

Cooperation over finite horizons: a theory and experiments*

Attila Ambrus[†] and Parag A. Pathak[‡]

First Version: July 2006 This Version: September 2009

Abstract

This paper shows that the presence of different types of players – those who only care about their own material payoffs, and those who reciprocate others’ contributions – can explain the robust features of observed contribution patterns in public good contribution games, even without the presence of asymmetric information. We show that if reciprocity functions satisfy some regularity conditions, then generically there is a unique perfect equilibrium, in which contributions are decreasing. In this equilibrium, selfish players contribute to induce subsequent contributions by reciprocal players, and this incentive diminishes as the end of the game approaches. The model explains the puzzling restart effect and is consistent with various other empirical findings. We also report the results of a series of experiments, using a probabilistic continuation design in which after each set of 10-period games, the group is restarted with low probability. We find specific support for the theory in our data, including that selfish players (identified exogenously) stop contributing earlier than reciprocal players, as directly implied by the model.

*We would like to thank Satoru Takahashi for a careful reading of the paper, and Colin Camerer, Ernst Fehr, Drew Fudenberg, Thomas Konig, Imran Rasul, Al Roth, and participants of the SITE meeting on Experimental Economics for useful comments. We would also like to thank Jim Andreoni and two anonymous referees for helping us improve the paper. The experiment was programmed and conducted with z-Tree; we are grateful to Urs Fischbacher for making this software available. Pablo Guillen and the staff of the Computer Lab for Experimental Research at Harvard Business School were helpful in conducting the laboratory sessions. For financial support, Pathak thanks the National Science Foundation and the Division of Research at Harvard Business School.

[†]Department of Economics, Harvard University, Cambridge, MA 02138, ambrus@fas.harvard.edu, <http://post.economics.harvard.edu/faculty/ambrus/ambrus.html>.

[‡]Department of Economics, MIT, Cambridge, MA 02142, ppathak@mit.edu <http://econ-www.mit.edu/faculty/ppathak>.

1 Introduction

There are many economic and social relationships in which participants interact with each other repeatedly, but for only a commonly known finite time. For instance, consider manager-employee relationships in which the employee has a nonrenewable fixed period employment contract, or faces mandatory retirement at a known time. In these situations, if at every stage participants face a myopic incentive to choose a noncooperative action, then no matter how long the interaction lasts, backwards induction implies that the unique subgame perfect equilibrium outcome involves playing the noncooperative action every period.

In contrast to this theoretical prediction, there is considerable experimental evidence that repeated interaction facilitates cooperation, even with a known finite horizon. Moreover, in many different environments, the same pattern of play emerges: there is a relatively large amount of cooperation in early rounds, which eventually breaks down, and by the end of the interaction, there is very little cooperation. This phenomenon is documented in finitely repeated prisoner's dilemma games (Selten and Stoecker (1986), Andreoni and Miller (1993), Kahn and Murnighan (1993)), repeated duopoly games (Selten, Mitzkewitz, and Uhlich (1997)), centipede games (McKelvey and Palfrey (1992)), and repeated voluntary public good contribution games (in many papers, starting with Isaac, Walker, and Thomas (1984) and Kim and Walker (1984)). Outside the laboratory setting, Bandiera, Barankay, and Rasul (2006) find evidence from a field experiment that temporary workers working together for a fixed period of time cooperate under a relative incentive scheme by holding back their efforts, provided that they can monitor each other's decisions.

In this paper, we present a theory that explains the documented pattern of cooperation in games with finite horizon, and we report the results of a series of experiments designed to test specific implications of the theory.

We conduct the analysis in the context of voluntary public good contribution games.¹ Players in these games decide how much of their endowment to allocate to a public good and how much they retain for private investment. Payoffs are such that, in a one-shot game, any contribution to the public good is a dominated strategy, but the aggregate benefit of contributing to the public good exceeds the individual benefit of a private investment.

In finitely repeated versions of the game, the decreasing cooperation phenomenon is prevalent: in final rounds contributions are very small.² The latter observation also holds for games

¹These games are widely studied in other social sciences, too. Fiske (1992) and Field (2002) are relatively recent references in sociology and anthropology, respectively. See also the references in Chapter 6.1 of Plott and Smith (2008). For a survey of the early economics literature, see Ledyard (1995).

²According to Fehr and Schmidt (2002), in the final period, roughly 75% of the subjects contribute nothing to the public good, and the rest contribute very little. Andreoni (1988) and Andreoni and Petrie (2004) find that roughly 70% of the subjects contribute nothing, while the rest of the subjects contribute about 28% of their endowment, on average. Contributions in final rounds are often concentrated in a few groups, indicating that there are a small number of groups with many or all reciprocators.

played by experienced subjects.³ Furthermore, it remains true even in games that follow a “surprise restart” announcement. Namely, at the end of a repeated public good contribution game, if a surprise announcement is made that the same group of subjects will play another repeated game, contributions initially jump back to a relatively high level and then decrease again over time. This “restart effect” was first reported by Andreoni (1988) and was later confirmed by several other studies.⁴ In contrast, existing evidence suggests that in one-shot games, contributions diminish when players become more experienced.⁵

These findings seem to be incompatible with the assumption that all subjects are rational in the traditional sense, that they care only about their own material payoffs.⁶ Theories proposed so far to explain the observed patterns of contributions can be divided into three rough categories: (i) players make mistakes, but they learn over time; (ii) players do not just maximize their own monetary payoffs, but their preferences involve altruism, reciprocity, or “warm-glow effects” (Andreoni (1989)); (iii) reputational considerations along the lines of Kreps and Wilson (1982) and Kreps, Milgrom, Roberts and Wilson (1982).⁷ However, it is not clear whether any of the above theories can explain all the robust findings of the experiments. Theories of altruism, reciprocity, or warm-glow effects can explain positive contribution levels, but do not imply a decreasing pattern of contributions. Learning can explain positive contribution levels and the decreasing level of contributions, but not the restart effect: if players learn not to contribute by the end of the game, then contributions should not jump back to a high level. The same holds for models of reputation, as long as reputation completely wears out before the restart (that is, selfish players towards the end of the finitely repeated game play differently than unselfish ones, and hence identify themselves).⁸

Our theory explains the initially positive but decreasing pattern of contributions as an equilibrium phenomenon. The basic feature of the model is that we assume the presence of both selfish players (in the sense of maximizing only their own material payoffs), and players who reciprocate contributions by others. This is in accordance with the observations of Fischbacher, Gächter, and Fehr (2001), Brandts and Scram (2001), Palfrey and Prisbey (1997), Ledyard (1995), and Saijo and Yamaguchi (1992) that roughly half of the subjects in public good experiments maximize individual payoffs, while 40-50% of them are conditional cooperators.⁹

³See Isaac and Walker (1988).

⁴See for example Croson (1996).

⁵See Isaac, Walker, and Thomas (1984) and Andreoni (1988).

⁶Fudenberg and Levine (1997) argue that in several well-known experiments, subjects’ observed behavior is consistent with self-confirming epsilon-equilibrium, but they mention that in voluntary contribution experiments, players’ losses seem to be too large for this to be the case.

⁷See Camerer and Weigelt (1988) for a finite borrower-lender game with reputation and Camerer (2003) for an overview of some related experiments.

⁸It is possible that the restart effect can be explained by a model in which reputation does not completely wear out by the end of a finitely repeated contribution game, although we are not aware of any work along this line in the literature.

⁹See for example Fehr and Schmidt (2002).

Motivated by existing experimental evidence, we assume that the reciprocal players reciprocate both past realized contributions and current expected contributions by others.¹⁰ The reciprocal preferences we adopt can be derived from various underlying motives, including fairness considerations, conditional altruism, or following some social norm.

In the resulting model, selfish players can influence future contributions of reciprocal players, and the more periods are left, the higher the increment they can induce on these contributions. This makes it worthwhile for them to contribute more of their endowment to the public good at the beginning of the game. In equilibrium, reciprocal players correctly anticipate these high contribution levels in early periods, which induces them to also contribute. As the game progresses, selfish players have less incentive to contribute, and in equilibrium, their contributions to the public good decrease. Lastly, decreasing contributions by the selfish players imply decreasing contributions by the reciprocal players.

To keep the model simple, we abstract away from asymmetric information and learning, and show that the main experimental findings can be explained even in a model of complete information. The conditions that reciprocity functions has to satisfy for this are arguably strong. However, they can be substantially relaxed to obtain weaker results (such as that ultimately the contribution pattern becomes decreasing).

Although we do not extend our model to an asymmetric information environment, we expect that similar qualitative conclusions would hold as in our model, possibly with weaker assumptions on the reciprocal players' preferences. This is because introducing asymmetric information facilitates selfish players to mimic reciprocal players at the beginning stages of a finitely repeated game. This reinforces high initial contributions and a decreasing pattern afterwards.¹¹

Our model also explains the restart effect. Since a major factor in determining equilibrium contributions is the number of remaining periods in the game, a surprise announcement of playing additional periods increases equilibrium contributions. In addition, the model predicts no restart effect if the restart does not come as a surprise (if the subjects know all along that a second session will be played), which is confirmed by Burlando and Hey (1997). Related to this point, the model we present implies that in a longer game, contributions to the public good are more persistent, another empirical result reported by several papers.¹²

The experiments we conducted were designed to test various assumptions and implications of our model. Each experiment is based on repeated sets of 10-period linear public good

¹⁰Sonnemans, Scharm, and Offerman (1999) (also Keser and Winden (2000)) find both forward-looking and backward-looking (adaptive) behavior in public good experiments.

¹¹These considerations are consistent with the results of Muller et al. (2008), who in a design involving soliciting contingent action plans at the beginning of the game, show that both strategic considerations and learning contribute to the gradual decline of contributions - although strategic reasons seem to have a much greater effect.

¹²See for example Isaac, Walker, and Williams (1994).

contribution games. To be able to conduct multiple “surprise” restarts, we used a design in which after each set of 10-period games, it was randomly decided whether the group stayed together and played a “restarted” game (25% probability) or whether players in the group were randomly reshuffled and played the next set of 10-period games with a new group (75% probability).

To conform the complete information assumption of our model, in all of our empirical investigation we focus on repeated games in which players are already experienced with each other, either because the game is a restarted one with the same participants, or because before the game we show group members previous choices of each other in a decision revealing reciprocal behavior. We find that, in accordance with our assumptions in the theory part, most selfish subjects start out a new repeated game by contributing positive amounts, even if they were already revealed to be selfish to their group members before the game started. This suggests that they expect their contributions to be reciprocated, even though other players know that their motives are selfish.

First we investigate whether the behavior of experienced players can indeed be described as an equilibrium of the game, by testing whether experienced players correctly foresee the contribution pattern in the game. We conducted a session where we solicited player’s forecasts of others’ subsequent contributions before the start of a 10-period game. Experienced players’ forecasts on average closely track the average of actual play, with the median forecast near the median of actual play. In particular, before the start of the restarted game, 69% of subjects anticipate that there will be no contribution in the tenth game, yet each of these subjects contribute a positive amount in the first game, contributing more than half their endowment on average.

Another crucial component of our model is that initially even completely selfish players find it worthwhile to contribute positive amounts, because it induces future contributions by others. To demonstrate this, we regress players’ second-period contributions on others’ first-period contributions. In restarted games, in our OLS specification, a unit increase in average contributions of others increased a player’s second-period contribution by a statistically significant 0.73 units. This result is compatible with the implication that in equilibrium, completely self-interested players start out contributing.

Finally, we test a direct implication of our model, namely that selfish players stop contributing earlier than reciprocal players. We do this using data from sessions in which the 10-period public good contribution games are preceded by a sequence of gift-giving games. We used these gift-giving games to identify selfish and reciprocal players based on their offer as second proposers. We find that the average reciprocal player stops contributing between period 7 and 8, while the average selfish player stops contributing between period 5 and 6, and this difference is statistically significant.

2 Related literature

A number of papers examine strategic interaction between selfish and reciprocal players. Offerman, Sonnemans, and Schram (1996), Anderson, Goeree, and Holt (1998), and Cox (2007) analyze public good contribution games, but without dynamic strategic considerations.¹³ Fehr and Schmidt (1999), and Andreoni and Samuelson (2006) investigate two-period games with multiple types. Brown, Falk, and Fischbacher (2004) study a finite-horizon contracting game in which a high enough fraction of workers are fairness-minded. Duwfenberg and Kirchsteiger (2004) extend reciprocity equilibrium (Rabin (1993)) to extensive form games.

There are various explanations of cooperation over finite horizon in which the focus is not on the dynamic interaction of different types of players. Radner (1980, 1986) shows that cooperation can be maintained for a while in a repeated oligopoly game and in a repeated prisoner's dilemma if players only care about maximizing their payoff up to epsilon precision, even for small values of epsilon. A somewhat similar argument is presented by Klumpp (2004) for repeated public good contribution games. He shows that even a relatively small amount of altruism can generate large contributions at the beginning of the game. The scope of these explanations are limited by the fact that in games with discrete action spaces, small departures from maximizing individual payoffs cannot explain any amount of cooperation, unless the number of periods is very large, while large deviations do not explain the breakdown of cooperation in the end. Neyman (1985) shows that cooperation can be achieved in equilibrium of a finitely repeated prisoner's game if players can only use strategies with bounded complexity. Jehiel (2005) presents a behavioral solution concept for multi-stage games, analogy-based expectation equilibrium, in which some players cannot distinguish between different stages of the game, and just best-respond to the average behavior of other players. The paper does not analyze repeated public good contribution games, but in somewhat similar contexts (like centipede games) it shows that the solution concept allows for initial cooperation between players. Finally, there are various explanations which relax the assumption that the fundamentals of the game are common knowledge among players. Kreps and Wilson (1982), Kreps, Milgrom, Roberts, and Wilson (1982), Sobel (1985), and Fudenberg and Maskin (1986) show how a small amount of uncertainty about payoffs (reputation) can induce cooperation in games with finite horizon, while Neyman (1999) points out that a small departure from the length of the game being common knowledge can lead to cooperative outcomes.

¹³See also Levine (1998). Outside the class of public good contribution games, Cox et al. (2007) introduce a parametric model of other-regarding preferences, and structurally estimate it for simple games like ultimatum games.

3 Model

We consider a T -period public good contribution game with $N \geq 2$ players. The model below could be extended in a straightforward manner to games in which for any strategy profile by others, a player's individual best response action and the action that maximizes the joint payoff of the players are always ordered the same way (for example, the former is always smaller than the latter). This way our qualitative results can be extended to a broad class of games, including repeated prisoner's dilemmas and repeated oligopoly games. Nevertheless, for ease of exposition we stick to the framework of public good contribution games.

Besides denoting the number of players, we also use N to denote the set of players whenever it does not cause confusion. In each period, each player has an endowment of 1 unit. Players in each period simultaneously decide how much of their endowment to contribute for public investment and how much to retain for private investment. Let $x_i^t \in [0, 1]$ denote player i 's contribution to the public investment in period t . After each period each player observes the contributions by all players.

The material payoff of player i in period t is the amount of endowment she retains for herself plus her share from the aggregate returns to the public investment:

$$(1 - x_i^t) + \frac{A}{N} \sum_{j \in N} x_j^t, \quad \forall i \in R.$$

Public investment yields a constant marginal return A , which is divided equally to all players. We assume $A > 1$ but $\frac{A}{N} < 1$.

Players $i = 1, \dots, S$ maximize the sum of their per period material payoffs. From now on, we refer to them as selfish players. Let S also denote the set of selfish players and R denote the rest of the players. Players $R = \{S + 1, \dots, N\}$ are reciprocal. Their payoffs are determined through their reciprocity functions. It is convenient to think about these functions as specifying target contribution levels. The arguments of f_i^t , the period- t reciprocity function of player $i \in \{S + 1, \dots, N\}$, are past and current contributions to the public good by others. We assume that every reciprocity function is nondecreasing in all other players' contributions and takes values in $[0, 1]$.

To keep the analysis tractable, we only consider reciprocity functions which are additively separable with respect to contributions made at different periods, and which, within the same period, are additively separable with respect to contributions made by different players:

$$f_i^t((x_j^1)_{j \in N/\{i\}}, \dots, (x_j^t)_{j \in N/\{i\}}) = \sum_{k=1}^t \sum_{j \in N/\{i\}} f_{i,j}^{t,k}(x_j^k)$$

where $f_{i,j}^{t,k}(\cdot)$ is nondecreasing for $i \neq j$ and $f_{i,j}^{t,k}$ is defined only for $k \leq t$. Furthermore,

throughout the paper, we assume that $f_{i,j}^{t,k}$ is concave and differentiable for every $i \in R$, $j \in N/\{i\}$ and $t, k \in \{1, \dots, T\}$.

The t -period payoff of player i is $g(x_i^t - f_i^t()) = -(x_i^t - f_i^t())^2$.¹⁴ Reciprocal players maximize the sum of these per period payoffs.

The game is a standard extensive form game. Reciprocal players differ from the selfish players in that their payoffs do not just depend on their own material payoffs. We assume that the game is of complete information. Players know how many selfish and how many reciprocal players there are in the group, and they know the reciprocity functions.¹⁵ This corresponds to our intention to analyze play in games in which the players are experienced and become familiar with each other.

Reciprocity functions provide a simple and tractable way to model players whose preferences are influenced by how much other players contribute. Our framework allows for the underlying motivation behind reciprocal preferences to come from many sources. One possibility is that reciprocal players are conditionally altruistic. Another source can be a desire to follow social norms: a player with such considerations would contribute more if she thinks others contribute more (or if she observes that others contributed a lot in past rounds). Yet another possibility is that reciprocal players care about fairness or equality, as in Fehr and Schmidt (1999) or Bolton and Ockenfels (2000).¹⁶

Our model specification is also consistent with the consideration that reciprocal players care not only about how much others contribute, but also about the types, or preferences, of the other players. In particular, we allow reciprocity towards other players to be asymmetric. For example, reciprocity functions are allowed in our model to be more responsive to other reciprocal players' contributions than to contributions made by selfish players. Therefore, as long as intentions of a player depend only on his or her preferences, our model allows for reciprocity to depend on the intentions behind contributions.

4 Theoretical results

In this section we impose some regularity conditions on reciprocity functions, and investigate the subgame perfect Nash equilibrium of the game.

4.1 Assumptions on reciprocity

We impose the following five assumptions on reciprocity functions.

¹⁴Instead of the quadratic loss function specified above, we could consider any strictly quasiconcave function $g()$ which attains its maximum at 0.

¹⁵More precisely, all this information is common knowledge among players.

¹⁶A player who is concerned about fairness reciprocates others' contributions in a public good contribution game because a contribution increases others' flow payoff at the expense of one's own payoff.

A1: Linear Reciprocity toward Reciprocal Players: $f_{i,j}^{t,k}(x_j^k) = \alpha_{i,j}^{t,k} x_j^k \forall i \in R, j \in R/\{i\}$ and $t, k \in \{1, \dots, T\}$, where $k \leq t$.

A2: Nonincreasing Total Impact of Contributions over Time: Suppose $i \in R, j \in N/\{i\}$, $t \in \{1, \dots, T-1\}$, and x_j^1, \dots, x_j^{t+1} is such that $x_j^k \geq x_j^{k+1} \forall k \in \{1, \dots, t\}$. Then $\sum_{k=1}^{t+1} f_{i,j}^{t+1,k}(x_j^k) \leq \sum_{k=1}^t f_{i,j}^{t,k}(x_j^k)$.

A3: Nonincreasing Marginal Impact of Contributions over Time: For any $x \in [0, 1]$, $i \in R$, $j \in N/\{i\}$, $t < t'$, $k \geq 0$, and $t' + k \leq T$, $\frac{\partial f_{i,j}^{t+k,t}(x)}{\partial x} \geq \frac{\partial f_{i,j}^{t'+k,t'}(x)}{\partial x}$.

A4: Positive Initial Reciprocity: For every $i \in R, j \in N/\{i\}$, and $t \in \{1, \dots, T\}$, there exists $t' \in \{1, \dots, t\}$ such that $\left. \frac{\partial f_{i,j}^{t,t'}(x)}{\partial x} \right|_{x=0} > 0$.

A5: No Overreciprocation: $\sum_{k=1}^t \sum_{j \in N/\{i\}} \frac{\partial f_{i,j}^{t,k}(x_j^k)}{\partial x_j^k} \leq 1 \forall t \in \{1, \dots, T\}$, $i \in R$, and $(x_j^k)_{j \in N/\{i\}}^{k=1, \dots, t} \in [0, 1]^{(N-1)t}$.

A1 guarantees that the impact of a marginal contribution by a selfish player on subsequent contributions by reciprocal players does not depend on the history of contributions. This both greatly simplifies the analysis and helps avoid multiplicity of equilibria. A2 is a condition on the level of reciprocity: it states that a reciprocal player does not reciprocate in a strictly increasing manner a nonincreasing sequence of contributions by any other player. This assumption in particular implies that reciprocity toward a given (one-time) contribution weakly decreases over time. A3 is a statement on marginal reciprocity: it requires that the marginal impact of a contribution on reciprocity k periods later (a contribution at t on reciprocity at $t+k$, a contribution at $t+1$ on reciprocity at $t+k+1, \dots$) does not increase over time.¹⁷ A4 imposes that reciprocity is initially strictly positive toward every other player. A5 implies that at any period t , a unit increase in contributions by other players up until t increases the value of a reciprocity function by not more than a unit. This, besides being a natural requirement, is imposed in order to avoid multiplicity of equilibria resulting from reciprocal players either having optimistic expectations with respect to each others' contributions and contributing more, or having pessimistic expectations and contributing less.

The assumptions above are strong, but each of them are necessary to obtain the strong result we derive in the next section, namely that generically the resulting game has a unique subgame-perfect Nash equilibrium that implies a decreasing pattern of contributions. However, weaker results along the same lines can be established even when dropping some of the assumptions

¹⁷We note that although the above assumptions on the time structure of reciprocity are sufficient to establish our main results, the reciprocity functions that are most appealing to us are ones in which reciprocity is in some sense constant over time. One way to formalize constant reciprocity over time is strengthening A2 by requiring that $\sum_{k=1}^{t+1} f_{i,j}^{t+1,k}(x_j^k) = \sum_{k=1}^t f_{i,j}^{t,k}(x_j^k)$ whenever $x_j^k = x_j^{k+1} \forall k \in \{1, \dots, t\}$ and $j \in N/\{i\}$.

above. For example, since for any reciprocity functions it holds that selfish players do not contribute in the very last period, A2 and A3 can be substantially weakened to show that the equilibrium contribution pattern is ultimately decreasing (even if initially it might not be). Similarly, A5 can be dispensed with if the goal is to show the existence of some subgame-perfect Nash equilibrium implying a decreasing pattern of contributions (as opposed to the uniqueness of such equilibrium).

4.2 Decreasing pattern of contributions in equilibrium

Our main result is that if the assumptions from the previous subsection hold, and there is at least one selfish player in the game, then generically there is a unique subgame perfect Nash equilibrium, and that it exhibits a decreasing pattern of contributions.

Theorem 1: If $S \geq 1$ and A1-A5 hold, then for generic A , the public good contribution game has a unique subgame perfect Nash equilibrium, and this equilibrium exhibits a weakly decreasing pattern of contributions. If reciprocity functions are strictly concave in selfish players' contributions, then the above statement holds for all A .

For the formal proof of this result, as well as all formal statements in our paper, see the online appendix Ambrus and Pathak (2009). The brief intuition for the result is as follows. Since reciprocity functions are concave, no overreciprocation and positive initial reciprocity imply that the impact of a contribution on future contributions by reciprocal players is uniquely defined. Linearity in other reciprocal players' contributions implies that this impact is independent of contributions made in other periods or by other players. Then for generic values of A (if reciprocity functions are strictly concave in selfish players' contributions, then for all values of A), selfish players' contributions are uniquely determined in subgame perfect equilibrium at every period. Nonincreasing marginal impact of contributions over time then implies that the marginal return of contributions at earlier periods, when more periods are left to be played, is higher. This establishes that selfish players' contributions are weakly decreasing over time. Finally, nonincreasing total impact of contributions over time implies that the reciprocal players' contributions are weakly decreasing over time, too.

4.3 Consistency with existing experimental findings

Our model is consistent with a series of results in the existing literature on public good contribution experiments. Here we summarize these implications of our model informally. For corresponding formal results, see the online appendix Ambrus and Pathak (2009).

First, consider increasing the return of the public investment (A), which brings the individual return from contributing to the public good ($\frac{A}{N}$) closer to the return from private investment (1). It is well-documented in experimental settings that this increases players' contributions

to the public good (see for example Isaac and Walker (1988), and Isaac, Walker, and Williams (1994)). This effect is captured by our model: increasing the marginal individual return from contributing increases contributions by the selfish players. Through positive reciprocity this also increases contributions by the reciprocal players (Proposition 1 of Appendix).

Another robust experimental finding is that increasing the number of periods in public good contribution games is shown to result in a longer period of positive contributions and in higher aggregate contribution levels (see Isaac, Walker, and Williams (1994)). This effect is also implied by our model. In a longer game, selfish players have more incentives to contribute, because there are more future periods in which reciprocal contributions are affected. In equilibrium, all players end up contributing more (Proposition 2 of Appendix).

Our model is also consistent with the famous restart effect first shown in Andreoni (1988). If players treat the restarted game as a new game, then our model immediately explains the restart effect: contributions jump back to the initial high level on the equilibrium contribution path. Even if players do not treat the new session as a new game, but as an extension of the first game (and therefore aggregate contributions in the previous game become a history in a longer game), the model is compatible with the restart effect. If it is unexpectedly revealed that more periods are to be played than previously thought, selfish players have increased incentives to contribute. This unambiguously increases the contributions of selfish players.

Finally, in a one-shot game assumptions A1-A5 guarantee that if there is at least one selfish player then there is a unique Nash equilibrium, in which all players contribute zero. The result applies to restarted one-shot games as well. The intuition behind the result is simple. Since there is no continuation, all selfish players contribute zero. Then A4 and A5 imply that the only fixed point of the expectations of reciprocal players is when they expect zero contributions from each other (Proposition 3 of Appendix). This result corresponds to the empirical finding that although initially subjects contribute significantly positive amounts in one-shot games in an experimental setting, contributions seem to go to zero with learning. In a setting in which after each round of play, subjects get randomly assigned to a new group (called the “strangers” treatment in the literature), contributions to the public good diminish over time (see Andreoni (1988), Croson (1996)).

4.4 Contributions by selfish versus reciprocal players

Our model does not give a clear prediction for the relative magnitudes contributions by a selfish and a reciprocal player. However, there is a general implication of the model for the timing of contributions of different players: every reciprocal player contributes positive amounts for at least as many periods as any selfish player. The reason is that selfish players have purely forward-looking considerations in contributing, while reciprocal players are partly backward-looking; hence selfish players’ contributions tend to be relatively more concentrated on earlier

periods than reciprocal players' contributions.

Theorem 2: Suppose assumptions A1-A5 hold and the game has a unique subgame perfect Nash equilibrium s . Then if $s_i^t > 0$ for some $i \in S$, then $s_j^k > 0 \forall j \in R$ and $k \in \{1, \dots, t\}$.

5 Experimental design

5.1 Hypotheses

Besides confronting our model with existing experimental evidence, we conducted a set of experiments ourselves, testing various assumptions and predictions of our model.

Since our model assumes complete information, we designed experiments in which subjects could get experienced both with the game, and with their group members. In one type of treatment we achieved the latter by adopting a probabilistic restart scheme. That is, after every 10 periods, we assigned subjects to new groups with a high probability, but with the remaining small probability we kept subjects in the same group, and restarted the game among them, for an extra 10 periods. This treatment allowed us to investigate how players played the game after multiple restarts, and still enabled us to treat the restart as a surprise. We regarded subjects in later rounds of the experiment who played in a restarted groups as players experienced with both the game and each other. In another treatment first we asked subjects to participate in gift-exchange games. Afterwards, we conducted public good contribution games, and after giving subjects opportunity to get experienced with the game, we assigned them to new groups and showed them the histories of their group members' play as respondents in the gift-exchange games. That is, in this treatment instead of a surprise restart, we gave information to subjects directly about their group members.

In the games with experienced players, we investigate three hypotheses. The first one is that, as opposed to many of the other explanations for cooperation over finite horizons, we claim that positive contributions (and other robust features of contribution patterns) can be an equilibrium phenomenon. Equilibrium in particular implies that players' expectations are fulfilled. Hence, we test how correct experienced players' expectations concerning the contribution pattern of others are, before a 10-period game. Second, in order to explain positive contributions, our model requires that even completely selfish players have incentives to contribute at the beginning of the game. Therefore, we test how contributions in the first period of a game affect contributions of others in the next period. Finally, we test a direct prediction of our model, corresponding to Theorem 4, that selfish players stop contributing earlier than reciprocal ones. We use the treatment in which subjects were first asked to play gift-exchange games to identify selfish and reciprocal type players.

The three hypotheses we formally test are:

- H1: Experienced players' forecasts are close to actual play in restarted 10 times repeated games.
- H2: With experienced players, contributions positively affect other players' subsequent contributions.
- H3: With experienced players, selfish players start contributing zero earlier than reciprocal players.

5.2 Treatments

Table 1 describes the three treatments we used. In each treatment, subjects play a public good game in groups of four, for ten periods with the same group members. The stage game shares features with other experiments who have examined repeated public good contribution games. Each player receives an endowment of 20 tokens, and must simultaneously decide how many tokens to contribute to the public project. After subjects made their contributions, we informed them of the contribution of each of the group member's to the public project and their income from the stage game.¹⁸ The income of a player is the amount that is not contributed to the public project plus the 1.6 times the average contribution of the group. In a "10-period game," the stage game is repeated 10 times with the same group of four players. We will refer to the first set of 10-period games as games 1-10, the second set as games 11-20, and so on.¹⁹ At the end of the session, tokens were converted to dollar amounts.

Experienced Restart: This session consisted of six sets of 10-period public good games where after the 10th period, subjects were reshuffled with probability 0.75. If a subject is reshuffled, then he or she is randomly assigned to a new group. If a subject is not reshuffled, then he or she stays in the same group. A group could not stay together for more than two consecutive 10-period games. Subjects were informed of how reshuffling takes place, and the probability of being reshuffled. In a typical session with 32 subjects, 6 out of the 8 10-period games would be reshuffled on average when the probability is 0.75.

Experienced Restart with Forecast: The only difference with this session and the Experienced Restart session is that at the beginning of every set of 10-period games, we asked subjects to forecast what they expected the average contribution of their three opponents to be in the first, fifth, and tenth period among the 10-period game. In this session, we did not

¹⁸We also ran experiments in which after every period we only told subjects the total contributions by the group, not the individual contributions by each group member. The results, which we reported in the previous working paper version of this paper, are very similar to the results presented here.

¹⁹In each session, all subjects answered control questions which intended to verify their understanding of the instructions. The actual instructions and control questions are available from the authors upon request.

pay an additional amount to subjects based on their forecasts to avoid interfering with the incentives provided by the public good game.

Identifying types: In this session, we used a gift-exchange game to identify player types and then sorted the players into groups while they played the 10-period repeated public good game. In Part I, players were randomly paired for each of six gift-exchange games. Each player took the role of first proposer and second proposer three times and no subjects were ever paired with the same opponent more than once. The first proposer started the game by deciding what fraction of her endowment of 10 tokens to give to the second proposer. This amount was doubled and was given to the second proposer. The remaining portion of the first proposer’s endowment was kept by her (but was not doubled). Following this first offer, the second proposer responded by deciding how much of her endowment of 10 tokens to give to the first proposer, who received double this amount. The remaining portion was kept by the second proposer (but was not doubled).

Play in the gift-exchange games was used to identify the degree of reciprocity in players. We constructed a measure of reciprocity based on how players responded to positive proposals when playing as the second proposer. The index of reciprocity was the ratio of their response to the first proposer’s offer when the first proposer’s offer was nonzero. According to this index, the higher the ratio, the more reciprocal a player is. In particular, we labeled the eight subjects with the highest reciprocity ratios as “reciprocal players” and labeled the eight subjects with the lowest reciprocity ratios as “selfish players.” We labeled the rest of the players in the experiment “unidentified.”

Subjects were not informed of the second part of the session until after completion of the first part. In part II, we asked players to play two sets of ten-period public good games in groups of four, where groups were randomly reshuffled, so that they become experienced with the game. Then, we sorted players into groups with three reciprocal players and one selfish player, and groups with one reciprocal player and three selfish players.²⁰ Before the start of the 10-period game, subjects were informed of the histories of all their opponents as second proposers in the gift exchange game. After the tenth game, subjects were resorted into different groups of three reciprocal players and one selfish player, and groups with one reciprocal player and three selfish players, and played another round of the 10-period game. Before playing a 10-period game, subjects were shown the histories of everyone in the group as second proposers in part I were shown to the players before the game.

5.3 Experimental results

All lab sessions were conducted with zTree (Fischbacher (2007)) at the Computer Lab for Experimental Research at Harvard Business School. Subjects were recruited from the greater

²⁰The unidentified players played among each other in the remaining groups.

Boston area, with a large fraction of university students. No subject was allowed to participate in multiple sessions. Each session lasted approximately 1.5 hours and involved either 32 and 36 subjects. The average earnings across sessions was between \$21 and \$23. We ran one session of Experienced Restart, one session of Experienced Restart with Forecast, and two sessions of Identifying Types.²¹ In total, 140 subjects participated in the sessions.

Main features of contribution paths

Before presenting the results from testing the hypotheses we derived from the model, we describe the basic features of contribution paths in the games we examined (games played by experienced players). The short summary is that we find exactly the same features as other studies: that contributions are significantly different from zero at the beginning, that contributions decrease over time, and that there is a significant restart effect.

Figure 1 plots the average contribution paths in restarted games for the Experienced Restart and Experienced Restart with Forecast sessions.²² In Figure 1, darker lines in the figure refer to sets of ten games played earlier in the session, so the darkest line corresponds to Games 11-20, the first set of restarted games. The figure shows that the path of contributions displays an overall downward trend, with the average contribution in the first period always much larger than the last game. Table 2 reports the mean and median contribution over the course of the 10-period restarted games for these two sessions. In Panel A, we report the distribution of the average time path of individual contributions from the last two sets of 10-period games. In the Experienced Restart treatment, there are 4 restarted groups (N=16), playing a restarted game among Games 41-50 and Games 51-60, while in the Experienced Restart with Forecast treatment, there are 5 restarted groups (N=20). We consider players who are restarted in these games to be experienced players. At this stage in the session, these players will have played at least 4 sets of 10-period games. For both sessions, the mean individual contribution in the first period is over half the endowment, while in the last period the mean individual contribution in the Experienced Restart session is 0, while in the Experienced Restart with Forecast session it is 2.3. For both sessions, in the first period, 16% of subjects contribute zero, while in the last period 88% contribute zero.

In Panel B of Table 2 we present the mean and median individual contributions for the two sets of 10-period games prior to the last set of 10-period games. For both the Experienced Restart and Experienced Restart with Forecast session, we see the same pattern as in Panel A.

²¹In one of the Identifying Types sessions, after each stage game, we displayed the contributions of each of the individuals in the group and the other we displayed the group average. We pool the observations so that we have enough observations to make meaningful statistical comparisons.

²²In Figure 1, we pool observations from restarted games played in both of these sessions because they follow the same qualitative pattern. When we formally test whether the distribution of individual contributions among experienced players in the first game and last game of the 10-period game are from the same distribution using an F-test, we find no difference for restarted Games 41-50 and Games 51-60 ($p=0.51$).

Both panels show that there is a downward path of contributions, and that the majority of the contributions of the individual players are zero in the last period of the restarted game.

Next, we examine the restart effect. Figure 2 plots the pattern of play for a group that is not reshuffled both in the ten periods before the restart, and the ten periods after the restart. The figure shows that there is a restart effect in these games, as the average contribution in the last period of the previous game is much smaller than the average contribution in the first period of the restarted game.

To verify the presence of a restart effect in our probabilistic continuation design, in Table 3, we formally test for the restart effect. Table 3 reports summary statistics of the restart effect across the two sessions by displaying the average contribution before and after restarts. That is, we compare the last period in the set of 10-period games with the first period in the restarted 10-period game.

Table 3 presents two types of comparisons: Panel A reports the last two sets of 10-period games (Games 40 vs. 41 and Games 50 vs. 51), while Panel B reports the intermediate set of 10-period games (Games 20 vs. 21 and Games 30 vs. 31). For instance, in the Experienced Restart session, all players in the 10th period prior to being reshuffled contribute 0, while in the first period of the restart with the same group, the mean individual contribution is 10.94, and only 19% of subjects contribute zero. Both the individual contribution level and number of players who contribute zero are statistically different using a Wilcoxon two-sided test ($p < 0.01$).²³

The next two rows of the table report the same comparison for the Experienced Restart with Forecast session. Here, 60% of subjects contribute 0 in the 10th period before the restart, while only 15% contribute 0 in the first restarted game. Both the individual contribution level and the number of players who contribute zero are statistically different ($p < 0.01$).

In Panel B of Table 3, we report the same comparisons. For the Experienced Restart session, the individual contribution amount and number who contribute zero are statistically different ($p < 0.01$). For the Experienced Restart with Forecast session, the individual contribution levels are statistically different ($p = 0.03$), as are the fraction who contribute zero ($p = 0.03$). Note that for this comparison, there are 3 restarted groups and hence only have 12 data points. We expect higher significance levels with a larger sample size.

The significance of the restart is also seen in comparisons at the group level. To examine this more formally, we pool the data from the two sessions presented in Panel A and run a regression of group contribution on indicators for the session and the first period, and find that the group average contribution in the first restarted game is 7.33 higher than in the previous game ($t = 3.97$). Likewise, when we pool the data presented in Panel B, we find that the group average contribution in the first restarted game is 11.00 higher than in the previous game ($t = 5.35$).

In summary, Tables 2 and 3 as well as Figures 1 and 2 demonstrate that there is a declining

²³All tests reported here, unless otherwise noted are two-sided Wilcoxon rank sum tests.

average contribution path even for experienced players in restarted games and that the restart effect is statistically significant for experienced players.

Next, we consider summary statistics from the identifying types session. There are 8 groups of 3 reciprocal players and 1 selfish player, and 8 groups of 1 reciprocal player and 3 selfish players, for a total of 640 individual observations. The groups with three selfish players exhibit little cooperation relative to groups with three reciprocal players. The average contribution per player for the groups with 3 reciprocal players and 1 selfish player is 9.78, while the average contribution per player for the groups with 1 reciprocal player and 3 selfish players is 4.87 ($p < 0.01$). In the former group type, in each of the 8 groups, the last positive contribution is made by a reciprocal player. In the latter group type with 3 selfish players, in 3 out of 8 groups, the last positive contribution is made by a reciprocal player even though there are 3 times as many selfish players. Across groups, the average selfish player contributes 5.41, while the average reciprocal player contributes 9.24 ($p < 0.01$). In selfish groups, the average selfish player contributes 4.83, while the average reciprocal player contributes 5.01 ($p = 0.57$). In reciprocal groups, the average selfish player contributes 7.18, while the average reciprocal player contributes 10.65 ($p < 0.01$).

Figure 3 shows the time path of contributions for the identifying types sessions for the two sets of 10-period games after the gift exchange game histories are made public to group members. The lines report the average contribution by group composition for each set of 10-period games. A selfish group in the figure is one with 3 selfish players and 1 reciprocal player, while a reciprocal group is one with 3 reciprocal players and 1 selfish player. For all groups, the average group contribution in the first game is much larger than the last game.²⁴ The figure also shows that the average group contribution in groups with 3 reciprocal players is larger than the average group contribution in groups with only 1 reciprocal player.

The summary statistics for the identifying types session and Figure 3 suggest that heterogeneity in player types is an important consideration to understanding the dynamics of the average contribution path in public good contribution games.²⁵

Testing the hypotheses

Table 4 examines how closely player forecasts match actual play. The table focuses on restarted games with experienced players. In these games, players expect the average contribution to be 10.87 in the first of the ten games, and on average, the actual average contribution of the three opponents is 10.44. For the fifth period, the average forecast is 7.06, and the average actual play is 4.82. For the 10th period, the average of player forecasts is 2.19, while the average

²⁴Although not shown, when we plot the average contribution for the groups with unidentified players, we also find a downward average contribution path and restart effect.

²⁵In independent work, Fischbacher and Gächter (2006) document the presence of player types based on directly eliciting individuals' cooperation preference.

actual play is 1.63. Note that for the 10th period, the averages are skewed by the presence of one player who contributes his entire endowment and forecasts that the average contribution is 20 for the 10th period. The difference between the forecast and actual play for the median player is zero, and half of the player’s forecasts are within one unit of the average contribution of their opponents. Moreover, 11 out of 16 players forecast no contribution in the tenth period. However, each of these players contribute a positive amount in the first game, contributing more than half their endowment on average. The findings together indicate that a large fraction of subjects correctly anticipate at the beginning of the game that although initial contributions in the game will be high, contributions at the end of the game will be close to zero. In short, players foresee the declining pattern of contributions.

We next regress forecasted contribution on average contribution of opponents for players in the last two sets of restarted games. We find that for the first game, the coefficient on average contribution of opponents is 0.92, with T-statistic of 5.06 and $R^2 = 0.63$. This suggests that forecasts and actual play are highly correlated. When we regress forecasted play on actual play for game five, the coefficient is 1.05 (T-stat=4.12, $R^2 = 0.53$). In the tenth game, 11 out of 16 players forecast no contribution, and 6 subjects are paired with opponents whose average contribution is zero. The large fraction of zeros weakens our ability to estimate a precise relationship and causes our coefficient estimate to be insignificant. These regressions show that there seems to be no large systematic error in expectations for experienced players in restarted games, which leads to our first conclusion:

Conclusion 1: Experienced players on average correctly anticipate the pattern of contributions in a 10-period game. In particular, they foresee that contributions will be close to zero by the end of the game.

Turning to the second hypothesis, Table 5 considers a measure of the importance of the strategic incentive to contribute. For the 10-period public good games in the experienced restart, and experienced restart with forecast session, the table reports regressions on the correlation of the average first period contribution on subsequent play. The three columns regress an individual’s contribution in the second game on the average contribution in the first period. The table reports the estimated coefficient and T-statistic in brackets of the impact of the average first period contribution. Each row corresponds to a different specification: OLS is simply an ordinary least squares regression with an intercept; session fixed effects include a dummy variable for the session; the next specification includes a fixed effect for each set of 10-period games; and lastly we report estimates from a random-effects tobit model. A dummy variable for the session allows for session specific differences in the estimate, a fixed effect for each set of 10-period games allows for differences across time, while the random effects model adjusts for multiple observations per individual. In the first specification, a unit increase in first period contribution of a player increases others’ contributions in the next period by 0.49

units in the first specification. The magnitude of the effect is larger in restarted games, where a unit increase in the first period contribution of a player increases others' contributions in the next period by 0.73 unit. For each of the models we estimate we find that the coefficient is larger in restarted games than not restarted games.

The positive and significant coefficient suggests the existence of conditional reciprocity even when players are experienced. We note that the amount of responsiveness we found in second period contributions is not enough by itself to induce selfish players to contribute in the first period. However, presumably first-period contributions have an effect on contributions in periods after the second one, too, increasing the selfish players' incentives to contribute at the beginning. We cannot estimate the latter effects from our data, because players' contributions are endogenous in all previous contributions of others and therefore in their own contributions up until two rounds preceding the current round.²⁶

Conclusion 2: For experienced players, contributions in the first round positively affect contributions in the subsequent round.

Lastly, we examine when selfish and reciprocal players stop contributing positive amounts. We have already mentioned that in the groups with a majority of reciprocal players, the last positive contribution is always by a reciprocal player. We can formally test Theorem 4 by comparing the last period in which selfish and reciprocal players give positive contributions. If a player never gives a positive contribution, then the period where she last gives a positive contribution is set to 0. We find that the average reciprocal player stops contributing in period 7.41, while the selfish player stops contributing in period 5.19. The difference is statistically significant ($p < 0.01$). This supports the third hypothesis.

Conclusion 3: In 10-period games played by experienced players, selfish players stop contributing earlier than reciprocal players.

6 Conclusion

This paper shows that all documented findings from finitely repeated public good contribution games can be captured by a model in which there are both selfish and conditionally reciprocal players, even in the absence of asymmetric information. We provide conditions on the preferences of reciprocal players for this to be the case. The model can be extended to various other games, like centipede games and finitely repeated oligopoly games, in which selfish players have an incentive to cooperate at early stages if they anticipate that the cooperation might be reciprocated. One caveat is that in different settings, the domain and the range of

²⁶In restarted games, the correlation coefficient between the total number of units a player contributes in periods 2-10 and the average number of units the others contribute in the first game is 6.37.

reciprocity might change. For example, in centipede games in which at every stage, players face a binary decision whether to continue or terminate the game, reciprocity might be defined on continuation probabilities. We hope to investigate these issues in future work.

7 References

AMBRUS, A. and P. PATHAK (2009): "Supplementary appendix to Cooperation over finite horizons: a theory and experiments," *mimeo, Harvard University*.

ANDERSON, S., J. GOEREE and C. HOLT (1998): "A theoretical analysis of altruism and decision error in public goods games," *Journal of Public Economics*, 70(2), 297-323.

ANDREONI, J. (1988): "Why free ride? Strategies and learning in public goods experiments," *Journal of Public Economics*, 37(3), 291-304.

ANDREONI, J. (1989): "Giving with impure altruism: applications to charity and Ricardian equivalence," *Journal of Political Economy*, 97, 1447-1458.

ANDREONI, J. and J. MILLER (1993): "Rational cooperation in the finitely repeated prisoner's dilemma: experimental evidence," *Economic Journal*, 103, 570-585.

ANDREONI, J. and R. PETRIE (2004): "Public Goods Experiments Without Confidentiality: A Glimpse Into Fund-Raising," *Journal of Public Economics*, 88, 1605-1623.

ANDREONI, J. and L. SAMUELSON (2006): "Building rational cooperation," *Journal of Economic Theory*, 127, 117-154.

BANDIERA, O., I. BARANKAY and I. RASUL (2005): "Social preferences and the response to incentives: evidence from personnel data," *Quarterly Journal of Economics*, 120: 917-962.

BOLTON, G. E. and A. OCKENFELS (2000): "ERC: A theory of equity, reciprocity, and competition," *American Economic Review* 90, 166-193.

BRANDTS, J. and A. SCHRAM (2001): "Cooperation and noise in public goods experiments: applying the contribution function approach," *Journal of Public Economics*, 79, 399-427.

BROWN, M., A. FALK and E. FEHR (2004): "Relational contracts and the nature of market interactions," *Econometrica*, 72, 747-780.

BURLANDO, M. and J. HEY (1997): "Do Anglo-Saxons free-ride more?," *Journal of Public Economics*, 64, 41-60.

CAMERER, C. (2003): *Behavioral Game Theory*, Princeton, NJ: Princeton University Press.

CAMERER, C. and K. WEIGELT (1988): "Experimental Tests of a Sequential Equilibrium Reputation Model," *Econometrica*, 56, 1-36.

COX, J. (2004): "How to identify trust and reciprocity," *Games and Economic Behavior*, 46, 260-281.

COX, J. (2007): "On modeling voluntary contributions to public goods," *Public Finance Review*, 35, 311-332.

COX, J., D. FRIEDMAN and S. GJESTARD (2007): "A tractable model of reciprocity and fairness," *Games and Economic Behavior*, 59, 17-45.

CROSON, R. (1996): "Partners and strangers revisited," *Economics Letters* 25, 25-32.

- DUWFENBERG, M. and G. KIRCHSTEIGER (2004): "A theory of sequential reciprocity," *Games and Economic Behavior*, 47, 268-298.
- FALK, A. and U. FISCHBACHER (2006): "A theory of reciprocity," *Games and Economic Behavior*, 54: 293-315.
- FEHR, E. and K. SCHMIDT (1999): "A theory of fairness, competition, and cooperation," *Quarterly Journal Of Economics* 114, 817-868.
- FEHR, E. and K. SCHMIDT (2002): "Theories of fairness and reciprocity - evidence and economic applications," IN *Advances in Economics and Econometrics - 8th World Congress, Econometric Society Monographs*, Cambridge, Cambridge University Press.
- FEHR, E. and E. TOUGAREVA (1995): "Do high monetary stakes remove reciprocal fairness? Experimental evidence from Russia," *Institute for Empirical Research in Economics, University of Zürich, Working Paper No. 120*.
- FIELD, A. (2002): *Altruistically inclined? The behavioral sciences, evolutionary theory and the origins of reciprocity*, Ann Arbor: UMichigan Press.
- FISCHBACHER, U. (1999): z-Tree - Zurich Toolbox for Readymade Economic Experiments - Experimenter's Manual, Working Paper 21, Institute for Empirical Research in Economics, University of Zurich.
- FISCHBACHER, U. and S. GÄCHTER (2006): "Heterogenous Social Preferences and the Dynamics of Free Riding in Public Goods," IZA Discussion Paper 2011.
- FISCHBACHER, U., S. GÄCHTER and E. FEHR (2001): "Are people conditionally cooperative? Evidence from a public goods experiment," *Economics Letters*, 71, 397-404.
- FISKE, A. P. (1992): "The four elementary forms of sociality: framework for unified theory of social relations," *Psychological Review*, 99 (4), 689-723.
- FUDENBERG, D. and D. LEVINE (1997): "Measuring players' losses in experimental games," *Quarterly Journal of Economics*, 112 (2), 507-536.
- FUDENBERG, D. and E. MASKIN (1986): "The folk theorem in repeated games with discounting and incomplete information," *Econometrica*, 54: 533-554.
- HOFFMAN, E., K. McCABE and V. SMITH (1996): "On expectations and monetary stakes in ultimatum games," *International Journal of Game Theory*, 25, 289-301.
- ISAAC, M. and J. WALKER (1988): "Group size effects in public goods provision: the voluntary contributions mechanism," *Quarterly Journal of Economics*, 103(1), 179-99.
- ISAAC, M., J. WALKER and S. THOMAS (1984): "Divergent evidence on free riding: an experimental evidence of possible explanations," *Public Choice*, 43(1), 113-49.
- ISAAC, M., J. WALKER and A. WILLIAMS (1994): "Group size and the voluntary provision of public goods: experimental evidence utilizing large groups," *Journal of Public Economics*, 54, 1-36.
- JEHIEL, P. (2005): "Analogy-based expectation equilibrium," *Journal of Economic Theory*, 123, 81-104.
- KAHN, L. and J. MURNIGHAN (1993): "Conjecture, uncertainty, and cooperation in the finitely repeated prisoner's dilemma: experimental evidence," *Journal of Economic Behavior and Organization*, 22, 91-117.
- KESER, C. and F. WINDEN (2000): "Conditional cooperation and voluntary contributions to public goods," *Scandinavian Journal of Economics*, 102(1), 23-39.

- KIM, O. and M. WALKER (1984): "The free rider problem: experimental evidence," *Public Choice*, 18, 3-24.
- KLUMPP, T. (2004): "Finitely repeated provision of a public good," *mimeo Indiana University*.
- KREPS, D. and R. WILSON (1982): "Reputation and imperfect information," *Journal of Economic Theory*, 27, 253-279.
- KREPS, D., P. MILGROM, J. ROBERTS and R. WILSON (1982): "Rational cooperation in the finitely repeated prisoners' dilemma," *Journal of Economic Theory*, 27, 245-252.
- LEDYARD, J. (1995): "Public goods," IN: *Handbook of Experimental Economics*, ed. by Kagel, J. and A. Roth, Princeton University Press.
- LEVINE, D. (1998): "Modeling altruism and spitefulness in experiments," *Review of Economic Dynamics*, 1, 593-622.
- McKELVEY, R. and T. PALFREY (1992): "An experimental study of the centipede game," *Econometrica*, 60, 803-836.
- MULLER, L., M. SEFTON, R. STEINBERG and L. VESTERLUND (2008): "Strategic behavior and learning in repeated voluntary-contributions experiments," *Journal of Economic Behavior and Organization*, forthcoming.
- NAGEL, R. and F. TANG (1998): "An experimental study on the centipede game in normal form an investigation on learning," *Journal of Mathematical Psychology*, 42, 356-384.
- NEYMAN, A. (1985): "Bounded complexity justifies cooperation in the prisoner's dilemma," *Economic Letters*, 19: 227-229.
- NEYMAN, A. (1999): "Cooperation in repeated games when the number of stages is not commonly known," *Econometrica*, 67: 45-64.
- OFFERMAN, T., J. SONNEMANS, & A. SCHRAM (1996): "Value orientations, expectations, and voluntary contributions in public goods," *Economic Journal*, 106, 817-845.
- PALFREY, T. and J. PRISBEY (1997): "Anomalous behavior in linear public goods experiments: how much and why?," *American Economic Review*, 87 (5), 829-46.
- PLOTT, C. and V. SMITH (2008): "Handbook of experimental economics results, Volume 1," New York, Elsevier North-Holland Publishing.
- RABIN, M. (1993): "Incorporating fairness into game theory and economics," *American Economic Review*, 83, 1281-1302.
- RADNER, R. (1980): "Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives," *Journal of Economic Theory*, 22: 136-154.
- RADNER, R. (1986): "Can bounded rationality resolve the prisoner's dilemma?" in Andreu Mas-Collel and W. Hildenbrand (eds.), *Contributions to Mathematical Economics*, North-Holland: Amsterdam, 387-399.
- ROTH, A. E., V. PRASNIKAR, M. OKUNO-FUJIWARA and S. ZAMIR (1991): "Bargaining and market behavior in Jerusalem, Ljubljana, Pittsburgh and Tokyo: an experimental study," *American Economic Review*, 81, 1068-95.
- SAIJO, T. and H. YAMAGUCHI (1992): "The "spite" dilemma in voluntary contribution mechanism experiments," *University of Tsukuba mimeo*.
- SELTEN, R., M. MITZKEWITZ and G. UHLICH (1997): "Duopoly strategies programmed by experienced players," *Econometrica*, 65, 517-555.

SELTEN, R. and R. STOECKER (1986): "End behavior in sequences of finite prisoner's dilemma supergames: a learning theory approach," *Journal of Economic Behavior and Organization*, 7, 47-70.

SLONIM, R. and A. ROTH (1997): "Financial incentives and learning in ultimatum and market games: an experiment in the Slovak Republic," *Econometrica*, 65, 569-596.

SOBEL, J. (1985): "A theory of credibility," *Review of Economic Studies*, 52: 557-573.

SONNEMANS, J., A. SCHRAM, and T. OFFERMAN (1999): "Strategic behavior in public good games: when partners drift apart," *Economic Journal*, 35-41.

TAKAYAMA, A. (1985): *Mathematical economics*, Cambridge: Cambridge University Press.

Table 1— Treatments

1) <u>Experienced Restart</u>	2) <u>Experienced Restart with Forecast</u>
Six sets of 10-period public good contribution games with groups reshuffled with known probability 0.75 after each set	Same as Experienced Restart Session except at start of each set of 10-period games, solicit forecast of path of contributions

3) Identifying Types

Six gift exchange games (play as first and second proposer equal number of times)

Two sets of 10-period public good contribution games, group composition changes randomly after first set

Construct groups based on gift exchange play; inform subjects of gift exchange play of fellow group members
One 10-period game

Construct groups based on gift exchange play; inform subjects of gift exchange play of fellow group members
One 10-period game

**Table 2— Contribution Path
in Restarted Games**

Panel A: Last Two Sets (Games 41-60)				
Period	Experienced Restart (N=16)		Experienced Restart w/ Forecast (N=20)	
	Mean	Median	Mean	Median
1	10.9	10.0	10.1	8.5
2	11.8	11.0	8.7	5.0
3	11.4	10.5	8.6	8.0
4	10.3	8.5	7.7	3.5
5	8.6	7.5	6.1	1.0
6	5.9	4.0	5.6	0.5
7	2.8	0.0	5.3	0.0
8	1.6	0.0	5.9	0.5
9	1.3	0.0	4.5	0.0
10	0.0	0.0	2.3	0.0

Panel B: Intermediate Sets (Games 21-40)				
Period	Experienced Restart (N=24)		Experienced Restart w/ Forecast (N=12)	
	Mean	Median	Mean	Median
1	13.9	16.0	8.4	3.0
2	13.9	15.5	7.1	3.5
3	13.6	15.5	6.7	4.0
4	14.8	17.0	6.7	3.5
5	14.9	19.0	6.8	3.0
6	13.8	20.0	5.4	1.0
7	10.8	12.5	4.4	1.5
8	5.5	0.0	4.9	5.0
9	3.2	0.0	5.2	2.0
10	0.1	0.0	0.6	0.0

Table 3— Restart Effect in 10-period Games Across Sessions

Games		Number of Obs.	Mean	Median	Fraction who contribute zero
Panel A: Last Two Sets of 10-period Games					
Experienced	Last before Restart (40,50)	16	0.00	0.00	1.00
Restart	First after Restart (41,51)	16	10.94	10.00	0.19
Experienced	Last before Restart (40,50)	20	5.60	0.00	0.60
Restart	First after Restart (41,51)	20	10.05	8.50	0.15
w/ Forecast					
Panel B: Intermediate Sets of 10-period Games					
Experienced	Last before Restart (20,30)	24	0.75	0.00	0.83
Restart	First after Restart (21,31)	24	13.88	16.00	0.08
Experienced	Last before Restart (20,30)	12	1.67	0.00	0.92
Restart	First after Restart (21,31)	12	8.42	3.00	0.42
w/ Forecast					

**Table 4— Forecasts and Actual Play in
Restarted Games of Experienced
Restarts with Forecasts Session**

Games	N	<u>Forecasted Play</u>		<u>Actual Play</u>		
		Mean	Median	Mean	Median	
Forecast of First Game:						
41, 51	16	10.87	10.00	10.44	10.83	
Forecast of Fifth Game:						
41, 51	16	7.06	5.00	4.82	4.17	
Forecast of Tenth Game:						
41, 51	16	2.19	0.00	1.63	0.33	

Table 5— Impact of First Period Average Contribution^a of Opponents on Future Play

Dependent variable:	<u>Contribution in 2nd Game</u>		
	Pooled (1)	Not restarted (2)	Restarted (3)
Value of β : ^b			
OLS	0.49 [6.45]	0.31 [3.18]	0.73 [6.00]
session fixed-effects	0.43 [5.62]	0.22 [2.26]	0.68 [5.68]
fixed-effect for each set of 10-period games	0.49 [6.32]	0.26 [2.57]	0.49 [2.70]
random-effects tobit	0.26 [4.00]	0.13 [1.75]	0.71 [5.79]
N	408	320	88

^aNotes: T-statistics are in brackets under estimated coefficients. All 10-period games in the Experienced Restart and Experienced Restart with Forecasts sessions are included.

^bThe regression equation is contribution in the second game = $\beta \cdot$ average contribution of opponents in first game + controls + ϵ_{it} . The controls in the OLS specification are an intercept, in the session fixed-effects specification they are session dummies, in the fixed effect for each set of 10-period games there is a dummy for the first set of 10-period games, the second set of 10-period games, etc., while the random-effects tobit specification includes a random-effect to account for multiple observations on the same individual (implemented using the STATA xttoit command).

Figure 1: Average Contribution in Restarted Games for each Set of 10-period Games in Experienced Restart and Experienced Restart with Forecast Sessions




Figure 2: Average Contribution Pattern within Restarted Groups in Experienced Restart and Experienced Restart with Forecast Sessions




Figure 3: Average Contribution in Identifying Types Session by Group Type

