frac{c^*}{C}\right)^{W-1}\left(\frac{C-c^*}{C}\right)^{N-W} \tag{2}$$

--For all experimental parameters with W < N, there is exactly one solution to (2) with $0<c^*<1$. This solution is "globally expectationally stable" (Palfrey and Rosenthal, 1991b), while the $c^* = 0$ equilibrium is unstable.

--If W = N, then $c^*=0$ is the unique globally expectationally stable equilibrium for the experimental values of C.

Summarizing, for all parameter values of the games reported in this paper, there is a unique stable Bayes-Nash equilibrium to the one-shot game.

THE REPEATED GAME

We analyze the repeated game as if it is an infinitely repeated game with a discount factor of .9. While the random stopping rule we used does not exactly correspond to this, it is very close. The fact that the probability of stopping was equal to 0 for the first 19 plays of the game simply means that any supergame payoffs that could be supported with a discount factor of .9 could also be supported if there is no discounting for the first 19 rounds, followed by a constant discount factor $\delta=0.9$. We assume everyone is risk neutral. This assumption was well supported, in terms of aggregate earnings, in the one-shot experiments.

Optimal Rotation Schemes.

First, we show that at $\delta=0.9$, it is possible to support the payoffs associated with an optimal rotation scheme. This is done by demonstrating that a player with a token value of C is better off spending rather than keeping in the current round, if spending leads to a perpetual continuation of an optimal rotation scheme, but keeping leads to the worst possible one-shot Bayes-Nash equilibrium in every future period. Such a punishment scheme is subgame perfect.

Case 1: $W = 1 \quad N = 3$

In the optimal rotation scheme, an individual will have to contribute every third game, and will give up a token value worth an expected $\frac{C}{2}$. In the other two games, a payoff of 1 will be earned. If a player contributes when it is his turn to do so, his current payoff will be $1 - c_i$, which is always greater than or equal to $1 - C$. If the player does not contribute, then his payoff will be equal to 0 in the current game, and the value of the worst one-shot Bayesian equilibrium in every future period.

Therefore, the value of contributing if $c_i = C$ equals:

$$V_C = 1 - C + \delta\left[\frac{1}{1-\delta^3}\right]\left(1+\delta+\delta^2\left(1-\frac{C}{2}\right)\right)$$

and the value of not contributing equals:

$$V_{NC} = 0 + \delta\left[\frac{1}{1-\delta^3}\right](1+\delta+\delta^2)V_0$$

where V_0 is the expected value of the worst (one-shot) Bayes-Nash equilibrium.

From (1) and (2), we get

$$V_0 = \frac{c*}{C}\left(1-\frac{c*}{2}\right) + \left(1-\frac{c*}{C}\right)(1-c*)$$

where $c*$ is the unique solution to (1) and (2) between 0 and 1. In our experiments, $C = 1.5$ and $\delta = .9$. Substitution into the formulas for V_c and V_{NC} shows $V_C > V_{NC}$.

The analysis is comparable, but simpler, when $W > 1$, since the value V_{NC}, of deviating from the equilibrium path equals 0.

Case 2: W = 2 N = 3

In this case, alternation requires each member of the group to spend two thirds of the time and keep one third of the time. So, a typical pattern of contribution will go: SSKSSKSSK...Therefore, V_c is smallest in a round where i's current token value equals C, and he must spend the next round as well. Simple algebra gives

$$V_c = \frac{1}{1-\delta^3}\left[1-\frac{C}{2}+\delta\left(1-\frac{C}{2}\right)+\delta^2\right] - \frac{C}{2}$$

$$= \frac{1}{1-\delta^3}\left[(1+\delta+\delta^2)-\frac{C}{2}(1+\delta+1-\delta^3)\right] > 0 \iff C < 2\left[\frac{1+\delta+\delta^2}{1+\delta+1-\delta^3}\right]$$

For $\delta = .9$, this reduces to the condition:

$$C < \frac{5.41}{2.17}$$

which is satisfied for the experimental values of C.

Case 3: W = 3 N = 3

As before, we evaluate the value of contributing if the token value is C.

$$V_c = 1 - C + \frac{\delta}{1-\delta}\left(1-\frac{C}{2}\right)$$

Thus $V_c > 0 \iff \delta > 2 - \frac{2}{C}$. This holds for the experimental values of C.

Case 4: W = 2, N = 4.

Here, individuals must contribute every other round. This means:

$$V_c = \frac{1}{1-\delta^2}\left(1-\frac{C}{2}+\delta\right) - \frac{C}{2} \geq 0$$

$$\Rightarrow C \leq \frac{2(1+\delta)}{2-\delta^2}$$

and substituting δ=0.9 gives:

$$C \leq \frac{3.8}{1.19}$$

This is satisfied for the experimental value, C = 2.25.

Optimal Asymmetric Contribution Rates Without Rotation

The optimal rotation schemes described above have the property that, even though the payoffs to the players are not perfectly symmetric (because the first contributor earns a little less on average than the second and third contributors), they are very nearly equal. In fact, if a die could be rolled to determine the rotation order, then they would be "fair" in the sense of giving all members of the group the same ex ante payoff.

The expected group payoff from rotation can be equaled if the same subset of W members of the group always spends in every round, effectively dividing the group into two subgroups, which we call "activists" and "free-riders". Of course, the activists in the group would earn much less than the free riders in the group.[14] In the W=2, N=3 game, an equilibrium with two activists, who both always spend, and one free rider, can be supported for the experimental parameter C=1.5, but not for C=2.25. In the W=1, N=3 game, one cannot support an equilibrium that involves one activist and two free riders, but an infinite alternation between two activist players with the third player always being a free rider can be supported in the supergame for $\delta=.9$. A similar solution can be supported in the W=2, N=4, C=2.25 game with 1 free rider and 3 activists.

One may ask the more general question of what the optimal asymmetric arrangement is, that is also stationary, in the sense that each player's strategy along the equilibrium path depends only on his current token value, and not upon the current time period. It is fairly easy to show that such strategies always take the form of a vector of cutpoints, one for each of the players. For the W=1, N=3 games, the W=2, N=4 games, and the W=2, N=3, C=1.5 games, the optimal asymmetric solution is the one just given. For the remaining parameter condition (W=2, N=3, C=2.25), we do not derive the optimal asymmetric cutpoints, but conjecture that for these treatments, there is no equilibrium asymmetric arrangement that does better than the optimal symmetric cutpoints discussed in the next section.

Since with the exception of the last condition, these optimal asymmetric vectors of cutpoints divide the members of the group into one set of players who always free ride and a second set of activists, the equilibrium strategies can be monitored under both the "reveal" and the "no-reveal" treatments. This is true of the rotation schemes, as well, since the trigger strategies only require knowledge of the past moves of the other players in the group.

Optimal Symmetric Contribution Rates Without Rotation

Finally, we show that in the "reveal" treatment, another kind of cooperative solution can be supported with repeated play.[15] For any given set of parameters, one can compute a unique value c**, called the efficient group cutpoint, that represents the best symmetric cutpoint rule for the group as a whole. In other words, if everyone uses the cutpoint c**, then, ex ante, each individual earns more than he would under any other rule that assigns a common cutpoint to all players.

In the case of W=N this obviously coincides with the optimal rotation scheme. However, when W < N, this rule will yield lower total expected payoffs to the group than the optimal rotation scheme[16], but higher payoffs than the best of the Bayes-Nash equilibria to the one-shot game. These arrangements also generally lead to lower payoff than the optimal asymmetric cutpoint rules. Table 2 summarizes these optimal symmetric cutpoints for all the experimental parameters. It is an easy exercise to show that, for any of our experimental parameters, they are supportable as supergame equilibria for $\delta=.9$, if the punishment phase involves reverting to the worst one-shot Bayes-Nash equilibrium.

[Table 2 about here.]

However, if the cutpoint arrangement calls for a cutpoint lying strictly between 0 and C, then a monitoring issue arises. When a player fails to contribute in a round, the other members of the group need to know his token value in order to ascertain whether the arrangement has been violated, since punishment is appropriate only for "keepers" with token values below the cutpoint. This is not a problem in the "reveal treatment, since players are provided with the necessary information. But in the "no-reveal" treatment, the simple trigger strategies we have been considering will not work. Thus, the analysis here only holds for the "reveal" treatment. We have not worked out the optimal asymmetric contribution rates that can be supported when token values are not revealed after each round. The one limited observation we can make is that the "no reveal" repeated games typically cannot support group payoffs as high as the "reveal" treatment.

4. HYPOTHESES

PAYOFFS

The central issue addressed in this paper is whether the play of a supergame leads to more efficient outcomes than in the one-shot play of the stage game. Accordingly, the first set of hypotheses focuses on differences between (normalized) per-subject earnings in the one-shot game sessions and the repeated sessions.

H1. Earnings are higher in the repeated sessions than in the one-shot sessions.

Second, we have a hypothesis, suggested in the previous section, relative to the secondary reveal/ no reveal treatment. This hypothesis is motivated by the idea that a greater ability for players to monitor the strategies of the other players in the group will lead to more cooperation:

H2. Earnings are higher in reveal treatments than in no reveal treatments, holding parameters constant.

Finally, there are several hypotheses about how well the theoretical models predict. The equilibria computed in the last section lead to some very specific hypotheses about earnings. For the one-shot games there is a single hypothesis:

H3. Earnings in the one-shot sessions are consistent with the predictions of the one-shot Bayesian equilibrium model.

For the repeated-game sessions, there are several natural possibilities to focus on. One is that repetition makes no difference, and that the one-shot Bayesian equilibrium predicts the experimental data in the repeated sessions, just as it is predicted to do in the one-shot sessions. Two other possibilities are indicated from the repeated game analysis in the last section: (1) rotation; and (2) optimal symmetric cutpoints. These are summarized in hypotheses 4-6.

H4. Earnings in the repeated sessions are consistent with the predictions of the one-shot Bayesian equilibrium model.

H5. Earnings in the repeated sessions are consistent with the predictions of the optimal rotation equilibrium.

H6. Earnings in the repeated sessions are consistent with the predictions of the optimal symmetric cutpoint equilibrium.

SPENDING

At first glance this may seem like the same question as efficiency, but it is not. While our theoretical model leads directly to 6 hypotheses that parallel H1-H6. with the word "earnings" replaced by "spending rates", tests of the spending hypotheses are not equivalent to tests of earning hypotheses. In our public goods game, an increase in spending that does not result in the threshold being met decreases efficiency as does an increase in spending that results in contribution above the threshold level. Even when the threshold is met exactly, if W<N, efficiency is greater if the low token value individuals are the spenders.

A well-known theoretical feature of the supergame is that there exist an infinity of equilibria that can support a large range of payoffs to the players. A glimpse of this variety is evident in the previous section. On the other hand, in the one-shot game there is a unique prediction of a stable equilibrium. Thus an additional hypothesis concerns dispersion of spending rates:

H7. Spending rates are more dispersed in the repeated sessions than in the one-shot sessions.

This hypothesis is motivated by at least three theoretical considerations:

(1) Because there are multiple equilibria to the repeated game, group and individual earnings may be dispersed, even if all groups exhibit equilibrium behavior, since different groups may select different equilibria.

(2) Individuals in a group may adopt strategies that are part of some equilibrium, but not all individuals choose the same equilibrium.

(3) The search for coordination in the repeated game leads to substantial out of equilibrium behavior, with different players trying different strategies.

(4) There exist asymmetric equilibria in the supergame.

STRATEGIES

There are several alternative hypotheses about the kinds of strategies individuals might be adopting, and about how these strategies might differ between the one-shot and repeated sessions. With respect to the one-shot sessions, three predictions follow immediately from the Bayesian equilibrium model and these are stated (in order of how strong the prediction is) as hypotheses 8-11.

H8. (DOMINANCE) Individuals never spend tokens if the token value exceeds 1.

H9. (CUTPOINTS) Individuals adopt cutpoint strategies, which may vary across parametric treatments.

H10. (SYMMETRY) All individuals adopt identical cutpoints within a given parametric treatment.

H11. (EQUILIBRIUM) Individuals adopt the cutpoint strategies equal to the ones predicted by the Bayesian equilibrium model.

Strategies in the repeated game are potentially more complicated because the strategy space is richer. Also, the dominance hypothesis is less compelling, since spending a token with a value greater than 1 is no longer a dominated action. For example rotation schemes will have players contributing when their token values exceed 1.

Similarly, the arguments for cutpoints are less compelling as well. In the one-shot game, best responses to *any* strategies by the other members of your group are always in the form of a cutpoint strategy. This is not true in the repeated games, with the obvious examples being the use of rotation schemes and punishment schemes as part of an equilibrium strategy. Nevertheless, such strategies are compelling partly because they constitute very simple decision rules.

For the same reasons that were given for Hypothesis 7, we expect more variation in cutpoint strategies in the repeated sessions. Finally, since there are so many possible equilibria besides the one-shot Bayesian equilibrium, we expect the one shot Bayesian equilibrium model to perform less well with the repeated-game data than with the one-shot data. Thus, compared to the one-shot games, we expect Hypotheses 8-11 to have less support in the repeated games.

Finally, there are some hypotheses about individual behavior that are specific to the repeated sessions. Do subjects adopt rotation strategies? Do some subjects act as complete "free riders", as the optimal asymmetric solution sometimes predicts? Are some subjects "activists"? Do subjects adopt cutpoints equal to the optimal symmetric cutpoint? These final hypotheses are stated below as hypotheses 12-14.

H12. (ROTATION) Subjects follow rotation strategies.

H13. (ACTIVISTS) In the K=1, N=3 and the K=2, N=3, C=1.5 parametric treatments, subjects separate into two groups: activists and free riders.

H14. (OPTIMAL SYMMETRIC MODEL) Cutpoints in the repeated game sessions are closer to the optimal symmetric cutpoint than to the one-shot Bayesian equilibrium cutpoint.

5. DATA

Efficiency: Does Repetition Help?

In order to compare efficiencies across parametric treatments, we convert the actual earnings by subjects into an efficiency index. This is done by subtracting the endowments (token values) from earnings paid at the end of the session, and then dividing these net earnings by the (net) earnings that would have resulted (conditioning on the actual token value draws) if everyone had used the group-optimal symmetric cutpoint.[17] This results in a normalization that will equal 0 if no one ever contributes, and will equal 1 if everyone contributes using the group-optimal symmetric cutpoint strategy.

The averages are taken over rounds 6 to 20. The first 5 rounds were excluded to control for inexperience and learning. Rounds 21 and higher in the repeated-game treatments were excluded to preserve comparability between the one-shot and repeated experiments.[18]

To address H1, that earnings are higher in the repeated-game sessions, we aggregate all earnings within a parameter treatment. A comparison of the predictions of the different theoretical models is summarized in Table 2. The observed efficiencies, comparing one-shot games and repeated-games are reported in Table 3.

[Table 3 about here.]

Except for the W=1, N=3 game, the results clearly show that repetition helps. This is especially strong in the W=3, N=3, C=1.0 game and in the W=2, N=3, C=1.5 game. The treatment differences for W=1, N=3 are very small and statistically insignificant but in the wrong direction, in that the one-shot sessions were more efficient.

There is a plausible ex post rationale to explain why this might have happened. This game is the most competitive, in the sense that the gains from cooperation are the least. This fact is reflected in the small difference between predicted per capita earnings in the one-shot Bayesian equilibrium (1.35) and per capita earnings with a group-optimal cutpoint (1.44).[19] With so little separation between the predictions for the cooperative vs. noncooperative outcomes, it is not surprising that we measured essentially no difference between the one-shot and repeated-play efficiencies in the W=1, N=3 game. The differences are greater in magnitude and in the right direction for all of the other parametric treatments. In fact, all these other differences are significant at the .01 level.

The hypothesis **H2** is tested by comparing the efficiency measures for the reveal and the nonreveal treatments as shown at the bottom of Table 3. Revelation of token values after each play raised efficiency in all four parameter treatments where we ran reveal as well as non-reveal sessions. But one feature of these results is perplexing. Revelation makes its biggest impact in 3 of 3 games. But, as explained above, theory predicts that monitoring should have no effect in these games.

A more detailed examination of the 3 of 3, C=1.0 games suggests that the magnitude of the reveal effect is not large. We ran (see Table 1) six sessions with these parameters. Four of these sessions occurred in two triple sessions. In the two 8/3/89 sessions, earnings were actually slightly less in the reveal treatment. In the two 10/29/91 sessions, earnings were virtually identical in the two treatments. So the difference between Reveal and No Reveal rests solely on the 6/11/91 session, in which all three groups coordinated fully from round three onward, having much higher earnings than the 9/4/91 session, when there was a very low degree of coordination. Thus it is possible that there cohort effects may be magnifying the reveal effect. On the other hand, there is some support for the increased efficiency from token revelation in the data for the W=N=3, C=1.5 experiments. The two reveal sessions were on 6/11/91. Subjects in these sessions not only did better than the subjects in the No Reveal setting of 9/4/91 but also bettered the 10/29/91 subjects who, as just reported above, achieved a high degree of coordination in the unanimity game with C=1.

Data that bear on hypotheses **H3-H6** are summarized in Figures 1 and 2 and Table 4. On the horizontal axis of Figure 1, we plot, for the random group treatment, the theoretical prediction of earnings, aggregated for a session, and using the actual token values drawn in the session. Each of the 22 sessions are plotted twice, once for the one-shot cutpoint prediction and once for the group optimum cutpoint prediction. If a theory were precisely correct, at the aggregate level, all the points would lie on the "Predicted=Actual" (45^{o}) line. To grasp the scale of the figure, the normalized value of C varies from 1 to 2.25 across treatments, so average endowments vary from 0.5 to 1.125. Thus an increase in earnings above endowments of 0.5 is quite substantial.

[Figures 1 and 2 about here]

The Bayesian one-shot predictions cluster quite closely about the 45^{o} line. The Root Mean Square Error (RMSE) of the theoretical predictions is only 0.09. (See Table 4.) The average deviation of the theoretical prediction from actual earnings is only -0.02, showing that, over all the sessions, there is only a slight tendency to under predict earnings. In contrast, the optimal symmetric cutpoints are, as expected, much worse for random groups. The average deviation is 0.261, showing that the cooperative model overpredicts earnings substantially. Results for the Rotation model are slightly worse (For our parameters there is no difference in predicted earnings between the Optimal and Rotation predictions for W=N=3 and only slight differences for some other parameters.)

[Table 4 about here]

Figure 2 reports the repeated games data in a similar way. With respect to fixed groups, the one-shot model no longer provides a tight match to the data. The RMSE more than doubles. Underprediction is more substantial; the average deviation increases to -0.13. As can be seen in Figure 2, the failure of the one-shot model here is largely the reflection of three sessions where the theoretical prediction is zero increase over the endowment but where the subjects actually gained nearly 0.5. These all represent W=N=3, C=1 experiments where all groups perfectly or nearly perfectly coordinated in Rounds 6 through 20. For the other parameters, the one-shot model is still reasonably accurate.

Indeed, the optimal symmetric cutpoint model does not fit the data better than the one-shot model. It continues to overpredict earnings, although by less than with random groups, but the RMSE is in fact slightly higher than that for the one-shot model. Again, rotation is the worst of the three models. These results again indicate that repeated play increases earnings, but only by a small fraction of the theoretically possible gain. Moreover, while, game theory is a highly accurate model of aggregate behavior in one-shot games, no game theoretic model adequately captures repeated play.

Repeated play might produce results closer to the optimum predictions if the Reveal treatment were used in all settings. (But see the caveat in the discussion above.) When we regress actual earnings against those predicted with the optimal cutpoint, a constant, and a dummy variable for Reveal, we find that earnings are, certeris paribus, 0.090 normalized units higher in the Reveal treatment (t-statistic = 1.879, p-value =0.03). But, as can be seen visually in Figure 2, even if subjects do better in Reveal treatments, they still tend to earn less than would be available with the optimal symmetric cutpoint.

Spending: Does Repetition Lead to More Contribution?

For all 6 of our parameters sets, repetition clearly leads to more contribution, as shown in Tables 3 and 5. Five of the six sets show statistically significant differences over rounds 6 to 20. The data also support **H7**, that spending rates also tend to be more dispersed with repeated play, particularly for W=N=3 sessions (see Table 5), where appropriate F-tests are highly significant.

In the W=3, N=3, C=1.0 treatment, the median subject always contributed. Most groups in that treatment achieved nearly perfect cooperation. Some groups failed to coordinate; subjects in those groups typically had contribution rates near zero. On the other hand, the "no contribution" equilibrium was never reached. Although contribution tended to diminish over time in groups that had not coordinated, sporadic contributions continued. These sporadic contributions lowered efficiency, particularly in the 8 groups that always or nearly always failed to obtain three contributions in a round.

For W=3, N=3, C=1.5, most subjects contributed very little, but a few subjects had very high contribution rates. While some groups coordinated sporadically, no group sustained coordination throughout the experiment, and some groups did fully or nearly lock in on everyone always keeping.

Table 5 shows that aggregate spending rates for some parameters are quite distant from either the Bayesian or optimal symmetric predictions. For example, for W=3, N=3, C=1.5 there is substantial spending in one shot scenarios where the Bayes prediction is for no spending (zero cutpoint). However, the spenders tend to have low token values. Thus, these individual deviations in spending lead to relatively small deviations from the predicted level of aggregate earnings.

[Table 5 about here]

Strategies

Dominance

Our first hypothesis concerning strategies (H8) stated that dominated strategies would not be used in one shot situations. This, as shown in Table 6, is supported. A dominated strategy was used only in under 2 percent of the possible opportunities. In contrast, dominated strategies should be used often in rotation schemes and even in the optimal symmetric cutpoint strategies. While they were used more in repeated play than in random groups, they were used much less, if at all, than might be expected in these cooperative schemes (see the statistical tests in Table 6).

Indeed, the only cases where we observe full coordination of some groups is where both unanimity is required and all endowments are below the value of the public good. Thus, achieving efficiency apparently depends not just on having a unique efficient equilibrium. It also depends on not having a temptation to defect in the current period.

[Table 6 about here]

When the endowments can be very large, the C=2.25 case, we observe almost no contribution when subjects have endowments greater than the value of the public good. Indeed there is little difference compared with random groups.

For more moderate endowments, C=1.5, contributions of tokens over the endowment level is greater than with random groups (again see the tests in Table 6), but remains very small. In the 1 of 3 case, there are 5 contributions for fixed groups, whereas none should occur with the optimal symmetric cutpoint. The five contributions were made by just 2 of the subjects, one of whom was someone who tried a rotation strategy, contributing on rounds 1, 4, 7, etc.

In the 2 of 3 and 3 of 3 experiments for C=1.5, there should be substantial contribution from subjects with token values above the public good value. While far more contributions occur here than with random groups, the level is still less than one-fourth of the theoretical level.

The Use of Cutpoint Strategies

The most striking feature of individual behavior is the extensive use of cutpoint decision rules by subjects to decide when to spend and when to keep. Table 7 shows how pervasive these strategies are, lending strong support for H9. That table reports the classification errors that result if one chooses, for each individual, the cutpoint that would minimize the total number of classification errors in that group. We also report results if we impose the restriction that all individuals in a session use the same cutpoint rule (the rows marked "session" in Table 7). For the repeated sessions we report classification errors under the weaker restriction that all members of the same group use the same cutpoint rule (the rows marked "group").

[Table 7 about here]

On average, we are able to assign a cutpoint to an individual and correctly classify nearly 95% of the spending decisions.[20] The low standard deviations for the averages show that there are few individuals who are classified very poorly.

In the repeated game treatment, we classify groups nearly as well as we do individuals, even though we buy back two or three degrees of freedom. This suggests that members of groups tend to adopt similar cutpoint strategies. Except in the unanimity (W=3,N=3) games, there is also little deterioration in fit if we assign a common cutpoint to the session. This suggests that group effects are weak except in unanimity games.

For the W=3, N=3 games, results deteriorate substantially if a common cutpoint is assigned to the session. This is an indication of strong group effects. Particularly with C=1, groups bifurcate into groups that almost always produce the good and groups that almost never produce the good.

The classification analysis is similar in the one-shot treatments except that there is a less marked deterioration in fit when the estimated cutpoints are constrained to be the same for all subjects in a session. An exception to this is the W=3, N=3, C=1 condition for random groups, where only 74% of spending decisions can be accounted for using a common cutpoint for each session.

In this case, in addition to the stable zero equilibrium, there is an unstable equilibrium where subjects all contribute. The 41 percent contribution rate suggests that some subjects made (unrewarded) efforts to obtain the high contribution equilibrium.[21] Our results suggest that repetition allows some, but not all groups, to arrive at this equilibrium. On the other hand, members of groups that fail do not "waste" their tokens as frequently as participants in random group sessions.

The comparison of the constrained and unconstrained classification analysis provides a partial answer to H10 (symmetry). While there is apparently some variation in the specific cutpoint strategies individuals pursue, this variation is not large in magnitude, except in the unanimity games. This is also reflected in the standard deviation of the estimated cutpoints reported in Table 3. As hypothesized, there is more heterogeneity in the repeated-game condition than in the one-shot condition.

The hypothesis (H11) that cutpoints are well-predicted by the Bayesian Nash equilibrium is reasonably well supported in the one-shot condition, with one caveat. For the parametric treatments in this paper average estimated cutpoints are systematically above the Bayesian equilibrium prediction. (See Table 3.) There are a number of possible explanations for this, which are explored in Palfrey and Rosenthal (1991b).

In the repeated-game condition, average estimated cutpoints are even higher. In Table 3 we present significance tests of the null hypothesis that the mean cutpoints are equal in the two conditions vs. the alternative that cutpoints are greater in the repeated condition than in the one shot. Four of the six parameter sets show a statistically significant difference. The small differences for the other two sets (in one case opposite to the theoretical prediction) are not significant. Overall, the estimated cutpoints in the repeated condition are closer to the levels represented by H14.

Finally, there is little support for the hypothesis that in the repeated game subjects divide into activists, who always contribute, and free-riders, who never contribute in either the W=1 condition or the W=2 conditions. In the W=1 and W=2 treatments, none of the subjects were activists in always contributing in rounds 6-20 and only 2 subjects were free-riders in that they never contributed. In fact only 12 of the 144 W=1 and W=2 subjects contributed two or fewer times in rounds 6-20. The two pure free riders did not succeed in "bluffing" the other members of their groups into high contribution rates. One free rider was in a 1 of 3 game, the other in a 2 of 4. The remaining subjects contributed at rates far below that expected in either the symmetric group optimal reduced game (1 of 2 or 2 of 3) and, a fortiori, what would be optimal for the entire group, including the free rider Thus there is little support for H13. To the extent that heterogeneity in cutpoints existed[22], it was not that extreme.

Rotation

As discussed previously, the rotation model finds no support in the analysis of efficiency. Here we investigate to what extent some individual subjects may be adopting rotation strategies. To check for the simplest forms of rotation by specific individuals, we looked for the following patterns for each subject in each experiment, with "..." indicating indefinite repetition, **S** indicating spending, and **K** indicating keeping.

W=3,N=3	W=1,N=3	W=2,N=3	W=2,N=4
S...	SKK...	SSK...	SSKK...
K...	KSK...	SKS...	SKSK...
	KKS...	KSS...	SKKS...
			KSSK...
			KSKS...
			KKSS...

For each subject, we computed the number of correct classifications of their spending decisions for each of the above patterns. We then ascribed to each subject the maximum classification accuracy over the relevant possible patterns.

We performed the analysis for all rounds, rounds 1-20, and rounds 6-20. We carried out the analysis of all patterns for each experiment because it is possible, especially with the rounds 6-20 analysis, that a high degree of classification can be achieved by fortuitous cycles of randomly drawn token values.

Table 8 reports, over all rounds and experiments, the results for the best fitting rotating scheme and the results for the marginals, and clearly shows that subjects did not use rotation strategies.[23] Only in the W=3, N=3, C=1.0 games, where rotation coincides with optimal symmetric cutpoint strategies, does the rotation model work. In that case, **S**... (a degenerate form of rotation) does quite well. In fact, the median classification percentage is 100.00, showing that subjects typically contribute from the beginning of the experiment to the end.

[Table 8 about here]

We also checked whether specific individuals rotated. We identified all individuals whose spending decisions could be correctly classified 80 percent or better with the rotation model. Using 80 percent is tilted toward finding rotation, since 20 percent error is quite high for the cutpoint analysis.

For the W=2 N=4 sessions, we found only 1 of 12 subjects close to rotating (classified at 87% in rounds 6-20 and 80% for all rounds). In the W=2 N=3 sessions, we found only 4 of 96 subjects at 80 % or better in rounds 6-20. Of these 4 subjects, only the perfect rotator mentioned earlier bettered 80 percent for all rounds. In the W=1, N=3 condition, we found only 5 of 36 at 80% or better in rounds 6-20. Only 1 of these 5 (again a perfect rotator) bettered 80 percent for all rounds. There was scant evidence for true rotation.

Behavior was more systematic in the W=3, N= 3 experiments, where no true rotation was required. With C = 1, 48 of 69 subjects contributed at least 80 percent of the time, and 10 of the remaining 21 contributed no more than 20 percent of the time over rounds 6-20. With C = 1.5, 22 of 42 never contributed more than 20 percent of the time, and 9 of the remaining 20 contributed at least in 80 percent of the rounds.

The two cases of perfect rotation we encountered, one in a W=1, N=3 game and the other in a W=2, N=3 game, in fact involved the same subject, subject #3 in the 25 July 1989 session. The first treatment in this triple session was the 1 of 3 game.

The entire history of the repeated game played by this subject's group is given in Figure 3. This game lasted for 21 rounds. The other two subjects essentially used a common cutpoint rule for a total of 2 errors in 42 decisions. The public good was produced in all but 6 of the 21 rounds. Thus the rotator did not have a strong incentive to deviate from rotation.

The next treatment of this triple session was W=2, N=3, C = 2.25. Subject #3 continued to rotate, contributing on rounds 3,4,6,7,9,10, etc. Rotation was reinforced by the fact that the good was produced frequently at the beginning of the experiment and that subject #3 drew low token values whenever it was "his turn" from rounds 18 onward. On the other hand, when the experiment ended the good had been produced in only 12 of 27 rounds.

[Figure 3 about here.]

The failure to produce the good frequently in the second session perhaps underlies subject #3 becoming a cutpoint player in the last game of this session, W=2, N=3, C=2.25. Using a cutpoint of 41 francs for #3 results in only 2 classification errors over all 61 rounds. Rotation requires contribution on two successive rounds. This happened only once, when 3's token values were very low.

6. Conclusion

Most of the theoretical hypotheses were supported to some degree in the data. Repetition leads to more cooperative behavior (more spending) and improves efficiency, and better monitoring appears to have a similar effect. Game theory can account for many of the qualitative features of the data, particularly in the one shot games. Subjects avoid dominated strategies and adhere to cutpoint decision rules that are, on average, very close to the ones predicted by Bayesian equilibrium.

On the other hand, subjects are unable to come even close to fully exploiting the opportunities for coordination and cooperation in the repeated games. The only environment where full efficiency is achieved is in the unanimity games which actually have a one-shot Bayesian equilibrium (albeit unstable) which yields full efficiency.

Thus, while the theory correctly predicts the qualitative effects of repetition, the observed magnitudes are much smaller than predicted. We conjecture that one reason for this is that the theory lacks a good dynamic model of how players might **reach** an efficient equilibrium in a repeated game. This void in the theory is particularly evident when one considers rotation schemes, which are non-stationary and require a specific timing structure. In order for such schemes to ever come about, players in a group will have to go through some sort of "groping about" process in the early plays of the game, reminiscent of the process modeled by Crawford and Haller (1990). Apparently even relatively simple rotation schemes are very difficult to support, and unlikely to emerge spontaneously without explicit coordination devices, such as direct or mediated communication. We found essentially no evidence at all for implicit coordination of this sort.

Unanimity games are the only ones where a groping stage is not necessary. However, even in unanimity games, many groups fail to coordinate, with success or failure largely determined in the early plays of the game. Again, this indicates a need to model the dynamic process of how players adjust their behavior over the course of the repeated plays of the game. Yet a striking feature of the behavior displayed in the experiments is the apparent use of heuristic cutpoints by the subjects. The high rates of classification success we obtained suggests that, after some initial rounds, individual behavior is highly stable. At the same time, group behavior is highly variable, particularly in the unanimity games with C=1 where groups bifurcated into fully cooperating groups and "failed" groups. From this perspective, the interesting dynamics may be packed into the initial experiences of groups.

REFERENCES

Andreoni, James (1988) "Why Free Ride?: Strategies and Learning in Public Goods Experiments," Journal of Public Economics, 37, 291-304.

Alger, Dan (1987) "Laboratory Tests of Equilibrium Predictions with Disequilibrium Date," Review of Economic Studies, 54, 105-145.

Cooper, Russell, Doug DeJong, Robert Forsythe, and Thomas Ross (1989) "Communication in the Battle of the Sexes Game," Rand Journal of Economics, 20, 568-587.

__________. (1991) "Cooperation without Reputation" mimeo, Boston University.

Crawford, Vincent and Hans Haller (1990) "Learning How to Cooperate: Optimal Play in Repeated Coordination Games," Econometrica, 58, 571-95.

Davis, Douglas D. and Charles A. Holt (1990) "Capacity Asymmetries, Market Power, and Mergers in Laboratory Markets with Posted Prices," working paper, Virginia Commonwealth University.

Friedman, James W. (1967) "An Experimental Study of Cooperative Oligopoly", Econometrica, 35, 379-397.

Friedman, James W. (1969) "On Experimental Research in Oligopoly," Review of Economic Studies, 36, 399-415.

Fudenberg, Drew and Jean Tirole (1991) Game Theory, MIT Press: Cambridge.

Gallo Jr., Philip S. and Charles G. McClintock (1965) "Cooperative and Competitive Behavior in Mixed-Motive Games," Journal of Conflict Resolution, 9, 68-78.

Isaac, R. Mark and James M. Walker (1991) "Costly Communication: An Experiment in a Nested Public Goods Problem," in Thomas Palfrey (ed.), Laboratory Research in Political Economy, Ann Arbor: University of Michigan Press, 1991.

Isaac, R. Mark, James M. Walker, and Susan H. Thomas (1984) "Divergent Evidence on Free Riding: An Experimental Examination of Possible Explanations," Public Choice, 43, 113-149.

Marwell, Gerald and Ruth E. Ames (1979) "Experiments on the Provision of Public Goods I: Resources, Interest, Group Size, and the Free Rider Problem," American Journal of Sociology, 84, 1335-1360.

__________. (1980) "Experiments on the Provision of Public Goods II: Provision Points, Stakes, Experience, and the Free Rider Problem," American Journal of Sociology, 85, 926-937.

__________. (1981) "Economists Free Ride, Does Anyone Else?," Journal of Public Economics, 15, 295-310.

McCabe, Kevin, Stephen Rassenti, and Vernon Smith (1991) "Cooperation and the

Repeat Interactions of Anonymous Pairings", mimeo.

McKelvey, Richard D. and Thomas R. Palfrey (1989) "An Experimental Study of the Centipede Game," Working Paper 732, California Institute of Technology.

Ostrom, Elinor (1990) Governing the Commons, Cambridge University Press: Cambridge.

Ostrom, Elinor and James Walker (1991) "Communications in a Common:Cooperation Without External Enforcement," in Thomas Palfrey (ed.) Laboratory Research in Political Economy, Ann Arbor: University of Michigan Press, 1991.

Palfrey, Thomas R. (1985) "Buyer Behavior and the Welfare Effects of Bundling by a Multiproduct Monopolist: A Laboratory Investigation" in Vernon Smith (ed.) Research in Experimental Economics, JAI Press, Vol. 3, 73-104.

Palfrey, Thomas R. and Howard Rosenthal (1988) "Private Incentives and Social Dilemmas: The Effects of Incomplete Information and Altruism," Journal of Public Economics, 28, 309-332.

__________________. (1991a) "Testing for Effects of Cheaptalk in a Public Goods Game with Private Information," Games and Economic Behavior, 3, 183-220.

__________________. (1991b) "Testing Game-Theoretic Models of Free Riding: New Evidence on Probability Bias and Learning," in Thomas Palfrey (ed.) Laboratory Research in Political Economy, Ann Arbor: University of Michigan Press, 1991.

Prisbrey, Jeffrey (1991) "An Experimental Analysis of Two-Person Reciprocity Games," mimeo, California Institute of Technology.

Rapoport, Anatol and C. Orwant (1962) "Experimental Games: A Review," Behavioral Science, 1-37.

Roth, Alvin E. and J. Keith Murnighan (1982) "The Role of Information in Bargaining: An Experimental Study," Econometrica, 50, 1123-1142.

Schneider, Friedrich and Werner Pommerehne (1981) "Free Riding and Collective Action: An Experiment in Public Microeconomics," Quarterly Journal of Economics, 96, 689-704.

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

Walker, James, Elinor Ostrom, and Roy Gardner (1990) "Rent Dissipation in a Limited Access Common-Pool Resource: Experimental Evidence," Journal of Environmental Economics and Management, 19, 203-211.

NOTES

----

* California Institute of Technology and Carnegie-Mellon University, respectively. The authors are thankful for the research support of the National Science Foundation through grants # SES-8718650 and # SES-9011828. The research assistance of Mark Fey, Jessica Goodfellow, and Jeff Prisbrey is gratefully acknowledged for their help in conducting the experiments. Sanjay Srivastava was instrumental in developing the computer network used for the experiments. Work on this paper proceeded while Rosenthal was a Fellow at the International Centre for Economic Research and a Fellow at the Center for Advanced Study in the Behavioral Science. He is grateful for financial support provided by National Science Foundation #BNS-8700864 during his stay at CASBS.

[1] See Fudenberg and Tirole, 1991, ch. 5, for an extensive bibliography

[2] This dates back to work in the 50's and 60's on repeated bimatrix games, which is surveyed in Rapoport and Orwant [1962] and Gallo and McClintock [1965]. Most of that research did not use financial incentives to induce the payoff matrices. More recently, there has been extensive work on finitely repeated games of cooperation in many settings, including public goods provision (Isaac and Walker [1991] and the articles they cite) oligopoly games (Friedman [1967,1969], Alger [1987], Davis and Holt [1990], common pool resource usage (Walker, Ostrom and Gardner [1990] and Ostrom and Walker [1991]) and other settings.

[3] Since we began running the experiments reported here, there has been a flurry of independent experimental research aimed at comparing one-shot play and repeated play in prisoner dilemma games and related environments. This includes work by Prisbrey (1991), Cooper et al. (1991), Andreoni and Miller (1991), and McCabe et al. (1991). The only earlier work we are aware of that was explicitly designed for a one-shot vs. repeated game comparison is by Andreoni [1988], who uses random matching schemes to emulate one-shot voluntary contribution games. His findings, based on a relatively small sample, are anomalous in that the one shot games he investigated led to greater cooperation than their finitely repeated counterparts.

[4] See, for example, Selten and Stoecker [1988] for evidence about the terminal period of a finitely repeated prisoners' dilemma game, or Isaac, Walker, and Thomas [1984] and subsequent work on finitely repeated voluntary contribution games.

[5] Most auction experiments have not used matching schemes.

[6] Ostrom (1990) emphasizes the enormous variety of ways groups find to overcome these free rider/coordination problems in naturally occurring versions of the related "commons" problem. These include examples of rotation schemes that she has documented, as well as other arrangements.

[7] Three of the nine-subject sessions were run as single sessions, rather than as triple sessions.

[8]Four of the sessions used 9 subjects. These exceptions are noted in Table 1.

[9]Appendix A contains a copy of the instructions for one of the sessions.

[10]No subject asked what we would have done if hours had passed and a 4 had not yet been rolled. Luckily the longest repeated game lasted for 67 periods. No triple session lasted more than two hours.

[11]We sometimes refer to i's token value as i's contribution cost.

[12]We use the terms "spend" and "contribute" interchangeably. In the actual experiment, we used the more neutral "spend/keep" terminology.

[13]See also Fudenberg and Tirole (1990 pp. 211-213).

[14]The analysis in this section does not apply in a natural way to the W=3, N=3 games, where the optimal solution is clearly for everyone to contribute.

[15]In the "no reveal" treatment, this kind of arrangement cannot be supported as a noncooperative equilibrium because of monitoring problems.

[16]In the W=1, N=3 games, the W=2, N=4 games, and the W=2, N=3, C=1.5 games, the optimal asymmetric solution gives the same expected group payoff as the rotation scheme.

[17]Table 2 shows the expected group earnings (net of endowment) for each of parameter sets in the one-shot equilibrium compared to the optimal symmetric cutpoint strategies.

[13]Inclusion of later rounds in the repeated-game sessions produces similar results. Table 5 gives an indication of the small magnitude of the differences depending on whether one uses all rounds or only rounds 6-20.

[19]See Table 2. Another comparison of the degree to which the game is competitive is simply to note that of all the games, the W=1, N=3 game is the only one in which there are no gains to preplay communication, or cheaptalk. In the other games, cheaptalk can lead to Pareto improving equilibria (Palfrey and Rosenthal [1991a]).

[20]Of course, more than 50% correct classification is guaranteed, even if subjects choose randomly. But Monte Carlo results in Palfrey and Rosenthal (1991a) indicate overwhelming rejection of chance models for the levels of classification reported here.'

[21]One subject, who had repeatedly spent his token, spontaneously remarked, while being paid, that he could not understand why the other subjects had been "so irrational".

[22]The F-tests in Table 3 indicate more variation in cutpoints for the repeated condition in three of the six parameter conditions.

[23]The results were similar for the rounds 1-20 and 6-20 analysis.The results are the same if one takes the classification percentage simply over all decisions rather than averaging the percentages for the subjects. The largest difference in these two percentages is 2.0%. The choice of the "best fitting" model was wholly insensitive to the use of either average or the median.

TABLE 1A

DESCRIPTION OF RANDOM GROUP EXPERIMENTS

Date	Site	w	N	Range	Cents	C	Subjs.	Rounds	Reveal	Sequence
7/13/88	CIT	1	3	90	0.5	1.5	12	30	No	-
2/15/89	CMU	1	3	90	0.3	1.5	12	20	No	2
2/16/89	CMU	1	3	90	0.3	1.5	12	20	No	2
7/31/89	CIT	1	3	90	0.3	1.5	12	20	No	2
8/8/89	CIT	1	3	90	0.3	1.5	12	18	Yes	2
1/22/87	CIT	2	3	90	1	1.5	9	20	No	-
12/3/87	CIT	2	3	90	1	1.5	9	20	No	-
12/20/87	CIT	2	3	90	1	1.5	9	20	No	-
2/15/89	CMU	2	3	90	0.3	1.5	12	20	No	1
2/16/89	CMU	2	3	90	0.3	1.5	12	20	No	3
7/31/89	CIT	2	3	90	0.3	1.5	12	20	No	1
8/8/89	CIT	2	3	90	0.3	1.5	12	20	Yes	1
4/27/89	CMU	2	3	90	0.7	2.25	12	20	No	1
5/2/89	CMU	2	3	90	0.7	2.25	12	20	No	3
8/3/89	CIT	3	3	204	0.1	.995	12	25	No	1
9/4/91	CIT	3	3	204	0.1	.995	12	20	No	1
7/21/88	CIT	3	3	90	0.5	1.5	12	30	No	-
2/15/89	CMU	3	3	90	0.3	1.5	12	20	No	3
2/16/89	CMU	3	3	90	0.3	1.5	12	20	No	1
7/26/89	CIT	2	4	204	0.1	2.22	12	17	No	1
7/31/89	CIT	2	4	90	0.3	2.25	12	20	No	3
8/8/89	CIT	2	4	90	0.3	2.25	12	20	Yes	3

TABLE 1B

DESCRIPTION OF REPEATED GROUP EXPERIMENTS

Date	w	N	Range	Cents	C	Subjs.	Rounds	Reveal	Sequence
7/25/89	1	3	204	0.1	1.5	12	21	No	1
8/9/89	1	3	90	0.3	1.5	12	20	Yes	2
7/27/89	1	3	204	0.1	1.5	12	47	Yes	3
7/25/89	2	3	204	0.1	1.5	12	27	No	2
7/27/89	2	3	204	0.1	1.5	12	34	Yes	1
7/27/89	2	3	204	0.1	1.5	12	29*	No	2
8/2/89	2	3	204	0.1	1.5	12	22	No	1
8/2/89	2	3	204	0.1	1.5	12	67*	Yes	2
8/2/89	2	3	204	0.1	1.5	12	20	No	3
8/9/89	2	3	90	0.3	1.5	12	20	Yes	1
7/25/89	2	3	204	0.1	2.22	12	61	No	3
8/3/89	3	3	204	0.1	.995	12	20	No	2
8/3/89	3	3	204	0.1	.995	12	31	Yes	3
6/11/91	3	3	204	0.1	.995	9	20	Yes	1
9/4/91	3	3	204	0.15	.995	12	30	No	2
10/29/91	3	3	204	0.12	.995	12	20	No	2
10/29/91	3	3	204	0.12	.995	12	29	Yes	3
6/11/91	3	3	204	0.1	1.5	9	27	Yes	2
6/11/91	3	3	204	0.1	1.5	9	20	Yes	3
9/4/91	3	3	204	0.15	1.5	12	20	No	3
10/29/91	3	3	204	0.12	1.5	12	34	No	1
8/9/89	2	4	90	0.3	2.25	12	20	No	3

Notes to Table 1.

Francs - This gives the upper end of the uniform (in integers) distribution of endowments in francs. This number and $\bar{c}$ can be used to calculate the franc value of the public benefit, B = Francs/$\bar{c}$.
Cents - The number of cents paid per franc earned in the experiment.
Subjects - The number of subjects in the experiment.
Rounds The number of (non-practice) rounds for the given (w, N, $\bar{c}$)parameters. After the first 20 rounds, a ten sided di was tossed to see if the game continued. The stopping probability was 0.1.
Reveal - When this parameter is "Yes", all token values were revealed after each round. Subjects could match the individual token values and the decisions of the other members of their group. Otherwise, token values were not revealed.
Sequence - Only one set of subjects was run on a given date. Subjects played three sets of parameters in sequence. The sequence shows the order in which the parameters were used. Some subjects were also used in experiments were groups were reformed at random on each round. This explains why only two sets of parameters are shown for the 9/4/91 and 8/3/89 experiments.
Site - All of the sessions in Table 1B were conducted at the CIT site. The sessions conducted on 9/4/91 used students enrolled at Pasadena Community College. Some of the 1989 CIT sessions included some subjects who were enrolled in a special summer program for high school students.
Other - The second session on 7/27/89 ended because of an unexpected computer crash following round 29. The second session on 8/2/89 had a computer crash in the eighth round, but was restarted and eventually lasted a total of 67 rounds. All three sessions on 8/9/89 were terminated immediately after round 20, without rolling a die.

TABLE 2.

THEORETICAL ANALYSIS RESULTS FOR THE EXPERIMENTAL PARAMETERS

TREATMENT PARAMETERS:							
	N	3	3	3	3	3	4
	K	1	2	2	3	3	2
	C	1.5	1.5	2.25	1	1.5	2.25
Bayes One Shot Symmetric Cutpoints		0.47	0.38	0.0	0.0	0.0	0.30
Group Optimal Symmetric Cutpoints		0.75	1.13	1.41	1.0	1.5	1.27
Earnings if Endowments Kept		0.75	0.75	1.13	0.5	0.75	1.13
Earnings for Bayes One Shot Symmetric		1.35	0.86	1.13	0.5	0.75	1.19
Earnings for Group Optimal Symmetric		1.44	1.17	1.20	1.0	1.0	1.48
Earnings for Optimal Rotation		1.5	1.25	1.75	1.0	1.0	2.13

Note to Table 2.
All computations with value of benefit (public good) equal to 1.0.

TABLE 3.

SUMMARY OF REPEATED VS. FIXED GROUP COMPARISONS

		TREATMENT					
PARAMETERS:	N	3	3	3	3	3	4
	K	1	2	2	3	3	2
	C	1.5	1.5	2.25	1	1.5	2.25
EQUILIBRIUM PREDICTIONS:	c^*_{os}	.47	.37	0	0	0	.3
	c^*_{∞}	.75	1.13	.94	1	1.5	1.27
	q^*_{os}	.31	.25	0	0	0	.13
	q^*_{∞}	.5	.75	.62	1	1	.57
OBSERVATIONS:	$\hat{c}_{os}$	.50	.55	.34	.42	.34	.60
	$\hat{c}_{\infty}$	.48	.68***	.56**	.73***	.60**	.60
	σ_{os}	.23	.22	.21	.24	.29	.27
	σ_{∞}	.25	.29††	.23	.34†	.54†††	.26
	$\hat{q}_{os}$	.28	.35	.18	.42	.25	.22
	$\hat{q}_{\infty}$	.34	.44	.28	.77	.41	.24
	n_{os}	60	75	24	24	36	36
	n_{∞}	36	84	12	69	42	12
	η_{os}	0.84	0.40	0.08	-0.18	-0.29	0.36
	η_{∞}	0.82	0.52	0.37	0.59	-0.04	0.48

Notes to Table 3.

The subscripts $_{os}$ and ∞ refer to the one-shot treatment and the repeated-group treatment, respectively. The reported values of $\hat{c}$ are computed as the average of the cutpoints that are estimated for each individual. Each of these was estimated by finding the hypothetical cutpoint that minimized the classification errors of a subject's contribution decisions. The reported values of $\hat{q}$ are the averages that appear in Table 4. The reported values of σ are computed as the standard deviation of the individual estimated cutpoints. The reported values of n are the number of individuals who participated in that treatment. The reported values of η, the efficiency measure, is the total payoff to the subjects in excess of their endowments, normalized so that earnings equal to the endowments have $\eta = 0$ and earnings predicted by the optimal symmetric cutpoint have $\eta = 1$.

Significance tests. See notes to Table 5 for parallel description and explanation.

TABLE 4

DEVIATION OF SESSION EARNINGS FROM THEORETICAL PREDICTIONS

	Random Groups			Fixed Groups		
	Model			Model		
	One Shot	Optimal	Rotate	One Shot	Optimal	Rotate
RMSE*	0.087	0.295	0.325	0.213	0.226	0.259
Ave. Dev.	-0.020	0.261	0.302	-0.129	0.194	0.231

Notes to Table 4.

* Root Mean Square Error.

All predictions calculated assuming subjects followed theoretical decision rules given their actual token values. Rotation predictions averaged over all possible rotation schemes. (For example, in 1 of 3 games, each player is allowed to be the contributor on Round 1.)

Units of analysis are sessions.

TABLE 5

CONTRIBUTION RATES

TREATMENT PARAMETERS:		N	3	3	3	3	3	4
		K	1	2	2	3	3	2
		C	1.5	1.5	2.25	1	1.5	2.25
Rounds 6-20								
Average		repeated	.34*	.44***	.28**	.77***	.41**	.24
		one-shot	.28‡	.35‡‡‡	.18‡	.42‡‡‡	.25‡‡‡	.22
Std. Dev.		repeated	.16	.16	.10	.34†	.37†††	.15
		one-shot	.13	.14	.11	.21	.16	.11
All Rounds								
Average		repeated	.33	.44***	.23	.76***	.38**	.26
		one-shot	.29‡	.35‡‡‡	.18‡	.43‡‡‡	.25‡‡‡	.24
Std. Dev.		repeated	.15	.14††	.10	.33††	.31†††	.12
		one-shot	.14	.10	.09	.19	.15	.10

Notes to Table 5.

Average refers to the percentage of tokens that were contributed for all subjects in a given treatment. The standard deviations (*Std. Dev.*) are computed by calculating a spending rate (proportion of tokens contributed) for each subject in a treatment, and then computing the standard deviation of these individual spending rates.

Significance tests.

The null hypothesis that the variance across subjects was equal in the two treatments vs. the alternative that the repeated variance was greater than the one-shot was tested by a standard F-test. The null hypothesis that the mean across subjects was equal in the two treatments vs. the alternative that the mean contribution rate was higher in repeated groups was tested by a standard t-test that allows for unequal variances. Using individual decisions rather than subjects as the units of observation, we also carried out the standard likelihood-ratio test of the null hypothesis that there was an equal contribution probability in the two populations. Because the effective number of observations was much higher in this test than in the t-test, p-values were lower for this test than for the t-test.

* One-tailed t-test (unequal variances) p-value ≤0.05.
** One-tailed t-test (unequal variances) p-value ≤0.01.
*** One-tailed t-test (unequal variances) p-value ≤0.001.
† F-test p-value ≤0.05.
†† F-test p-value ≤0.01.
††† F-test p-value ≤0.001.
‡ Likelihood-ratio test p-value ≤0.01
‡‡‡ Likelihood-ratio test p-value ≤0.001.

TABLE 6.

ENDOWMENTS AT LEAST EQUAL TO BENEFIT: CONTRIBUTIONS AND THEORY

Rounds 6-20

TREATMENT PARAMETERS:							
	N	3	3	3	3	3	4
	K	1	2	2	3	3	2
	C	1.5	1.5	2.25	1	1.5	2.25
Repeated Groups							
Actual Spends		5	23	1	n.a.	53	0
Theoretical Spends		0[a]	118***	30***	n.a.	222***	32***
Total Observations		160	420	92	n.a.	222	101
One-Shot Groups							
Actual Spends		3[n]	9†	4[n]	n.a.	9†††	5[n]
Theoretical Spends		0	0	0	n.a.	0	0
Total Observations		316	408	216	n.a.	188	309

Notes to Table 6.

The theoretical spends are calculated on the predictions of the symmetric cooperative cutpoint equilibrium for fixed groups and the actual token values used in the experiment. (For random groups, the non-cooperative prediction is always "keep" for endowments at least equal to the value of the benefit.)

*** The likelihood-ratio test that the actual spends were at a rate equal to the theoretical prediction was rejected at the .001 level or

a The test was not performed because the theoretical level is 0.

† The likelihood-ratio test of the hypothesis that the probability of spending a token at least equal to the benefit was equal in the repeated and one-shot treatments was rejected at the 0.05 level.

††† The hypothesis is rejected at the 0.001 level.

n The test was not carried out because the number of actual spends was small in both treatments.

TABLE 7

PERCENT CORRECTLY CLASSIFIED

Rounds 6-20 TREATMENT PARAMETERS:							
	N	3	3	3	3	3	4
	K	1	2	2	3	3	2
	C	1.5	1.5	2.25	1	1.5	2.25
Fixed Groups							
Average---Session		82	84	88	81	76	87
Group		86	87	90	94	91	90
Subject		92	92	93	96	97	94
Std. Dev.-Session		2	4	n.a.	18	8	n.a.
Group		5	8	6	8	9	4
Subject		8	11	7	8	6	7
Random Groups							
Average---Session		87	88	88	74	85	89
Subject		95	94	94	89	92	94
Std. Dev.-Session		4	4	4	0.4	6	2.
Subject		6	6	8	10	9	6

Note: "Session" refers to the classification percentages that result when a common cutpoint is calculated for the 12 (or 9) subjects in a given session for the parameters. "Group" refers to a common cutpoint for all subjects in a distinct group. "Subject" refers to a cutpoint for each subject. Cutpoints were chosen to minimize classification errors. In computing the standard deviations, the unbiased estimator formula (division by NOBS-1) was used. The standard deviation for sessions was computed with sessions as the units of observation. Accordingly,' "n.a." is shown for treatments with only a single session.

TABLE 8. RESULTS OF ROTATION ANALYSIS

	Parameters					
	3 of 3 C=1.0	3 of 3' C=1.5	2 of 3 C=1.5	2 of 3 C=2.25	2 of 4 C=2.25	1 of 3 C=1.5
Marginals						
Percent Correctly Classified, Averaged Over Subjects	76.1	61.8	56.4	76.6	74.1	66.8
Standard Deviation Over Subjects	33.1	31.3	13.9	10.5	12.0	15.2
Median Percent, Over Subjects	100.0	73.3	58.9	73.8	75.0	66.7
Best Fitting Rotation Model	3 of 3	1 of 3	1 of 3	1 of 3	1 of 3	1 of 3
Percent Correctly Classified, Averaged Over Subjects	76.1	60.8	61.7	63.2	66.7	63.2
Standard Deviation Over Subjects	33.1	11.4	6.8	4.8	6.5	8.2
Median Percent, Over Subjects	100.0	66.7	61.8	62.3	65.0	61.9
Number of Subjects	69	42	84	12	12	36

Note. The marginal model was Y for 3 of 3, C = 1.0, N for all other parameter sets. The best fitting rotation model was chosen from among the rotation models of 1 of 3, 2 of 3, 3 of 3, and 2 of 4.

Figure 1
Session Earnings, Random Groups

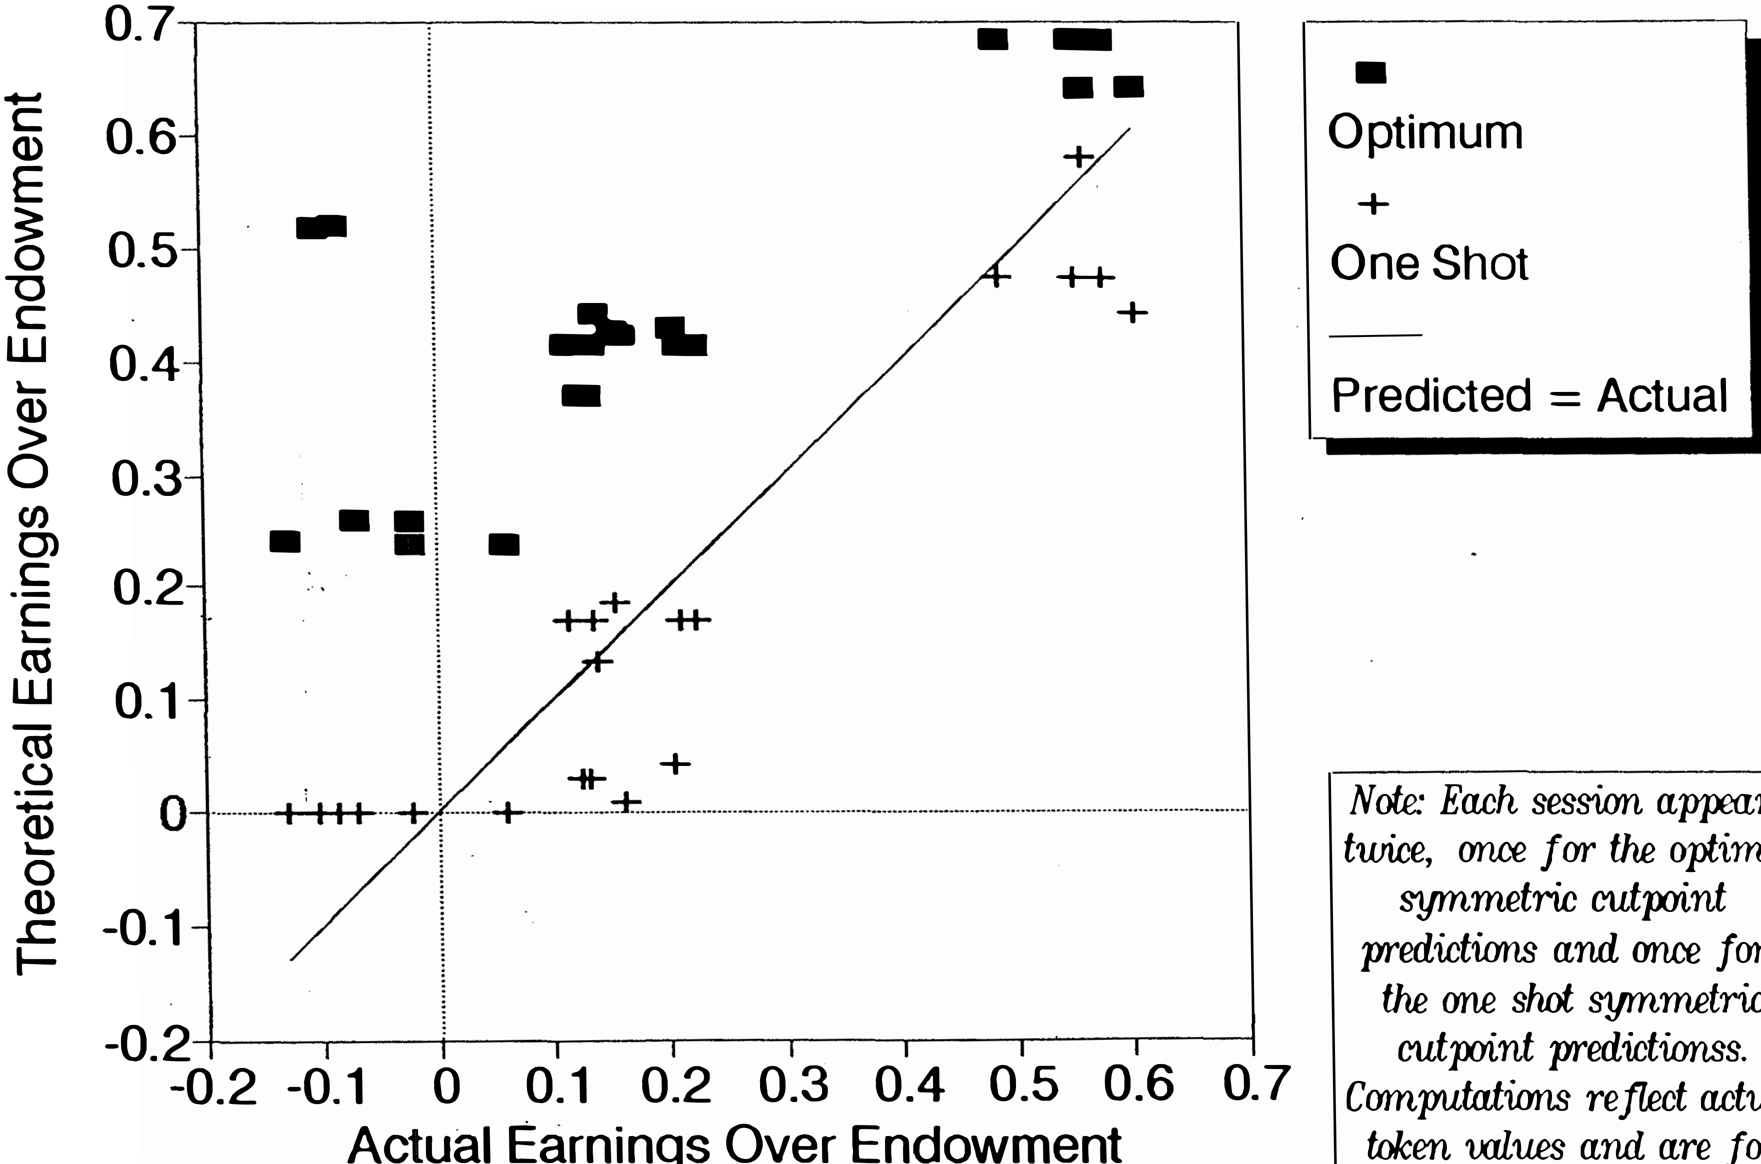

Note: Each session appears twice, once for the optimal symmetric cutpoint predictions and once for the one shot symmetric cutpoint predictionss. Computations reflect actual token values and are for ROUNDS 6–20.

Figure 2
Session Earnings, Fixed Groups

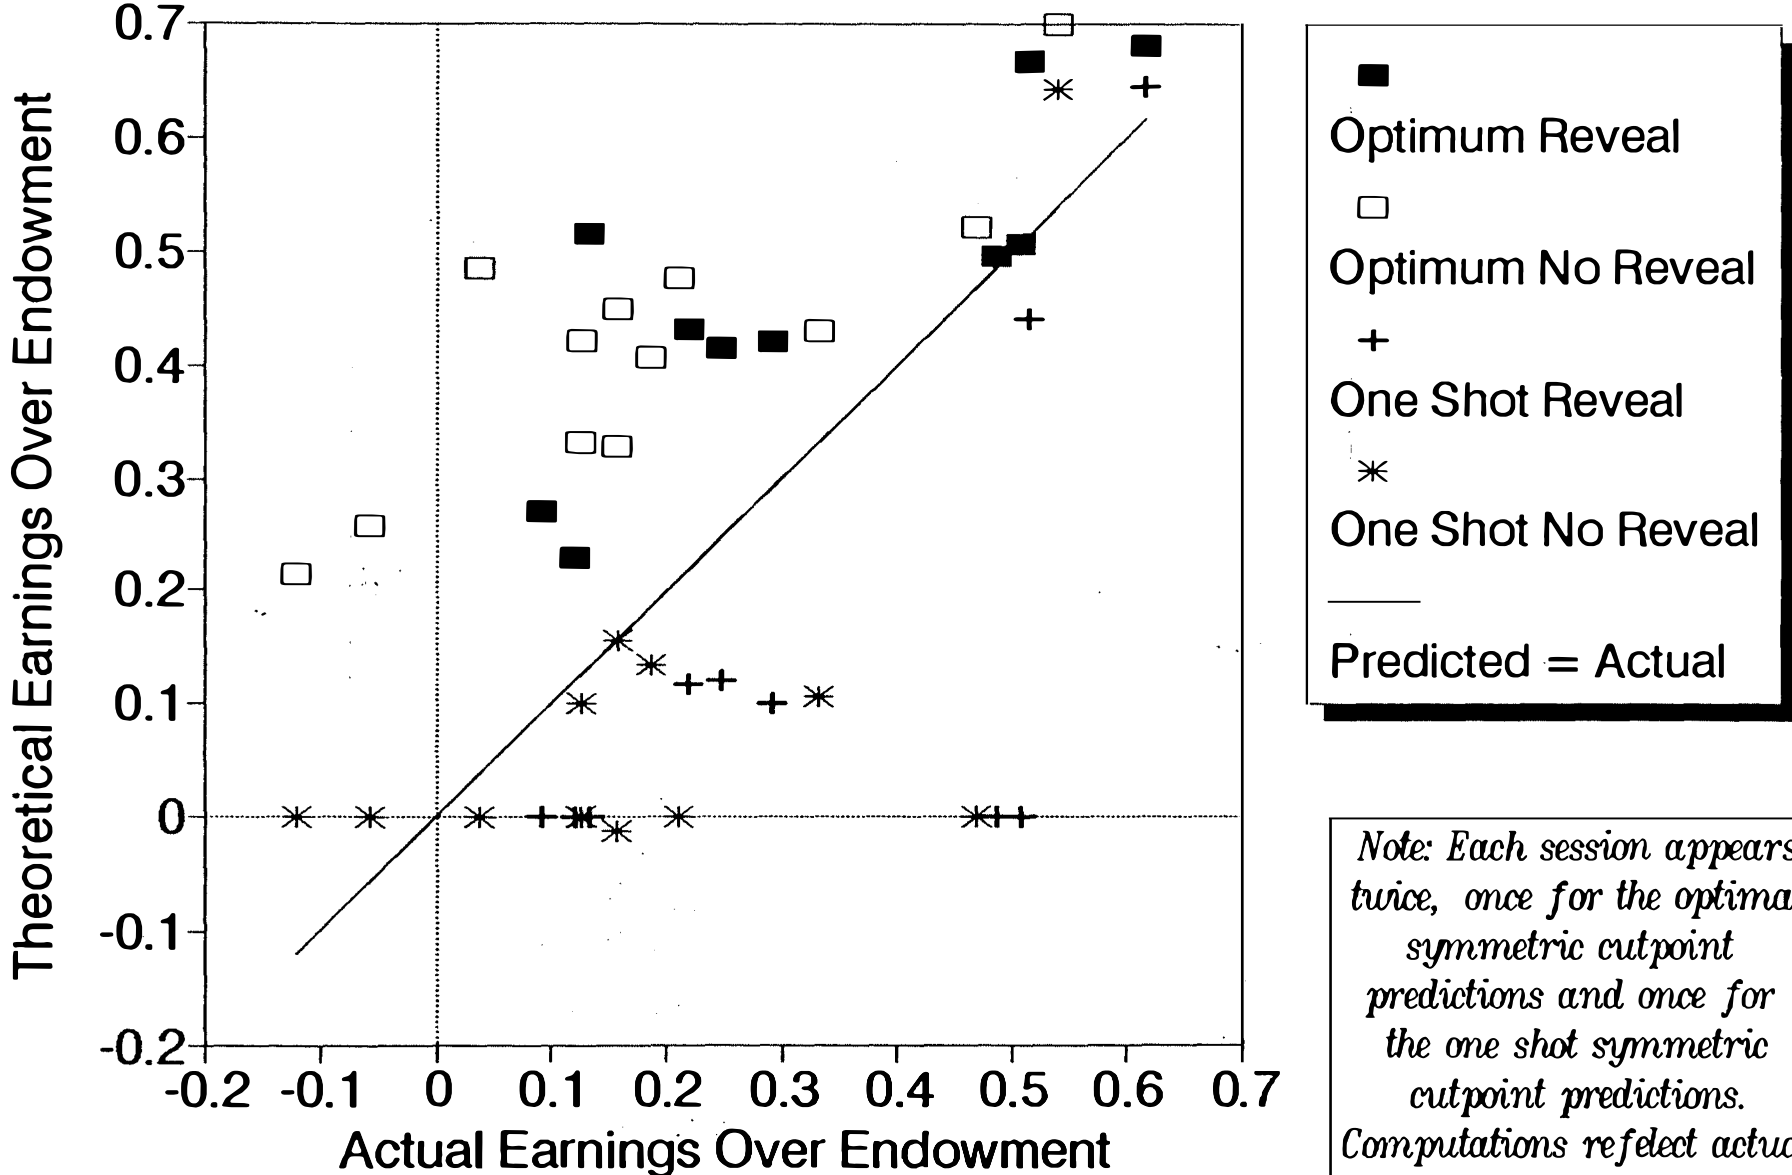

Note: Each session appears twice, once for the optimal symmetric cutpoint predictions and once for the one shot symmetric cutpoint predictions. Computations refelect actual token values and are for ROUNDS 6–20.

Figure 3

One of Three Repeated No Reveal Games
History of Group 1-1, CIT, 7/25/89

"E" Shows Classification Errors for Cutpoints
Benefit = 136
Max Cost = 204

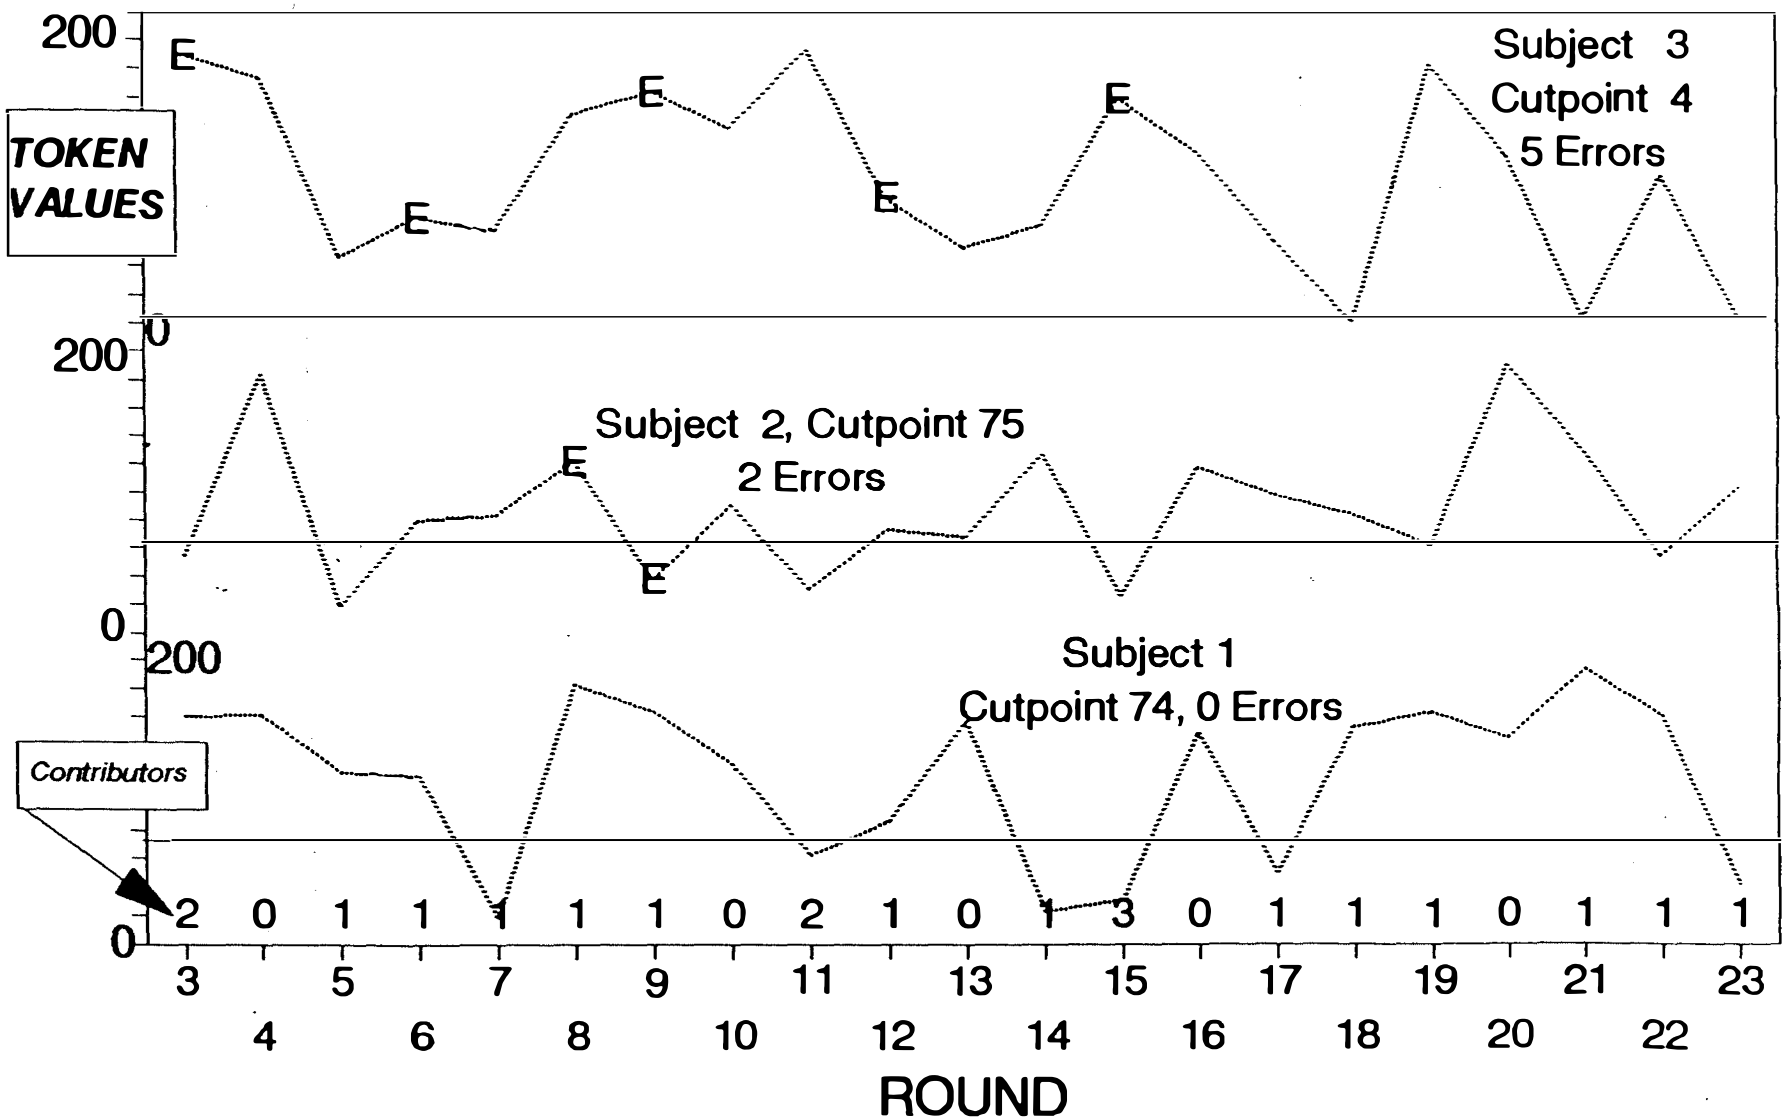

Appendix A. Sample Instructions (from 10/29/91 session)

Decision-Making Experiment

This is an experiment in decision making. You will be paid IN CASH at the end of the experiment. The amount of money you earn will depend upon the decisions you make and on the decisions other people make. It is important that you do not talk at all or otherwise attempt to communicate with the other subjects except according to the specific rules of the experiment. If you have a question, feel free to raise your hand. One of us will come over to where you are sitting and answer your question in private.

This session you are participating in is broken down into a sequence of three separate experiments.

EXPERIMENT 1:

This experiment is divided into many rounds, or periods. At the beginning of this experiment, you are assigned to a group with 2 other persons. You will not be told which of the other people in the room are in your group. The other members of your group will stay the same in every round of this experiment. Here are the rules that apply to every round of the experiment:

Each round in the experiment you have a single token to use in one of two ways: Option #1: Spend the token. Option #2: Keep the token.

The amount of money you earn in a round depends upon whether you keep or spend your token that round and how many others in your group spend their token. All money is denominated in FRANCS. At the end of the experiment, you will be paid $.12 for every 100 FRANCS you have accumulated during the course of all three experiments. Before each round begins, you will be told how many FRANCS your token is worth if you do not spend it. This amount, called your TOKEN VALUE, will change from round to round and will vary from person to person randomly. To be more specific, in each round, this amount is equally likely to be anywhere from 1 to 204 FRANCS. There is absolutely no systematic or intentional pattern to your token values or the token values of anyone else. The determination of token values across rounds and across people is entirely random. Therefore, everyone in your group will generally have different token values. Furthermore, these token values will change from period to period in a random way. You will be informed PRIVATELY what your new token value is at the beginning of each round and you are not permitted to tell anyone what this amount is.

Specific instructions:

At the start of each round you are told your token value for that round. After being told your token value, you must wait at least 10 seconds before making your decision to keep or spend. Your keyboard will be frozen for this period of time. After everyone in the room has made a decision, you are told which members of your group spent their token and what your earnings were for that round. You will never be told what the other members' token values were. The experiment will last at least 20 rounds. Beginning in round 21, we will roll a 10-sided die to determine whether to terminate the experiment or continue. We will terminate the experiment if and only if a 4 is rolled, otherwise we will conduct another round. This will continue until we roll a 4. Your total earnings in dollars will be your accumulation of FRANCS multiplied by the exchange rate of $0.12 = 100 FRANCS. [Write exchange rate on board.]

PAYOFFS

In each round, if all 3 members in your group choose to spend their token, every member in your group will each earn 136 FRANCS. Otherwise, the spenders in your group earn 0 and the nonspenders in your group earn their token value. Each group is completely independent: WHAT HAPPENS IN YOUR GROUP HAS NO EFFECT ON THE PAYOFFS TO MEMBERS OF THE OTHER GROUPS AND VICE VERSA. Therefore, in each round, you have exactly three possible earnings.

[write payoff matrix on board and explain]

Earnings Table

Your spending decision:	Number of Others Spending:	Your Earnings
Yes	2	136 FRANCS
Yes	0 or 1	0 FRANCS
No	0 or 1 or 2	Your Token Value

[Give quiz]
[Two practice rounds. Tell no one to press any keys unless instructed to do so. In round 1, make everyone spend. In round 2, make everyone keep. Explain the displays and the history screen.]

EXPERIMENT 2:

This is exactly the same as experiment 1 except for 2 things:

1. You are assigned to a new group with 2 other persons in the room. Neither of the other members in your new group were in your group in the last experiment. This new group assignment is fixed and will remain the same in every round of this new experiment.

2. If everyone in your group spends their token in a round, then everyone in your group earns 205 FRANCS in that round (instead of 136 in experiment 1).

[write new payoff matrix on board, and explain.]

[No quiz or practice rounds]

EXPERIMENT 3:

This is exactly the same as experiment 2 except for 2 things:

1. You are reassigned to a completely different group. The other two persons in your group now are different from the other two persons in your group in either experiment 1 or experiment 2. This new group assignment is fixed and will remain the same in every round of this new experiment.

2. At the end of each round, you are told what the token values of the other members of your group were in that round.

[No quiz or practice rounds]