

Cooperation without Reputation: Experimental Evidence from Prisoner's Dilemma Games*

RUSSELL COOPER†

Department of Economics, Boston University, Boston, Massachusetts 02215

DOUGLAS V. DEJONG AND ROBERT FORSYTHE

College of Business Administration, University of Iowa, Iowa City, Iowa 52242

AND

THOMAS W. ROSS

*Faculty of Commerce and Business Administration, University of British Columbia,
Vancouver, British Columbia, Canada V6T 1Z2*

Received September 16, 1992

This paper investigates cooperative play in prisoner's dilemma games by designing an experiment to evaluate the ability of two leading theories of observed cooperation: reputation building and altruism. We analyze both one-shot and finitely repeated games to gauge the importance of these theories. We conclude that neither altruism nor reputation building alone can explain our observations. The reputation model is inconsistent with play in both one-shots and finitely repeated games while the model with altruism is unable to explain observed play in the finitely repeated games. *Journal of Economic Literature* Classification Numbers: C72, C92. © 1996 Academic Press, Inc.

* We are grateful to the National Science Foundation, Price Waterhouse, the University of Iowa, and the Social Sciences and Humanities Research Council of Canada for financial support and to Zvi Eckstein, Charles Holt, Yong Gwan Kim, and the reviewers and editors for comments and suggestions. Comments by seminar participants at the 1990 meeting of the Economic Science Association, Harvard University, New York University, the Northwestern University Conference on "Cooperation and Dynamics in Games," the University of Pennsylvania Conference on "Behavior in Games," the University of Quebec at Montreal, and Queens University are appreciated.

† E-mail: rcooper@acs.bu.edu.

I. INTRODUCTION

There is an abundance of experimental evidence that subjects will often take actions that are not apparently in their best interest.¹ In many cases, these actions may be viewed as attempts to reach cooperative outcomes. While the empirical literature on cooperative play is quite extensive, the literature evaluating the predictive power of theories proposed to explain the observed cooperation is not. The purpose of this paper is to test two leading theories of cooperative play: altruism and reputation building.

Much attention has been paid in particular to cooperative play in the prisoner's dilemma (PD) game. In Game PD below, each player has a dominant strategy: she should fink regardless of her expectation regarding her rival's play and the outcome of (fink, fink) is therefore predicted. The important feature of this game is that this outcome is not Pareto-optimal. Indeed, the outcome in which both players cooperate Pareto-dominates (fink, fink) and maximizes joint payoffs.

		Player II	
		Fink	Cooperate
Player I	Fink	b, b	c, d
	Cooperate	d, c	a, a

Game PD: $c > a > b > d$

Experimental evidence on games of this form repeatedly reveals that some players cooperate.² While the design of these experiments has varied widely in terms of the frequency of play and the number of times a player faces the same opponent, the observed cooperative play is quite robust to these changes.³ In particular, cooperative play is observed in both repeated and one-shot environments.

One leading theoretical explanation of cooperative play is associated with Kreps *et al.* (1982), who argue that cooperation in finitely repeated

¹ To be precise, players take actions which appear to be against their self-interest *assuming* that the experimenter has controlled payoffs.

² See Dawes (1980) and Dawes and Thaler (1988) for an overview of this vast literature.

³ Cooperative play has also been observed in the play of coordination games (see, e.g., Cooper *et al.*, 1990, hereafter CDFR1), in ultimatum and dictator games (see, e.g., Forsythe *et al.*, 1994) and in the centipede game of McKelvey and Palfrey (1992).

PD games arises due to the presence of incomplete information regarding the true payoffs of a rival. The power of the Kreps *et al.* result is that a small belief that an opponent will cooperate is enough to support considerable cooperative play even when all players are purely self-interested. Common knowledge of rationality does not hold in this theoretical structure. A leading alternative to reputation theories of cooperation admits the possibility that some players actually are altruistic. In this way, cooperate is not a dominated strategy since true payoffs differ from those given in Game PD.

However, relatively little attention has been paid to the testing of theories that attempt to explain this behavior and that is the goal of this paper.⁴ Relative to the existing literature on cooperation in experimental games, this paper makes two contributions. First, we analyze both one-shot and repeated games to gauge the relative importance of reputation effects. In light of existing experimental evidence of cooperation in one-shot games, we know that the reputation arguments of Kreps *et al.* do not explain all observations. One open question is that of whether the reputation model is consistent with observed play in the finitely repeated game. Andreoni and Miller (1991) make a similar comparison between one-shot and finitely repeated games and conclude, based on aggregate play, that there is evidence in favor of the reputation model.⁵ Relative to one-shot games, we find that the amount of cooperative play increases when players interact for a finite number of plays but the pattern of individual play is not consistent with that predicted by Kreps *et al.* Further, the dynamic pattern of cooperation predicted by the theory does not match the observed aggregate pattern of play: cooperation rates do not fall nearly as quickly as predicted by the theory. Also, we evaluate alternative theories of cooperation in one-shot games based on altruism and find support for the hypothesis that a significant fraction of players in our cohorts are altruists. Using maximum likelihood methods, we estimate that approximately 12–13% of the subjects in our sample are altruists.⁶ Overall, we find that neither the altruism model nor the reputation model of Kreps *et al.* is consistent with our observations. The reputation model is inconsistent with observed cooperation in one-

⁴ Dawes and Thaler (1988) do confront some of the theories with evidence from selected studies. Our approach here is somewhat different, in that we are designing an experiment explicitly to test the theories and so will attempt greater control from treatment to treatment. Further, our goal is to test the theories for games outside of those that they were constructed to explain.

⁵ Kahn and Murnighan (1993) study the effects of experimenter-controlled uncertainty about rivals' payoffs on cooperation in finitely repeated PD and related games and find more cooperation than predicted by the theory. They do not consider models of altruism or the possible influence of uncertainty from uncontrolled sources.

⁶ In this we follow the lead of McKelvey and Palfrey (1992), who estimate that about 5% of their subjects playing the centipede game were altruists.

shot games, an observation that the model of altruism can explain. However, neither model is consistent with the frequency of cooperation observed in the finitely repeated game.

II. THEORIES OF COOPERATION

Two main types of theories have been offered to explain why players cooperate in PD games. The first applies only to agents playing the game repeatedly and involves history-dependent strategies. These theories, associated with Kreps *et al.*, maintain the assumption of self-interested players and rely on the repeated nature of the game to create incentives for cooperation. In these models, the key assumption is that players hold a small belief that their opponent is a cooperative player and this induces the self-interested players to cooperate in a finitely repeated PD game. The second type of theory postulates that at least some agents are not strictly self-interested and benefit from cooperation in a manner not reflected in the payoff matrix provided in PD experiments. We discuss the implication of these models for observed play of one-shot and finitely repeated PD games.

a. One-Shot Games

i. *Reputation.* The Kreps *et al.* model assumes that, while players believe that a fraction of their opponents are altruists, all players are in fact egoists. While these “irrational beliefs” have considerable power in generating cooperative play in finitely repeated games, it is equally clear that in a sequence of one-shot games, the theory of Kreps *et al.* predicts that cooperation rates will be zero.

ii. *Altruism.* In models with altruism, in contrast, there are assumed to be a subset of players for whom cooperate is not a dominated strategy. To study this, we restrict attention to a “warm glow” model in which a player receives an additional payoff by cooperating in the PD game. Consider the following payoff matrix where the entries correspond to the payoffs of Game PD except that the row player is assumed to be an altruist.

		Egoist	
		Fink	Cooperate
Altruist	Fink	b, b	c, d
	Cooperate	$\delta + d, c$	$\delta + a, a$

This is a “warm glow” model in that the payoffs of the row player in the event cooperate (C) is chosen are augmented by $\delta \geq 0$.⁷ In general, we will assume that δ is distributed across the population according to a cumulative distribution function $G(\delta)$. When $\delta = 0$ this game is the same as Game PD. We term players with δ less than $\min(b - d, c - a)$ *egoists* since fink (F) is a dominant strategy for them. If δ exceeds both $b - d$ and $c - a$, then cooperate becomes a dominant strategy for the row player. We term players with payoffs satisfying these restrictions *dominant strategy altruists*. If δ exceeds $c - a$ but is less than $b - d$, then cooperate is no longer a dominant strategy so that cooperative play could be rationalized only by a belief that a rival is cooperating with a sufficiently high probability. Players with these preferences are *best response altruists*.

As long as there are enough players with $\delta > c - a$, this framework rationalizes observed cooperative play in one-shot games. Clearly, if there are dominant strategy altruists, those players will cooperate in all periods of play. If there are only best response altruists, an equilibrium always exists in which altruists and egoists fink. However, if the proportion of altruists in the cohort is large enough ($G(c - a)$ is sufficiently small), then there will also exist an equilibrium in which the altruists cooperate and the egoists fink. Finally, when this equilibrium exists, there will also exist a third equilibrium in which the egoists fink and the altruists randomize between fink and cooperate. Thus, a model with best response altruists can have multiple Nash equilibria which can be Pareto-ranked.⁸ Note that absent learning, there is no reason for the distribution of play to change over time in a sequence of one-shot plays of this game.

b. *Finitely Repeated Games*

i. *Reputation.* It is well known that if two players play an infinitely repeated PD game, equilibria, in addition to those from the stage game, will exist. For example, consider the “grim strategy” where each player

⁷ More generally, one might consider specifications of preferences for cooperative players in which either cooperative is a dominated strategy or cooperate is a best response to cooperate and fink is a best response to fink. The preferences considered here are a simple way of modeling these two patterns of best responses for cooperative players. In fact, one could go further and allow for δ_i to depend on past plays. See Bergstrom *et al.* (1986), Andreoni (1989), and Andreoni and Miller (1991) for a discussion of alternative theories of altruistic behavior in public goods environments.

⁸ To see why these equilibria are Pareto-ordered note that all players receive a payoff of b in the equilibrium in which they all fink. In the equilibria with partial or complete cooperation by altruists, egoists earn a payoff exceeding b since there is a positive probability of being matched with a player cooperating and thus receiving c . Since any of the altruists could fink, their return in either the full cooperation or mixed strategy equilibrium must yield an expected payoff in excess of b . By a similar argument, it is easy to see that the payoffs of both types are higher in the full cooperation equilibrium than in the mixed strategy equilibrium.

cooperates until the other defects (i.e., finks in Game PD) and then finks forever after. If adopted by both players, this strategy yields a subgame perfect Nash equilibrium involving cooperation forever as long as players do not discount the future too much.⁹ With only a finite number of repetitions, however, backward induction arguments imply that the only subgame perfect Nash equilibrium will involve finking in every period. Given that most available experimental results relate to finitely repeated games, reputation effects associated with infinite horizon games are unlikely to be an explanation for the observed cooperation.^{10,11}

In an important contribution, Kreps *et al.* offered a theory of cooperation based upon an assumption of incomplete information in PD games. If a player assigns a positive probability that his opponent has adopted a tit-for-tat strategy, it may be optimal for him to cooperate in early rounds. In a tit-for-tat strategy a player cooperates in the first period, then in every subsequent period plays whatever his opponent played in the previous period. As Kreps *et al.* argue, there may be an equilibrium in which a purely self-interested player will cooperate in early periods to obtain the benefits from the (cooperate, cooperate) outcome of the stage game. Cooperation occurs early in the game when the short-run gains from defecting do not exceed the long-run costs since, following an initial defection, the equilibrium strategies dictate that both players fink forever after.

Alternatively, if a player allows for the possibility that his opponent is a cooperative player, for whom cooperate is a best response to cooperate, then cooperative play may also emerge. Kreps *et al.* argue that for this game there exists an equilibrium in which self-interested players disguise themselves as altruists by cooperating in early rounds of play. However, in order to induce self-interested players to cooperate, it is necessary that the "cooperative" players best respond to fink by finking in the stage game, i.e., they must be best response altruists. If, alternatively, cooperation is a dominant strategy for altruists, then there is no basis for reputation building in the Kreps *et al.* model. This places restrictions on models of altruism that we return to below. The most powerful result from the Kreps *et al.* model is that cooperation for some periods can arise even when the proba-

⁹ For a discussion of this result and others related to the Folk theorem, see Fudenberg and Maskin (1986) and the references therein.

¹⁰ While all experimental games are of finite length, it is possible to randomly choose an endpoint so that the results can be interpreted as the outcome of an infinite horizon game with discounting in which the probability of ending the game at a point in time is imbedded in the discount rate. See Roth and Murnighan (1978) for a discussion of this point.

¹¹ One might argue that players must learn about the logic of backward induction so that, in initial replications of finitely repeated games, cooperation might be observed. See Selten and Stoecker (1986) for evidence that, as experience with finitely repeated games increases, the amount of cooperative play falls.

bility of one's opponent being a tit-for-tat or a cooperative player is quite small.

In our analysis of the data from a finitely repeated version of game PD, we will be comparing the path of play against the predictions of the Kreps *et al.* equilibrium. To understand that comparison, we present an equilibrium path for the game in which self-interested players (egoists) attach a small probability to their opponents being best response altruists. We concentrate on the equilibrium with mixing by egoists so that cooperation rates fall along the equilibrium path, as is generally observed in finitely repeated PD games.

Suppose that players are matched for T periods where PD is the stage game. An algorithm for finding the mixed strategy (by the egoists) equilibrium in a T -period PD game is based on the following three equations, which hold for $t = 1, 2, \dots, T$:

$$\rho_t = \rho_{t-1}/(\rho_{t-1} + (1 - \rho_{t-1})\alpha_{t-1}),$$

$$(\rho_t + (1 - \rho_t)\alpha_t)(a - c + V_{T-t}(\rho_{t+1}) - V_{T-t}(0)) = (1 - \rho_t)(1 - \alpha_t)(b - d),$$

$$V_{T-t}(\rho_{t+1}) = (\rho_{t+1} + (1 - \rho_{t+1})\alpha_{t+1})(a + V_{T-(t+1)}(\rho_{t+2}) \\ + (1 - \rho_{t+1})(1 - \alpha_{t+1})(d + V_{T-(t+1)}(0))).$$

In this system, ρ_t is the period t probability that a player is an altruist conditional on that player having cooperated for the first $t - 1$ periods. The probability that an egoist cooperates in period t is α_t given that both players have cooperated in all previous periods.

The first equation is Bayes' rule, which updates beliefs about the type of an opponent conditional on cooperation by both players until period t . In the event that a player fails to cooperate in any period, $\rho_t = 0$ for all future periods. The second equation ensures that egoists are indifferent between cooperation and finking in period t . Thus the algorithm solves for an equilibrium in which egoists mix between cooperate and fink. Under this condition, $V_{T-t}(\rho_{t+1})$ is the expected payoff for the remaining $T - t$ periods of the game assuming both players cooperate through period t , where ρ_{t+1} is determined by Bayes' rule using the equilibrium probability of cooperation by an egoist. So, under the second condition, $V_{T-t}(0)$ is the value associated with the remaining $T - t$ periods of the game after a defection, with $\rho_t = 0$ in all future periods. The third condition relates the value with $T - t + 1$ periods remaining to the values for the subsequent period given beliefs and actions in period $t + 1$.

To characterize a candidate equilibrium in which an egoist cooperates in period t with probability α_t , one can solve the system of equations given

an arbitrary belief in the last period, ρ_T . Since $\alpha_T = 0$, if ρ_T is given, one can then calculate $V_1(\rho_T)$ and then solve for α_{T-1} and ρ_{T-1} from the egoist's incentive condition and Bayes' rule. Continuing in this fashion generates the entire time path of beliefs, egoists' decisions for each period, and an initial belief consistent with this path of play. Thus the time path of cooperation rates is parameterized by ρ_T or, equivalently, by the initial beliefs ρ_0 . In addition to this equilibrium with mixing by the egoists, there may also exist an equilibrium in which egoists cooperate with probability one until the last period.

The power of this model is the possibility of generating cooperative play with only a small initial belief that an opponent is altruistic. Another important feature of this equilibrium is the time path of cooperation. Since egoists are mixing along the equilibrium path, the probability of cooperation in a given period is falling over time. In this equilibrium, the dynamics are the consequence of strategic interactions not of learning.

To better understand this type of equilibrium, consider the case of $T = 2$. Let $E(x, y)$ denote an egoist's payoff from playing x in period 1 and y in period 2. The expected payoff from cooperating in the first period and playing fink in the second is given by

$$E(c, f) = \rho_1(a + c) + (1 - \rho_1)[\alpha_1(a + b) + (1 - \alpha_1)(d + b)],$$

where α_1 is the fraction of egoists who cooperate in period 1. The first term is the payoff to the egoist in the event that the other player is an altruist (a possibility that these agents entertain) and the second term represents expected payoffs when matched with an egoist randomizing between cooperate and fink. Instead, by finking in both periods, the egoist would obtain

$$E(f, f) = \rho_1(c + b) + (1 - \rho_1)[\alpha_1(c + b) + (1 - \alpha_1)2b].$$

For this two-period game, there are two types of equilibria. First, if $\rho_1 \geq (c - a)/(c - b)$ then there is an equilibrium in which egoists cooperate with probability one in the first period and fink with certainty in the second. This behavior by the egoists is supported by the beliefs that the fraction of altruists is sufficiently high. In addition, there are (ρ_1, α_1) pairs that satisfy the condition $E(c, f) = E(f, f)$ leaving $\alpha_1 \in (0, 1)$ so that egoists mix. A necessary condition for this is that $\rho_1 \geq (c - a)/(c - b)$.¹²

In our analysis of the 10-period game, we focus on mixed strategy equilibria such that ρ_9 , beliefs at the start of period 9, are at least as large as

¹² If $\rho_1 \geq (c - a)/(c - b)$, there will also exist an equilibrium in which all egoists cooperate with probability one until the last period.

$(a - c)/(b - c)$. If not, egoists would not cooperate in period 9 and hence would have no incentive to cooperate earlier. Of course, in the 10-period game, the beliefs held in period 9 reflect the mixing that takes place in the first 8 periods. Thus, and this is the power of the Kreps *et al.* model, very low initial beliefs can rise rapidly due to the mixing by egoists.

ii. *Altruism.* The model of altruism has some interesting implications for play in finitely repeated games. In particular, even if egoists do not cooperate in the repeated game (i.e., even if they are not building reputations), more best response altruists (those with lower values of δ) will choose to cooperate due to the repeated nature of play. Thus, increased cooperation can occur in repeated games relative to one-shots *without* reputation building by egoists.

To illustrate, consider Game PD and suppose, as in our experiment, that $a = 800$, $b = 350$, $c = 1000$, and $d = 0$. Further, assume there are three types of players: $\frac{3}{4}$ with $\delta = 0$, $\frac{1}{8}$ with $\delta = \delta_1$, and $\frac{1}{8}$ with $\delta = \delta_h > \delta_1$. For the one-shot game, if $\delta_h > 350 - (\frac{1500}{8}) > \delta_1$, then there is an equilibrium in which only the altruists with the high level of warm glow will cooperate. In fact, if $\delta_1 < 350 - (\frac{1500}{4})$, this will be the only equilibrium on the one-shot game.

Now, suppose that Game PD is played twice. As long as $\delta_1 > 200$, there will be an equilibrium in which all of the altruists cooperate in period 1 and cooperate in period 2 iff their opponent cooperates in period 1. Since the proportion of altruists is relatively small, egoists will fink in both periods. Thus we see that even without the reputation building effects stressed in Kreps *et al.*, repeated play can increase cooperation rates by inducing cooperation by more of the altruists.

More generally, for a T-period repeated game one can construct an equilibrium in which the egoists do not build reputations (i.e., play fink) and altruists cooperate in the first period and continue to cooperate iff their opponent cooperated in period 1 as well. In particular, choosing fink in every period will be optimal for an egoist, given that other egoists fink too, iff $\rho_1 \leq \rho^E$, where

$$\rho^E \equiv (b - d)/((b - d) + (T - 1)(a - b)).^{13}$$

Likewise, an altruist with warm glow of δ will follow the strategy given above if $\rho_1 \geq \rho^A$, where

$$\rho^A \equiv (b - d - \delta)/(Ta - (T - 2)b - c - d + (T - 1)\delta).$$

¹³ In deriving this condition, we make use of the fact that if an egoist were to cooperate in the first period and be matched with an altruist, the pair would cooperate until the egoist finks in the last period.

Using the fact that $\delta + a \geq c$, one can show that $\rho^E > \rho^A$. Further, both ρ^E and ρ^A fall with T . An equilibrium for a T -period game in which players select these strategies is characterized by a level of cooperation $\rho^*(T)$ satisfying two conditions. First, $\rho^E > \rho^*(T) > \rho^A$. Second, in this equilibrium, there is a critical level of warm glow, denoted $\delta^*(T)$, such that only altruists with $\delta \geq \delta^*(T)$ cooperate with $\rho^*(T) = 1 - G(\delta^*(T))$.

III. EXPERIMENTAL DESIGN

Each experiment was conducted using cohorts of players recruited from undergraduate, sophomore and above, and graduate classes at the University of Iowa. Players were seated at separate computer terminals and given a copy of the instructions. Since these instructions were also read aloud, we assume that the information contained in them is common knowledge. These instructions are reproduced in the Appendix.¹⁴

The matrix used was Game PD with $a = 800$, $b = 350$, $c = 1000$, and $d = 0$.¹⁵ We induced payoffs in terms of utility using the Roth–Malouf (1979) procedure. In the matrix games, each player's payoff was given in points which determine the probability of the player winning a monetary prize. At the end of each period of play, we conducted a lottery where “winning” players received a \$1 prize and “losing” players received nothing.¹⁶ This procedure is designed so that all utility maximizing players will maximize the expected number of points in each game regardless of their attitudes toward risk and ensures that these utility payoffs are common knowledge.¹⁷

a. *One-Shot Games*

Each player participated in a sequence of one-shot games against different anonymous opponents. One was designated the row player and the other the column player. All pairing of players was done through the computer. No player knew the identity of the player with whom he was currently paired or the history of decisions made by any of the other players.

¹⁴ The Appendix contains instructions for the basic PD game and the finitely repeated game.

¹⁵ Strictly speaking, these are the payoffs for egoists playing this game.

¹⁶ For altruists, view the payoffs that include warm glow effects as the actual payoffs and thus ignore the translation of points into dollars. In terms of our experiment, we think of cooperative players as obtaining utility from playing the cooperative strategy in addition to the points earned from the choice of this strategy. That is, suppose that the utility function for a cooperative player depends on the dollars earned and a warm glow term. One can then interpret the δ_i terms as the “point equivalents” of these warm glow effects.

¹⁷ The experiments took from one to one and one-half hours to complete. Payments ranged from \$4 to \$16.

In these games, we employed a matching design in which reputation effects were not feasible. Under this design, 40 players participated in 20 one-shot games. The players were split into two groups, red and blue. In each period, a red player was anonymously paired with a blue player. The pairings were constructed so that players were matched with a different player each period so that contagion effects, associated with Kandori (1988), and reputation effects from repeated play against a fixed opponent could not arise.¹⁸

The easiest way to understand the matching procedure employed here is to consider Townsend's (1980) turnpike model in which an agent belonging to a sequence of agents traveling east on a turnpike is matched with another agent belonging to a sequence of agents traveling west; all agents travel at the same speed. These agents are matched for one period and then move along in opposite directions to meet new opponents. The important property of this model is that actions taken by a pair in period t cannot influence the behavior of the agents these players will be matched with in the future: the turnpike ensures that histories lie behind a given player. Since our experiment consisted of a finite number of players, we converted the turnpike into a circle and constructed a matching matrix with the same non-contagion property of Townsend's turnpike.

This matching procedure was explained to the players and they were presented with the actual matrix used in this procedure as well as an example (see the Appendix). At no time could they ever identify an opponent: anonymity was maintained throughout. Players alternated between being row and column players during the treatment.¹⁹

b. *Finitely Repeated Games*

To explore the effects of repeated play on cooperation, 30 players, in three separate cohorts, each played two 10-fold repetitions of Game PD. In both repetitions, subjects were randomly and anonymously paired. Each

¹⁸ Kandori's work demonstrates that reputation effects might arise in infinitely repeated PD games in which self-interested players are randomly matched at the start of each period and histories are not known. Building on Kreps *et al.*, it would appear that contagion effects might also exist in finitely repeated games of incomplete information. Along the equilibrium path of the PD game, players cooperate. In the event of a defection by an opponent players defect forever after. While the player who initially defects (say, player A) is not necessarily punished the next period (unless A is matched with the same opponent) the effects of his defection will spread throughout the community. Eventually A will meet an opponent who is not cooperating because A had started a chain of defections. In this way, A is eventually punished for his defection. This contagion effect, which reminds us that random matching and rematching over several periods is not the same thing as a series of one-shot games, can support cooperative play for low enough discount rates and a high enough probability of being matched with someone not cooperating because of an earlier defection.

¹⁹ After each period, a different player was selected to draw the lottery ticket.

subject played the same opponent 10 times as either a row or a column player. Players were then anonymously matched with new opponents and play continued for 10 more periods. At the end of each period in a repetition, subjects were told the play of their opponent. However, when players were rematched, they were not told anything about the history of play of their new opponent. After the first repetition, row (column) players became column (row) players for the second repetition.

Because of our interest in cooperation levels across one-shot and finitely repeated games, we wanted to provide a comparable experience for the two treatments. As described under Results, in evaluating the data from the one-shot games, we concentrate on outcomes in the last 10 periods. Thus, to make the results in the finitely repeated treatment comparable to those from the one-shot treatment, the first 10-period repeated game was preceded by the play of 10 one-shot versions of the PD game. In this treatment we utilized the matching procedure from CDFR1.^{20,21}

c. *Treatments*

Thus we have results to report on the following treatments:

- (i) PD: This is the series of one-shot plays of Game PD.
- (ii) PD-FR: This is the 10-fold finitely repeated version of Game PD.

IV. RESULTS

We begin with an overview of the cooperation rates observed in PD and PD-FR and then turn to an evaluation of the models. These results are from the last 10 periods of PD and the 10-period repeated games from the PD-FR treatment.²² This comparison assumes that the 10 one-shot games

²⁰ With this matching procedure, there were three cohorts, each consisting of 11 players. In the one-shot session, players were anonymously paired for one period and then matched to another player in a subsequent period. Players alternated between row and column from one period to the next. Over 11 periods, each player was paired with every other player once and sat out once. The second and third sessions of the experiments consisted of the two 10-period repetitions of Game PD described above. The last subject (11th player) to arrive for the experiment in each cohort was told in the instructions that he/she would only play in the first session.

²¹ We claim that the match–rematch design itself has no effect on cooperation. The basis for this claim is data from a coordination game with a cooperative dominated strategy, Game 3 of CDFR1, where the treatments were the matching procedure, the Townsend turnpike design, and our match–rematch design (CDFR1). There were no differences across treatments.

²² We rejected independence across periods for all periods of the PD treatment but were unable to reject independence across periods of play for the last 10 periods of the PD treatment. Hence, those results are pooled. Further, there was no evidence that play in PD-FR differed across the two matches so that those results are pooled in Fig. 1. Table I reports on the cooperation rates in each of the matches from PD-FR.




FIGURE 1

that precede both the last half of PD and the repeated games in PD-FR provide comparable experience for our subjects.

Figure 1 displays the time pattern of play and Table I reports the frequency of cooperative play over the last 10 periods for these two treatments. Three striking features emerge from these treatments.²³

Observation 1. Cooperation rates are positive and generally declining over time in PD.

TABLE I

Rates of Cooperation	
Treatment	Proportion of cooperative play
PD, last 10 periods	0.22
PD-FR, first match	0.52
PD-FR, second match	0.57

²³ Similar patterns are observed by Andreoni and Miller (1991), though they offer a different interpretation of the observations.

Observation 2. Cooperation rates are positive and generally declining over time in PD-FR.

Observation 3. Cooperation rates are higher in PD-FR than in PD.

Observations 1 and 2 are consistent with the findings of many experiments examining cooperation in one-shot and finitely repeated PD games.²⁴ That is, cooperation has been observed in most experimental one-shot and finitely repeated PD games. Further, the decline in cooperation rates is frequently observed and, in fact, motivated the reputation model of Kreps *et al.* Observation 3 is somewhat novel in that most studies do not present a comparison of cooperation rates for finitely repeated and one-shot games.

The central issue is to determine the consistency of these observations with the models discussed above. In particular, what is the importance of reputation building relative to altruism in our results? To address this we first look in detail at the one-shots and then evaluate PD-FR.

a. *One-Shot Prisoner's Dilemma Game*

The strongest evidence that altruism is at least part of the explanation for observed cooperative play comes from these one-shot PD games. If players are egoists, cooperation will not be observed in one-shot PD games. This hypothesis is clearly rejected by the data: over the entire 400 observations (the last 10 periods of the 40 player treatment), cooperation rates were greater than 20%. Moreover, the amount of cooperative play was much larger in the initial periods, averaging 38% in the first 10 periods.

Figure 2 gives the distribution of cooperation rates in the second half of this treatment. The distribution reveals that almost all cooperative play comes from players who do not always cooperate: dominant strategy altruism does not explain the results. Thus we focus on evaluating best response altruism.

Our approach is to use our model of best response altruism to estimate the fraction of altruists and then to test the restriction that this fraction is zero. We estimate the fraction of altruists in our population using two different methods. The first uses the observed fractions of play at the aggregate level to estimate the fraction of altruists (ρ) in our sample. The second method looks at individual play and classifies players as altruists and egoists. We present the results from these two estimation strategies in turn.

The first approach to estimating the fraction of altruists is to choose the value of ρ that maximizes the likelihood of the observed cooperation rates,

²⁴ Again, see the discussion and references in Dawes (1980), Dawes and Thaler (1988), and Roth (1988).




FIGURE 2

given in the first row of Table I. Thus the initial estimate is simply $\rho = .22$. However, this procedure assumes that egoists fink with probability one, altruists always cooperate, and mistakes do not occur. Attempting to distinguish egoists' mistakes from altruism is impossible if we restrict attention to the aggregate cooperation rate since the fraction of observed cooperation does not allow us to identify both the fraction of altruists and the probabilities that they and the egoists make mistakes.

The alternative estimation strategy is to estimate the fraction of altruists at the individual level. This allows us to identify types as well as the probability of making a mistake for each type. Assume that altruists intend to cooperate and egoists fink, but that mistakes are possible. With this model, we can estimate the proportion of altruists and egoists in this sample and the probability of mistakes. To see how this is accomplished, consider the choice to fink (F) or cooperate (C) by an individual of type $\tau \in \{-1, 1\}$, where $\tau = 1$ denotes an altruist and $\tau = -1$ an egoist. Let X_t indicate the period t outcome of the agent's choice, where $X_t = \tau + \lambda u_t$. In this expression, u_t is a uniform random variable in the interval $[-1, 1]$ and λ parameterizes a mean preserving spread of this random variable. This decision (d_t) is observed and reflects the realized value of X_t . Assume that $d_t = F$ if $X_t < 0$ and $d_t = C$ otherwise. Thus the restriction that $\tau = 1$ for an altruist implies that this player is quite likely to choose C while egoists are more likely to choose F. Both types, however, can err and this is the point of introducing noise into the decision process.

At the individual level, the maximum likelihood estimates of each player's type (altruist or egoist) and his probability of making a mistake are easily determined. The estimated fractions of players of each type are given in

TABLE II

Proportion (%) of Player Types from PD (40 Players)		
Type	First 10 periods	Last 10 periods
Altruist	15	12.5
Egoist	62.5	85
Other	22.5	2.5

Table II. A player is classified as an altruist (egoist) if he plays cooperatively more (less) than 50% of the time. A player cooperating exactly 50% ("other" in Table II) of the time cannot be unambiguously identified. Note that we do not find that all observations of cooperative play are the consequence of mistakes by egoists.

While the proportion of altruists remained roughly the same in the first and second halves of this treatment, the egoist group grew in the second half as a number of players who had played cooperatively exactly half the time became less cooperative. Only one first half egoist became a second half altruist.

Given the identity of either an altruist or an egoist, the frequency that a player selects a strategy other than that indicated for his type provides an estimate of that player's individual mistake probability. This probability, labeled ν_i , is directly related to λ_i

$$\nu_i = (\lambda_i - 1)/2\lambda_i.$$

For these players, mistakes can arise either because agents fail to select their optimal action or because optimal actions depend on beliefs which we do not observe. Thus, if we observe a best response altruist choosing fink we term this a mistake even though that action might be rationalized by the beliefs of the player at that point in the treatment. Data on the distribution of these individual mistake proportions are given in Table III.

TABLE III

Type	Distribution of Individual Mistake Proportions, Last 10 Periods					Total
	0%	10%	20%	30%	40%	
Altruist	1	0	1	0	3	5
Egoist	14	3	7	9	1	34

Though the sample of altruists is rather small, there is evidence here that egoists are less prone to mistakes. The average mistake probability among egoists was 14% while for altruists it was 28% in the last 10 periods.²⁵ Since egoists have a dominant strategy, mistakes by altruists are more likely because, as noted above, a revision of beliefs can lead an altruist to fink, an event which we would count as a mistake. Further, using these estimates of the proportion of altruists and the mistakes probabilities by each type, mistakes by egoists are more costly than mistakes by altruists.

To get some measure of how well this model explains observed behavior, we evaluate the predictions of the theory, based on our estimates of individual types and mistake frequencies, to actual aggregated play. The probability of an outcome, say (fink, fink), equals the probability that two egoists meet and do not make mistakes, plus the probability that two altruists meet and both make mistakes, plus the probability that one altruist making a mistake meets an egoist not making a mistake, plus the probability that one egoist not making a mistake (or an altruist making one) meets an "other" type playing fink by chance, plus the probability that two "other" types meet and both fink:

$$\begin{aligned} \text{Prob (fink, fink)} &= \varepsilon^2(1 - \nu_e)^2 + \rho^2\nu_a^2 + 2\rho\varepsilon\nu_a(1 - \nu_e) \\ &+ \gamma\varepsilon(1 - \nu_e) + \gamma\varepsilon\nu_a + \gamma^2/4; \end{aligned}$$

where ρ is the proportion of altruists, ε is the proportion of egoists, γ is the fraction of "others," and ν_a and ν_e are the mistake probabilities of altruists and egoists, respectively. The probabilities of the other outcomes are expressed in a similar way.

Using data on the proportions of types from the second column of Table II and the mistake frequencies $\nu_a = 0.28$ and $\nu_e = 0.14$, as given above, we would predict the following frequencies, with the observed frequencies in parentheses:

$$\begin{aligned} \text{prob (fink, fink)} &= 0.606 \text{ (0.620)} \\ \text{prob (coop, coop)} &= 0.049 \text{ (0.065)} \\ \text{prob (coop, fink)} &= 0.172 \text{ (0.155)} \\ \text{prob (fink, coop)} &= 0.172 \text{ (0.160)} \end{aligned}$$

²⁵ For the first 10 periods of play, the average proportion of mistakes by egoists was over 24%, while that by altruists (27%) was about the same as that from the second half of the treatment.

These predictions are close to the actual observations. Define the statistic τ as $\tau = \sum (N_i - Np_i)^2 / (Np_i)$, where N is the total number of observations (here, 200), N_i is the actual number of observations in the i th cell where the cells correspond to the outcomes (fink, fink), (cooperate, cooperate), (fink, cooperate), or (cooperate, fink), and p_i is the theoretically predicted proportion of observations in cell i . If K is the number of cells, the statistic τ can be shown to be distributed as χ^2 with $K - 1$ degrees of freedom. The statistic τ equals 1.613 with three degrees of freedom. This value is so low as to make it impossible to reject the hypothesis that the model with mistakes explains these observations.

Thus, we conclude that our best response altruist model, with the possibility of mistakes, best explains the PD data. For the last half of play, we estimate the proportion of altruists to be 12.5%.²⁶ From the perspective of the competing theories, it is important to note that these estimates of ρ are different from zero. The maximum likelihood estimation could have attributed all cooperation to mistakes by egoists. But, given the high frequency of cooperation by some players, it was more likely that they were altruists.

b. *Finitely Repeated Prisoners' Dilemma*

The results from PD-FR provide an alternative means of evaluating the two theories. The model of Kreps *et al.* was constructed to roughly match observed cooperative play in finitely repeated PD games. The declining cooperation rate shown in Fig. 1 is certainly consistent with the reputation theory.²⁷ Thus, qualitatively, Observation 2 matches the prediction of the reputation model.²⁸

However, there are a number of inconsistencies between the observations and the reputation model. The reputation model can be evaluated by focusing on play at the pairs and individual level. Study of the pairs data reveals little support for a pure reputation explanation for cooperation. Excluding as inconsistent with the Kreps *et al.* theory any pair in which (i) cooperative play follows noncooperative play by the same player or his/her opponent

²⁶ Using this and the estimated type-contingent mistake probabilities given above, the critical value of δ such that altruists are indifferent between cooperate and fink is about 317.

²⁷ It should be noted that the cooperation rates are actually not strictly decreasing period by period. As noted later, this is a consequence of the fact that the strategy described in the Kreps *et al.* model is not followed by a majority of the players.

²⁸ Andreoni and Miller (1991) reach a similar conclusion based on aggregated data but do not focus, as we do below, on individual play. As discussed below, conclusions based on aggregate play are misleading since individual players do not generally play in a manner consistent with the equilibrium outlined by Kreps *et al.*




FIG. 3. Predicted and actual cooperation rates for PD-FR.

and/or (ii) either player cooperates in the last period, we find that only 10% (3 out of 30) of the pairs performed in a manner consistent with the theory. At the individual subject level, 25% (15 out of 60) of all players did nothing inconsistent with the theory. In the final period of play, cooperate was chosen in 14 out of 60 observations, almost 25%. Overall, while the level of cooperation and its decline over time is qualitatively consistent with Kreps *et al.*, there is little evidence of their equilibrium at the pairs and individual level.

Further, it is useful to compare the theoretical predictions for the finitely repeated game with the observations from PD-FR. Figure 3 plots the time path of cooperation rates predicted from a mixed strategy equilibrium of the Kreps *et al.* model, as discussed in Section II, and the actual cooperation rates from PD-FR.²⁹ The predicted cooperation rate line was generated by solving for an equilibrium path in which the probability an opponent is an altruist, conditional on cooperation by both players throughout, reaches $\frac{4}{13}$ after period 8. As explained earlier,

²⁹ We focus on the mixed strategy equilibrium since the equilibrium in which egoists cooperate with probability one is easily rejected.

this is an important equilibrium path since it is necessary to have the conditional probability of altruism rise to $\frac{4}{13}$ by the start of period 9 in order to obtain any cooperation by egoists. Note that in the early periods, the predicted cooperation rates are reasonably close to those actually observed in the PD-FR treatment.

The most striking feature of Fig. 3 is that the observed cooperation rate remains quite high relative to that predicted by the model. In particular, while the model predicts cooperation will be almost 70% in the first period, by period 5 cooperation rates should fall below 10%. This low cooperation rate is necessary to boost conditional beliefs up to the $\frac{4}{13}$ level by period 9 as conditional beliefs rise only to the extent that egoists place relatively low weight on cooperation in their mixed strategy. The observed cooperation rate, however, is above 50% through period 9 and is in excess of 20% in period 10. Thus the time pattern of play observed is rather different than predicted even though qualitatively the observed cooperation rates have the predicted pattern of decline over time. Note too that the observed cooperation rates are a concave function of time while the predicted cooperation rates are convex.

The predicted time path of play for the reputation model is parameterized by the beliefs that agents hold at the start of the game, ρ_1 . The predicted line in Fig. 3 assumes that $\rho_9 = \frac{4}{13}$, as this value leads to the highest cooperation rates over the finitely repeated game. While one could evaluate the likelihood that the observed pattern of play emerged from the mixed strategy equilibrium, the fact remains that the positive level of cooperation observed in the last period, as well as some of the deviations from the predicted play at the individual level, is inconsistent with the theory.³⁰

As argued earlier, repeated play can induce increased cooperation by altruists (players with $\delta > 200$) who might choose not to cooperate in one-shot games. Thus it is conceivable that the higher cooperation rates observed in PD-FR could occur without any reputation building by egoists. However, from the evidence, it seems quite unlikely that the observed cooperation rates in PD-FR are a consequence of cooperative play by altruists alone.

First, the equilibrium of the repeated game without reputation building by egoists has a particular structure: Altruists cooperate in the first period while all egoists fink. If both players cooperate in period 1, then they are both known to be altruists and they cooperate for the remainder

³⁰ That is, since this is a mixed strategy equilibrium, there will be a distribution of outcomes in the 10-period game which, in principle, could be evaluated. However, the cooperation in period 10 has zero probability and thus can only be explained by formally adding mistakes to this model. We are grateful to Paul Beaudry for conversations clarifying this point.

of the game. In all other cases, both players fink for the remainder of the game. Empirically, this equilibrium implies that cooperation rate should be highest in the first period (but less than 100% assuming there are some egoists) and at a lower, *constant* level for the remainder of the game. For PD-FR, Fig. 1 clearly shows that this is not the case since cooperation rates are not independent of time. Moreover, in the proposed equilibrium, there is no basis for the observed drop in cooperation rates in the last period.

Second, the cooperation rates in periods two through nine exceed 50%. For this to be consistent with the proposed equilibrium requires that the fraction of altruists exceeds 70%. However, this is well above the calculated value of $\rho^E(10)$ so that, according to our earlier discussion, egoists will cooperate which breaks the equilibrium.

V. CONCLUSIONS

The point of this paper was to evaluate competing theories of cooperation. To do so, we designed an experiment to compare cooperative play in one-shot and repeated environments. Two extreme models were considered, one with reputation building and one with altruism. Neither model alone is sufficient to explain observed behavior in the one-shot and finitely repeated PD games. The reputation models fail to explain observed cooperation in one-shot games and the presence of altruists without reputation building by egoists is insufficient to explain the higher cooperation rates and the time pattern of play in finitely repeated games. Thus both models fail to explain the observations off the “domain” that the models were constructed to match. Finally, the reputation model of Kreps *et al.* fails to match observed play in the finitely repeated games.

In principle, one could evaluate a mixed model in which the finitely repeated structure induces more altruists to cooperate and creates an incentive for egoists to build reputations. This is consistent with the higher cooperation rates observed in PD-FR relative to PD and with the drop in cooperation in the last period. With regard to individual pairs, about 30% are perfectly consistent with this equilibrium.³¹ This is higher than the 10% consistency reported for the Kreps *et al.* model though it is clearly much less than 100%. The mixed model also fails to explain the time path of play observed in PD-FR. Perhaps our inability to fit the observations better reflects our limited model of reputation building since, following Fudenberg

³¹ This figure was calculated in the same manner as the 10% figure reported in the evaluation of the reputation model in Section IVa except that condition (ii) is excluded.

and Maskin (1986), we know that a wide variety of equilibria can arise in this class of models if one searches more broadly for alternative types of “irrational” players.

There is an opportunity to take this approach further. First, we can use the data generated by our finitely repeated game to estimate the proportions of egoists and altruists. Second, we have not yet integrated the possibility of mistakes into the model of repeated play which, while quite complicated, may be quite interesting. Consider again the twice-repeated PD game described in Section II. Along the equilibrium path, it was necessary that altruists fink in the second period if their opponent chose fink in period 1. This provides the needed incentive for egoists to cooperate in the first period. Once the model is extended to allow for mistakes, the optimal choice of an altruist in the second period conditional on an opponent selecting fink in the first period will depend on the beliefs that the altruist holds regarding the type of his opponent. If, as our data suggest, egoists are less mistake prone than altruists, then an equilibrium in which all players cooperate in the first period may not be robust to the introduction of mistakes since an altruist may not have an incentive to punish a deviation from the path. A complete analysis of the equilibria for finitely repeated games with mistakes is beyond the cope of this paper, though this modification of the Kreps *et al.* equilibrium may be a useful direction for bridging the gap between theory and observation.

APPENDIX

Instructions for One-Shot PD Game

General

You are about to participate in an experiment in the economics of decision making. If you follow these instructions carefully and make good decisions, you might earn a considerable amount of money which will be paid to you in cash at the end of the experiment.

The experiment will consist of a series of separate decision making periods. Each period consists of two phases. In Phase I you will be paired with another person and, based upon your combined actions, you will be able to earn *points*. In Phase II, you will have the opportunity to earn dollars based upon the points you earn in Phase I. We begin by describing Phase II so that you understand how the points you earn affect the number of dollars you earn. Then, we describe Phase I in detail so that you understand how to earn points.

Phase II Instructions

At the end of Phase I, you will have earned between 0 and 1000 points according to the rules we will discuss below. The number of dollars you earn in Phase II will depend partly on the number of points you earned in Phase I and partly on chance. Specifically, we have

a box which contains lottery tickets numbered 1 to 1000. In Phase II, a ticket will be randomly drawn from the box. If the number on this ticket is LESS THAN OR EQUAL TO the number of points you earned in Phase I, you WIN \$1.00. If the number on this ticket IS GREATER THAN the number of points you earned in Phase I, you WIN \$0.00. For example, if you have 600 points, you will have a 60% chance of winning \$1.00. Notice that the more points you have, the larger will be your chance of winning the \$1.00 prize.

Phase I Instructions

In each decision making period, you will be paired with another person. One of you will be designated the Row player and the other will be designated the Column player. At the beginning of the period, both the Row player and the Column player must separately and independently select an action. The combined actions of the Row player and the Column player jointly determine the number of points earned by the Row player and the number of points earned by the Column player.

You will alternate from being the Row player to being the Column player from one period to the next. At the beginning of each period, you will receive a message on your terminal stating:

“FOR PERIOD _____, YOU ARE A ROW PLAYER.”

or

“FOR PERIOD _____, YOU ARE A COLUMN PLAYER.”

In your folder you will find a record sheet. On this sheet you will indicate, based on the message previously received on your terminal, whether you are a Row player or a Column player.

In the experiment we are going to conduct today, there are 40 participants who are divided into two equal groups of 20 players. The experiment will consist of 20 periods.

The pairings for the periods are constructed in such a way that the decision you make in any one period can have no effect on the decisions of individuals you will be paired with in later periods. To illustrate this, consider the following example with 8 players. The players are divided into two equal groups: the Red Group with players 1 to 4 and the Blue Group with players a to d. Players in the Red Group are called red players and players in the Blue Group are called blue players. There are 4 periods in the experiment and the pairings are as follows:

		Red Group Table			
		Period Number			
		1	2	3	4
Red	1	a	b	c	d
Group	2	b	c	d	a
Players	3	c	d	a	b
	4	d	a	b	c

Blue Group Table

		Period Number			
		1	2	3	4
Blue	a	1	2	3	4
Group	b	2	3	4	1
Players	c	3	4	1	2
	d	4	1	2	3

In each group's table, each player in the group is identified on the left. In the row corresponding to a player, the order in which he/she will play individuals in the other group is shown. Suppose that you are player 1. You would play player a in period 1, player b in period 2, player c in period 3, and player d in period 4. Similarly, you can determine how all other players will be paired.

From these pairing tables, you can see that the decision you make in any one period can have no effect on the decisions of individuals you play in later periods. Suppose you are a red player and that you try to use your current decision against a blue player to influence the decision of some other blue player whom you will be paired with in a future period. Since blue players never play against each other you will have to accomplish this indirectly. At the very least, this would require that your current decision affect a future decision of the blue player you are currently paired with in a particular way. Namely, this blue player must affect the decision of some other red player, who in turn affects the decisions of other blue players whom you will be paired with in the future. But this can never happen since the blue players you will be paired with in the future all play against other red players *before* the blue player you are currently paired with plays those same red players.

To see this in the above example, continue to assume that you are player 1, and notice that the decision you make when paired with player a in period 1 can have no effect on the decision made by player d when you are paired with player d in period 4. This is the case because after you are paired with player a, player a meets players 2, 3, and 4 after they have already been paired with player d. So anything that player a does cannot alter the way players 2, 3, or 4 play when they are paired with player d. Therefore, the decision which player d makes when playing you in period 4 cannot be affected by the way you played against player a.

You can similarly verify that a decision made by any player in any of the other pairings will have no effect on the decisions made by individuals he/she is paired with in the future.

The tables we will use for pairing the 40 players in today's experiment are attached at the end of these instructions (please turn to them now). Players numbered 1 through 20 are in the Red Group and players numbered 21 through 40 are in the Blue Group. The tables were constructed in the same way as the tables given above. You are free to verify that the same principle holds there as in the previous example—there is no way for you to make a decision that will affect the decisions made by individuals you are paired with in the future. Again this is because of the way that players you are paired with in the other group go on to meet other players in your group. In particular, after being paired with you, they are paired with

other members in your group only after the other members in your group have been paired with individuals you will play in the future.

Except for their identification number or letter, you will not be told the name of any person you are playing in any period. Similarly, nobody you are paired with will know your name in any period nor will you be told who these people are either during or after the experiment.

The points that you earn in each period will be determined by the rules given below.

Specific instructions to Row player. In this part of the instructions we will be referring to specific numbers of points. These numbers are the same as you will be using in today's experiment and those found on the session information sheet at the end of these instructions.

In those periods in which you are a Row player, you and the Column player must separately and independently decide on actions which will jointly determine the number of points earned by you and the number of points earned by the Column player. As the Row player, you may choose either action R1 or action R2. Similarly, the Column player may choose action C1 or action C2. The number of points earned by you is given by the following table for each pair of actions you and player S might select:

Number of points earned by Row player

		Column's Action	
		C1	C2
Row's Action	R1	350	1000
	R2	0	800

To read this table, suppose that you chose action R2 and the Column player chose action C1. You would then earn 0 points. Similarly, suppose that you chose action R1 and the Column player chose action C2. You would then earn 1000 points. In a like manner, you can use this table to determine the number of points you would earn for all other pairs of actions you and the Column player may select. Column players also earn points depending upon the type of action they select. These are given in the next section of the instructions.

When you select an action, enter the action chosen into the computer via your terminal and record the action chosen on your record sheet. Once both you and the Column player have selected your actions and entered them into the computer via your terminals, the computer will determine the number of points earned by you based on the table given above. The result is then sent to you via your terminal. The message will look like the one below:

PERIOD POINTS ARE _____.

At the end of the period, you are to record your point earnings for Phase I on your record sheet. Make sure you check your earnings in points against the computer's calculations. The computer will also inform you about the action taken by the Column player. Make sure you record this information on your record sheet.

Specific instructions to Column player. In those periods in which you are the Column player, you and the Row player must separately and independently decide on actions which will jointly determine the number of points earned by you and the number of points earned by the Row player. As the Column player, you may choose either action C1 or action C2. The number of points earned by you is given by the following table for each pair of actions you and the Row player might select:

Number of points earned by Column player

		Column's Action	
		C1	C2
Row's Action	R1	350	0
	R2	1000	800

To read this table, suppose that the Row player chose action R2 and you chose action C1. You would then earn 1000 points. Similarly, suppose that the Row player chose action R1 and you chose action C2. You would then earn 0 points.

When you select an action, enter the action chosen into the computer via your terminal and record the action chosen on your record sheet. Once both you and the Row player have selected your actions and entered them into the computer via your terminals, the computer will determine the number of points earned by you based on the table given above. The result is sent to you via your terminal. The message will look like the one below:

PERIOD POINTS ARE _____.

At the end of the period, you are to record your point earnings for Phase I on your record sheet. Make sure you check your earnings in points against the computer's calculations. The computer will also inform you of the action taken by the Row player. Make sure you record this information on your record sheet.

Phase II recording rules. After completing your Phase I record sheet for a given decision making period, you are to use your profit sheet to record the dollars you earn in Phase II. First, record your Phase I point earnings in the row corresponding to the number of the period that is currently being conducted. One player, a different one in each period, will then be asked to draw a lottery ticket from the box. Before he/she returns the ticket to the box, the number on the ticket will be announced. You should record the number of the ticket in the second column of your profit sheet. If the number drawn IS LESS THAN OR EQUAL TO the number of points you earned in Phase I, circle \$1.00 in the next column; otherwise circle \$0.00 in that column. Pay careful attention to what you circle. Any erasure will invalidate your earnings for the period. If you do make a mistake and circle the wrong number, call it to the experimenter's attention.

At the end of the session, add up your total profit in dollars and record this sum in row 21 of your profit sheet. All dollars on hand at the end of the session in excess of \$0.00 dollars are yours to keep. Subtract this number, which is on row 22, from your total dollars in row

21 and record this difference on row 23. This is the amount of dollars you have earned in this session.

In summary, your earnings in the experiment will be the total of the amounts you win in all Phase II lotteries. The amount of money you earn will depend partly upon luck and partly upon whether you have made good decisions in Phase I. Notice that the more points you earn in Phase I, the more likely you will win in Phase II. Are there any questions?

Instructions for Finitely Repeated PD Game

General

You are about to participate in an experiment in the economics of decision making. If you follow these instructions carefully and make good decisions, you might earn a considerable amount of money which will be paid to you in cash at the end of the experiment.

The experiment will consist of a series of separate decision making periods. Each period consists of two phases. In Phase I you will be paired with another person and, based upon your combined actions, you will be able to earn *points*. In Phase II, you will have the opportunity to earn dollars based upon the points you earn in Phase I. We begin by describing Phase II so that you understand how the points you earn affect the number of dollars you earn. Then, we describe Phase I in detail so that you understand how to earn points.

Phase II Instructions

At the end of Phase I, you will have earned between 0 and 1000 points according to the rules we will discuss below. The number of dollars you earn in Phase II will depend partly on the number of points you earned in Phase I and partly on chance. Specifically, we have a box which contains lottery tickets numbered 1 to 1000. In Phase II, a ticket will be randomly drawn from the box. If the number on this ticket is **LESS THAN OR EQUAL TO** the number of points you earned in Phase I, you **WIN \$1.00**. If the number on this ticket is **GREATER THAN** the number of points you earned in Phase I, you **WIN \$0.00**. For example, if you have 600 points, you will have a 60% chance of winning \$1.00. Notice that the more points you have, the larger will be your chance of winning the \$1.00 prize.

Phase I Instructions

In each decision making period, you will be paired with another person. One of you will be designated the Row player and the other will be designated the Column player. At the beginning of the period, both the Row player and Column player must separately and independently select an action. The combined actions of the Row player and the Column player jointly determine the number of points earned by the Row player and the number of points earned by the Column player.

Except for the last person to arrive for the experiment, player 11, all of you will be participating in three separate sessions during today's experiment. Player 11 will only participate in the first session. The first session will be 11 periods; you will be randomly paired with each person once —as either the Row player or the Column player—and you will not participate in one period. The second session will be 10 periods; however, you will be paired with the same person during the entire session. The third session will

also be 10 periods and you will be paired with the same person during the entire session. However, the person you are paired with in the third session will differ from the one you were paired with in the second session. In all sessions, you will not know the identification of the person you are playing against in any period. Similarly, nobody in your decision making pair will know your identification in any period. Further, you will not be told who these people are either during or after the session.

In the first or current session, you will alternate from being the Row player to being the Column player from one period to the next. Since there is not an even number of people participating in this experiment, you will occasionally be required not to participate during a particular period. When this is the case, you will receive a message on your terminal which states:

“FOR PERIOD _____, YOU ARE SITTING OUT.”

In the periods in which you are participating you will receive a message stating:

“FOR PERIOD _____, YOU ARE A ROW PLAYER.”

or

“FOR PERIOD _____, YOU ARE A COLUMN PLAYER.”

In your folder you will find a record/profit sheet. On this sheet you will indicate, based on the message previously received on your terminal, whether you are a Column player, a Row player, or sitting out this period.

The points that you earn in each period will be determined by the rules given below.

Specific instructions to Row player. In this part of the instructions we will be referring to specific numbers of points. These numbers are the same as you will be using in the first session of today's experiment.

In those periods in which you are the Row player, you and the Column player must separately and independently decide on actions which will jointly determine the number of points earned by you and the number of points earned by the Column player. As the Row player, you may choose either action R1 or action R2. Similarly, the Column player may choose action C1 or action C2. The number of points earned by you is given by the following table for each pair of actions you and the Column player might select:

Number of points earned by Row player

		Column's Action	
		C1	C2
Row's Action	R1	350	1000
	R2	0	800

To read this table, suppose that you chose action R2 and the Column player chose action C1. You would then earn 0 points. Similarly, suppose that you chose action R1 and the Column player chose action C2. You would then earn 1000 points. In a like manner, you can use this table to determine the number of points you would earn for all other pairs of actions you and the Column player may select. Column players also earn points depending upon the type of action they select. These are given in the next section of the instructions.

When you select an action, enter the action chosen into the computer via your terminal and record the action chosen on your record sheet. Once both you and the Column player have selected your actions and entered them into the computer via your terminals, the computer will determine the number of points earned by you based on the table given above. The result is then sent to you via your terminal. The message will look like the one below:

PERIOD POINTS ARE _____.

At the end of the period, you are to record your point earnings for Phase I on your record sheet. Make sure you check your earnings in points against the computer's calculations. The computer will also inform you about the action taken by the column player. Make sure you record this information on your record sheet.

Specific instructions to Column player. In those periods in which you are the Column player, you and the Row player must separately and independently decide on actions which will jointly determine the number of points earned by you and the number of points earned by Row player. As the Column player, you may either choose action C1 or action C2. The number of points earned by you is given by the following table for each pair of actions you and the Row player might select:

Number of points earned by Column player

		Column's Action	
		C1	C2
Row's	R1	350	0
Action	R2	1000	800

To read this table, suppose that the Row player *B* chose action R2 and you chose action C1. You would then earn 1000 points. Similarly, suppose that the Row player chose action R1 and you chose C2. You would then earn 0 points.

When you select an action, enter the action chosen into the computer via your terminal and record the action chosen on your record sheet. Once both you and the Row player have selected your actions and entered them into the computer via your terminals, the computer will determine the number of points earned by you based on the table given

above. The result is sent to you via your terminal. The message will look like the one below:

PERIOD POINTS ARE _____.

At the end of the period, you are to record your point earnings for Phase I on your record sheet. Make sure you check your earnings in points against the computer's calculations. The computer will also inform you of the action taken by the Row player. Make sure you record this information on your record sheet.

Phase II recording rules. After completing your Phase I record sheet for a given decision making period, you are to use your profit sheet to record the dollars you earn in Phase II. First, record your Phase I point earnings in the row corresponding to the number of the period that is currently being conducted. The person who sat out in this period will then be asked to draw a lottery ticket from the box. Before he/she returns the ticket to the box, the number on the ticket will be announced. You should record the number of the ticket in the second column of your profit sheet. If the number drawn IS LESS THAN OR EQUAL TO the number of points you earned in phase I, circle \$1.00 in the next column; otherwise circle \$0.00 in that column. Pay careful attention to what you circle. Any erasure will invalidate your earnings for the period. If you do make a mistake and circle the wrong number, call it to the experimenter's attention.

At the end of the session, add up your total profit in dollars and record this sum in row 23 of your profit sheet. All dollars on hand at the end of the session in excess of \$0.00 dollars are yours to keep. Subtract this number, which is on row 24, from your total dollars in row 23 and record this difference on row 25. This is the amount of dollars you have earned in this session.

In summary, your earnings in the experiment will be the total of the amounts you win in all Phase II lotteries. The amount of money you earn will depend partly upon luck and partly upon whether you have made good decisions in Phase I. Notice that the more points you earn in Phase I, the more likely you will win in Phase II. Are there any questions?

Session II

This session of the experiment will again consist of a series of separate decision making periods. Each period will gain consist of two phases. In Phase II you will be able to earn dollars based upon the points you earned in Phase I in exactly the same way you did in the first session. In Phase I you will again be paired with another person and, based upon your combined actions, you will be able to earn points.

However, Phase I in this session differs from Phase I in the previous session in the following ways:

- (1) There will be 10 periods in this session.
- (2) There will be 10 players in this session.
- (3) At the beginning of the session, each player will be assigned as either the Row player or the Column player. Once assigned as the Row player, he/she will be the Row player for the entire session. Once assigned as the Column player, he/she will be the Column player for the entire session.
- (4) At the beginning of the session, each player will be randomly paired with another player. These two players will be paired together for the entire session—10 periods.

Session III

This session of the experiment will again consist of a series of separate decision making periods. Each period will again consist of two phases. In Phase II you will be able to earn dollars based upon the points you earned in Phase I in exactly the same way you did in the first two sessions. Similarly, in Phase I you will again be paired with another person and, based upon your combined actions, you will be able to earn points.

Phase I in this session differs from Phase I in the previous session in the following ways:

(1) Row players in the previous session will be Column players in this session. Column players in the previous session will be Row players in this session.

(2) Each player will be randomly paired with a different player than he/she was paired with in the previous session.

REFERENCES

- ANDREONI, J. (1989). "Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence," *J. Polit. Econ.* **97**, 1447–1458.
- ANDREONI, J., AND MILLER, J. (1991). "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence," Working Paper, Social Systems Research Institute, University of Wisconsin.
- BERGSTROM, T., BLUME, L., AND VARIAN, H. (1986). "On the Private Provision of Public Goods," *J. Public Econ.* **29**, 25–50.
- COOPER, R., DEJONG, D. V., FORSYTHE, R., AND ROSS, T. W. (1990). "Selection Criteria in Coordination Games," *Amer. Econ. Rev.* **80**, 218–233.
- DAWES, R. M. (1980). "Social Dilemmas," *Ann. Rev. Psych.* **31**, 169–193.
- DAWES, R. M., AND THALER, R. (1988). "Anomalies: Cooperation," *J. Econ. Perspectives* **2**, 187–197.
- FORSYTHE, R., HOROWITZ, J., SAVIN, N., AND SEFTON, M. (1994). "Replicability, Fairness and Pay in Experiments with Simple Bargaining Games," *Games Econ. Behav.* **6**, 347–369.
- FUDENBERG, D., AND MASKIN, E. (1986). "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica* **54**, 533–554.
- KAHN, L., AND MURNIGHAN, J. K. (1993). "Conjecture, Uncertainty, and Cooperation in Prisoners' Dilemma Games: Some Experimental Evidence," *J. Econ. Behav. Organ.* **22**, 91–117.
- KANDORI, M., (1988). "Social Norms and Community Enforcement," mimeo, Stanford University.
- KREPS, D., MILGROM, P., ROBERTS, J., AND WILSON, R. (1982). "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma," *J. Econ. Theory* **17**, 245–252.
- McKELVEY, R., AND PALFREY, T. (1992). "An Experimental Study of the Centipede Game," *Econometrica* **60**, 803–836.
- ROTH, A. E., (1988). "Laboratory Experimentation in Economics: A Methodological Overview," *Econ. J.* **98**, 974–1031.
- ROTH, A. E., AND MALOUF, M. W. K. (1979). "Game-Theoretic Models and the Role of Bargaining," *Psych. Rev.* **86**, 574–594.
- ROTH, A. E., AND MURNIGHAN, J. K. (1978). "Equilibrium Behavior and Repeated Play of the Prisoner's Dilemma," *J. Math. Psych.* **17**, 189–198.

- SELTEN, R., AND STOECKER, R. (1986). "End Behavior in Sequences of Finite Prisoner's Dilemma Supergames," *J. Econ. Behav. Organ.* **7**, 47-70.
- TOWNSEND, R. (1980). "Models of Money with Spatially Separated Agents" in *Models of Monetary Economies* (J. Kareken and N. Wallace, Eds.). Minneapolis: Federal Reserve Bank of Minneapolis.