

Finitely Repeated Voluntary Provision of a Public Good*

Tilman Klumpp[†]
Emory University

July 2010

Forthcoming in *Journal of Public Economic Theory*

Abstract

This paper theoretically explores the voluntary provision of a public good when either one of the following holds: (i) agents' utility is the sum of their monetary earnings and a certain non-material component, or (ii) agents' exhibit satisficing behavior. We show that a small degree of either non-material payoffs or satisficing behavior can generate large contributions in a finitely repeated game, even if the incentive to free-ride on others' contributions calls for negligible public good provision in the static game. The equilibrium is characterized by a sharp decline in contributions toward the end of the game. Several comparative results regarding group size and technology are consistent with laboratory data obtained by Isaac and Walker (1988) and Isaac et al. (1994). The model also predicts the puzzling restart effect observed by Andreoni (1988) in an experimental study.

Keywords: Public Goods, Voluntary Contribution Mechanism, Folk Theorem, Social Preferences, Satisficing Behavior.

JEL code: C73, C92, D64, H41.

*I thank Vijay Krishna, Mattias Polborn, Richard Steinberg, Xuejuan Su, Jimmy Walker, seminar audiences at the University of Illinois and IUPUI, as well as two anonymous referees and the associate editor, for their helpful comments.

[†]Department of Economics, Rich Building 316, 1602 Fishburne Dr., Atlanta, GA 30322. E-mail: tklumpp@emory.edu.

1 Introduction

This paper examines the pattern of voluntary contributions to a public good over time. If agents have strictly dominant strategies to not contribute to the good in a one-shot setting, then clearly they should not contribute in the subgame perfect equilibrium of any finitely repeated version of the game. This prediction is at odds with a large body of empirical evidence. In a widely used experimental design, the linear voluntary contribution mechanism (VCM), subjects are given monetary endowments which they can allocate to either a private good or toward the provision of a public good. The parameters of the game are set so that the welfare maximizing allocation is for all agents to give their entire endowment to the public good. Similar to the prisoners' dilemma, however, it is a dominant strategy for each individual to allocate their entire endowment to the private good, so that one should expect zero contributions to the public good if this experiment is repeated any finite number of times. Strikingly, though, in virtually all laboratory studies subjects contribute significant fractions of their endowments.¹

One explanation for this observation is that subjects' preferences are to maximize something other than their earnings. For example, individuals may care about the well-being of others (altruism) or have a preference for giving (impure altruism or warm glow). Another explanation is that subjects are not fully rational. For example, individuals might simply imitate others, or be content with sub-optimal payoff levels (satisficing behavior). The aim of this paper is to demonstrate that, in dynamic settings, even a *small* degree of non-standard preferences or satisficing behavior can have dramatic effects on outcomes. The reason for this amplification effect is that, under certain assumptions on preferences or rationality, the static game possesses several Pareto-ranked equilibria. A folk theorem-type argument for finitely repeated games (Benoit and Krishna [1985], Radner [1980]) then shows that large contributions can be generated in equilibria of sufficiently long dynamic games. These contributions are supported by the threat to return to the zero level of contribution in case a player attempts to free-ride on others' contributions. Remarkably, such behavior can arise even when a player's utility function "almost" coincides with the player's earnings function, or when her tolerated deviation from the optimal payoff is "almost" zero.

Contributions along the payoff-dominant symmetric equilibrium path match several other empirical regularities on a qualitative level. Firstly, contributions are high in the beginning of the game but decline toward the end—the so-called endgame effect found in virtually all VCM experiments. In the experimental literature, the endgame effect is often attributed to players' learning the incentives of the game over time. In our framework, on the other hand, the endgame effect is entirely strategic and not due to learning. Secondly, the round in which the endgame effect sets in does not depend on how many rounds have

¹See Ledyard (1995) for a survey.

already been played, but on how many rounds are still remaining. This result is consistent with data presented in Isaac et al. (1994), namely that the rate of decay in contribution levels is inversely related to the length of the game. Thirdly, we characterize the effects that changes in observable parameters of the game, namely group size and technology, have on the equilibrium path. These comparative results, too, are in line with experimental observations: Isaac and Walker (1988) report that contributions increase in group size, or when production of the public good becomes more efficient; however, they decrease in group size when the efficiency parameter is adjusted so that the group’s “feasibility set” remains constant. Finally, the model predicts the restart effect reported by Andreoni (1988): After the last round of a repeated VCM experiment, subjects were told that there was enough time left to run a second set of repetitions to gather more data. Contributions in the early stages of the second set jumped up, after having declined in the late stages of the first set—a behavior that is difficult to explain through learning or reputation effects. On the other hand, a replay of the same equilibrium for a second set of rounds is certainly possible in a framework of complete information.

It is important to emphasize that the main point advocated in this paper is *not* that individuals may have a desire to contribute to public goods, or may fail to be perfectly rational. What this paper demonstrates is that it is possible to capture a rich set of phenomena in public good experiments by using only the basic tools of complete information repeated games and making *mild* assumptions on either preferences or rationality. Our results hence shed light on the possible role of strategic behavior in VCM experiments. For example, our framework suggests that the reason for the endgame effect may simply be that late in the game the possibilities for punishing deviators by withholding future contributions are necessarily limited. Alternatively, take the group size effect in Issac and Walker (1988). It may seem counter-intuitive that, *ceteris paribus*, larger groups should be more successful providing public goods than smaller groups. But larger groups can punish deviators more severely than smaller groups; therefore, less free-riding behavior should be observed in larger groups.

One may object that a well-designed experiment should eliminate these effects, for instance through random reassignments of subjects into groups after each round. Several experimental studies have shown that, even with randomly reassigned groups, public good contributions exhibit the same qualitative features as within a fixed group.² When players are randomly rematched, the results of our model are weakened, but it can still generate contributions above the myopic level. In particular, we show that a slight amount of satisficing behavior can still sustain large contributions in a repeated game with random rematching.

The remainder of the paper is organized as follows. Section 2 describes the voluntary contribution mechanism and reviews the main results from the experimental literature.

²See Andreoni and Croson (2008) for a survey.

Several competing explanations for these results will also be discussed in this section. In Section 3, we describe our model of non-standard preferences and rationality on the stage game level, and discuss its relation to existing models of reciprocity and fairness. Section 4 characterizes the subgame perfect equilibrium that arises in our framework, and relates the equilibrium behavior to the experimental results (in particular, we examine how the equilibrium path is affected by changes in group size and technology). In Section (5), we briefly discuss how the model extends to repeated games with rematching of players between rounds. Section 6 offers concluding remarks. Most proofs are in the Appendix.

2 Voluntary Public Good Provision: Theory and Evidence

2.1 The Voluntary Contribution Mechanism

Consider the following simple environment, known as the linear voluntary contribution mechanism (VCM): There are N players, each endowed with K units of a private good. A player's set of strategies is $[0, K]$, where $s_i \in [0, K]$ is player i 's contribution to the production of a public good. The profile of contributions across all players is denoted $\mathbf{s} = (s_1, \dots, s_N)$. Player i 's consumption of the private good is then $K - s_i$, and the amount of the public good produced is $\alpha \sum_{i=1}^N s_i$, where $N^{-1} < \alpha < 1$. Thus, a player's total consumption is

$$u_i(\mathbf{s}) = K - s_i + \alpha \sum_{j=1}^N s_j.$$

Call this one-shot game Γ_1 , and the T -fold repetition thereof Γ_T . We use the convention to denote by $t = 1$ the last stage, by $t = 2$ the second last stage, and so on, so that $t = T$ is the first time the mechanism is played. Denote by s_i^t player i 's contribution at time t , and let $\mathbf{s}^t = (s_1^t, \dots, s_N^t)$ be the contribution profile at t . Thus $(\mathbf{s}^1, \dots, \mathbf{s}^T)$ denotes the outcome of the repeated game Γ_T , and player i 's average consumption in this game is

$$U_i(\mathbf{s}^1, \dots, \mathbf{s}^T) = \frac{1}{T} \sum_{t=1}^T u_i(\mathbf{s}^t).$$

A pure strategy for player i in Γ_T is a mapping from histories, defined as usual, into actions at each stage.

The utility functions u_i and U_i are referred to as player i 's material payoffs. If the players' objective is to maximize their material payoffs, the voluntary contribution mechanism resembles a prisoners' dilemma: In Γ_1 , $s_i = 0$ is the dominant strategy for each player i , and the resulting dominant strategy equilibrium is a Pareto-dominated outcome. In particular, the full-contribution profile (K, \dots, K) is not an equilibrium, but maximizes group welfare. As a consequence of the uniqueness of the zero-contribution equilibrium in Γ_1 , $s_i^t = 0 \forall i, t$ is the unique subgame perfect equilibrium in Γ_T . As before, $s_i^t = K \forall i, t$ is the welfare maximizing outcome.

2.2 Experimental Studies

A large body of experimental research has examined the VCM game as described above, and an excellent survey of the seminal contributions to this literature is Ledyard (1995). Below we focus on a set of empirical regularities found by Isaac et al. (1984), Isaac and Walker (1988), Andreoni (1988), Isaac et al. (1994), and others. The common theme from these and other studies is contrary to the theoretical prediction:

Observation 1. *In repeated VCM games, contributions are typically positive but tend to decay as the game progresses.*

It is interesting to examine how the contribution pattern is affected by changes in the game’s parameters. Adopting Isaac and Walker’s terminology, we call N the *group size*, and α the *marginal per-capita return (MPCR)* from investing in the public good. Both of these variables can be controlled by the experimenter. Isaac and Walker (1988) compare two different group sizes, $N = 4$ and $N = 10$, and two different MPCRs, $\alpha_H = 0.75$ and $\alpha_L = 0.3$. The four possible combinations of N and α are called $4L$, $4H$, $10L$, and $10H$, respectively. Some of the findings that emerged from these experiments are described below.

The first finding concerns changes in group size. As a fraction of their per-period endowment, members of the $10L$ groups allocated significantly more resources to the public good than members of the $4L$ groups. While there was no significant difference between the $10H$ and $4H$ groups in either direction, we may still note that

Observation 2. *Contributions increase in group size: Holding MPCR fixed, larger groups tend to provide more of the public good per person.*

The second finding concerns changes in MPCR. Members of the $10H$ group provided more of the public good than members of the $10L$ group; the same observation was made when comparing the $4H$ group and the $4L$ group. We can thus note the following:

Observation 3. *Contributions increase in the efficiency of public good production: Holding group size fixed, a higher MPCR leads to higher contributions per person.*

Finally, one can compare groups with different sizes and MPCR, but with identical “feasibility sets.” Specifically, Isaac and Walker (1988) compare the contributions in the $4H$ vs. $10L$ groups. In each period, both groups can provide up to $3K$ units of the public good. However, the $4H$ group provided more than the $10L$ group. We therefore note:

Observation 4. *Contributions are lower in larger groups with smaller MPCR: When varying MPCR and group size simultaneously, holding αN constant, larger groups are less successful than smaller ones.*

Regarding the length of the game T , two observations will be noted. First, Issac et al. (1994) compare $T = 10$, $T = 40$ and $T = 60$, and observe the following:

Observation 5. *If the time horizon is increased, contribution tend to stay high for a higher number of rounds.*

Second, Andreoni (1988) reports the following surprising result. The time horizon $T = 10$ was common knowledge among subjects, but after stage 10 had been completed, subjects were told that there was enough time left to run an additional set of stages in order to generate more data. At the beginning of this second set, observed contribution levels were higher than those at the end of the first set, and almost as high as those at the early stages of the first experiment.

Observation 6. *Restart effect: An unanticipated restart of a repeated VCM game, after it has ended, increases contributions.*

2.3 Possible Explanations

The experimental results described above are clearly at odds with the theoretical prediction that contributions should be zero. Several explanations of these phenomena have been proposed, which can roughly be grouped in three categories. (The literature relating to the question why people cooperate, or why they contribute to public goods, is vast, and no attempt at an exhaustive survey will be made here.)

In the first category, it is argued that the monetary payouts provided by the experimental design accurately reflect preferences, but that people typically fail to behave rationally and do not maximize these preferences. Without a fully rational understanding of the game, subjects may use other decision making procedures such as imitation, or simply err and learn the incentives of the game only through repeated trials. There is a considerable body of literature, theoretical and experimental, that concerns imitation in related games such as the prisoners' dilemma (Eshel et al. (1998), Ahn et al. (2001), and others). Simple decision errors have been studied theoretically in the public good context by Anderson et al. (1998). They show that a quantal response equilibrium, involving the "correct" statistical distribution of errors, can exhibit many of the properties found in the experimental studies. Palfrey and Prisbrey (1996) argue that, because the equilibrium contribution is zero, any errors must necessarily manifest themselves as over-contributions. The decline in contributions can then be attributed to a simple reduction in the size or frequency of errors. However, such learning cannot explain either the restart effect (Observation 6) or the slow decay effect (Observation 5).

The second category of explanations is based on the assumption that subjects act rationally, or at least do not err systematically, but that their preferences are to maximize something other than their earnings in the experiment. Warm glow, altruism, and other forms of other-regarding preferences have already been mentioned in the introduction, and there is evidence that suggests that those motivations indeed influence subjects' decisions

to give to public goods. Andreoni (1995) contains a series of experiments aimed at disentangling altruistic motives (which belong to the second category) from confusion (the first category). Using public good experiments with interior equilibria, Keser (1996) and Isaac and Walker (1998) present evidence of what the latter authors call a “residual cooperative component” that can explain the contribution pattern in games with zero-contribution equilibria. Eliciting contribution schedules for a set of possible MPCR values, Brandts and Schram (2001) find that the decisions of some subjects are driven by other-regarding motives. The linear VCM game is explicitly modeled as a finitely repeated game in a recent paper by Ambrus and Pathak (2009), who show that the presence of players with reciprocal utility functions (i.e., interdependent utility across stages) can explain the declining contributions effect (Observation 1) and the restart effect (Observation 6).³

Finally, some have suggested incomplete information about each others’ preferences or rationality as a third source for cooperation. Palfrey and Rosenthal (1988) present a static public good model with altruism and incomplete information. Dynamic models with incomplete information are more complex, as reputational concerns arise in these settings. The theoretical groundwork for the study of how multi-sided incomplete information affects cooperation in repeated games is laid by Kreps et al. (1982) in the context of the prisoners’ dilemma, and is experimentally tested by Andreoni and Miller (1993). A reduced-form reputational model is developed in Brandts and Figueras (2003). While public goods problems share important features with the prisoners’ dilemma, modelling multi-sided incomplete information in repeated public goods games is more difficult, due to the larger action space and the need to find consistent beliefs for a larger set of out-of-equilibrium actions. In light of this difficulty, we abstract from possible incomplete information in this paper. We develop instead a framework that falls in the preference-based as well as the rationality-based category. This will be done in the next section.

3 Non-Material Utility and Satisficing Behavior

3.1 Non-Material Utility

Let us first consider players whose preferences are different from their material payoffs. Instead, assume that in Γ_1 each player maximizes

$$v_i(\mathbf{s}) = u_i(\mathbf{s}) + \beta\rho(\mathbf{s}),$$

where the function ρ measures non-material utility derived from a contribution profile. One can think of ρ as a representation of a player’s “social preferences,” that is, considerations

³Ambrus and Pathak (2009) also present experimental evidence that rational selfish players indeed contribute in the presence of reciprocal players. Charness and Rabin (2002) present tests designed to test for social preferences (outside the strict public good context) and find some evidence for reciprocal motivations.

of how the outcome \mathbf{s} is evaluated in the player's mind that are not described by the player's material payoff. The weight place on this non-material component is given by the scalar $\beta \geq 0$. In Γ_T , the average payoff for a player is then given by $V_i(\mathbf{s}^1, \dots, \mathbf{s}^T) = \sum_{t=1}^T v_i(\mathbf{s}^t)/T$.

For sake of simplicity, we will assume a particular functional form for the non-material payoffs:

$$\rho(\mathbf{s}) = (s_i)^{\gamma_1} (\bar{s}_{-i})^{\gamma_2}, \quad \gamma_1 > 0, \gamma_2 > 0, \gamma_1 + \gamma_2 < 1, \quad (1)$$

where $\bar{s}_{-i} \equiv (N-1)^{-1} \sum_{j \neq i} s_j$ is the average contribution among i 's opponents. With the functional form assumption for ρ , a player's best response to an average contribution of \bar{s}_{-i} among her opponents is

$$b(\bar{s}_{-i}) = \min \left\{ K, \left[\frac{\gamma_1 \beta}{1 - \alpha} (\bar{s}_{-i})^{\gamma_2} \right]^{1/(1-\gamma_1)} \right\}. \quad (2)$$

This best response has several appealing properties. Since $b(0) = 0$, players do not possess an unconditional preference for giving to the public good. On the other hand, if at least one other player contributes to the public good, then player i would like to contribute herself.⁴

3.2 Satisficing Behavior

Let us now consider players who do not maximize their payoffs but instead are content with reaching a payoff that is within a distance ε of the maximal possible level. Call a strategy profile in a normal form game an ε -*equilibrium* if no player can increase her payoff by more than ε (with $\varepsilon \geq 0$) by switching her prescribed strategy to any alternative strategy, given the strategies of the other players (see Radner [1980]). Formally:

Definition 1. A profile \mathbf{s} is an ε -*equilibrium* of Γ_1 if for all i and $s'_i \in [0, K]$, $u_i(\mathbf{s}) \geq u_i(s'_i, \mathbf{s}_{-i}) - \varepsilon$.

If $\varepsilon = 0$, the definition is that of Nash equilibrium. If $\varepsilon > 0$, however, the notion of ε -optimality captures the idea of *satisficing behavior* of players.

In repeated games, we will apply the ε -optimality criterion to the average payoffs in each subgame. This implies that players are permitted to make an ε -mistake in every

⁴We assume a functional form only because the paper is concerned with the dynamic strategic effects in the VCM game, and the utility representation of a player's non-material preferences does *not* matter for these dynamic effects as long as these preferences induce at least two Pareto-ranked equilibria in the stage game (see Lemma 1). To this end, we could have used any function ρ whose induced best response is continuous, twice differentiable on $(0, K]$, satisfies $b(0) = 0$, $b'(\bar{s}_{-i}) > 0$, $b''(\bar{s}_{-i}) < 0$, $\lim_{\bar{s}_{-i} \rightarrow 0} b'(\bar{s}_{-i}) > 1$, and $\partial b(\bar{s}_{-i})/\partial \beta \rightarrow 0$ as $\beta \rightarrow 0$.

round of the repeated game.⁵ The definition of ε -equilibrium then extends easily to the following:

Definition 2. A strategy profile is ε -subgame perfect equilibrium (ε -SGPE) of Γ_T if it induces an ε -equilibrium in every subgame of Γ_T .

Similar to the preference scaling factor β above, the size of ε provides a measure for the degree of non-optimality of players' choices. The smaller ε , the less the satisficing aspect of the players' behavior and the closer are their decisions to optimality.

3.3 Stage Game Equilibria

Fixing all other parameters, the framework developed above gives rise to a class of games described by the values for β and ε . Let us denote these games by $\Gamma_T(\beta, \varepsilon)$. Throughout the paper we will consider only games where at least one of β and ε is zero. We will refer to the game $\Gamma_T(0, 0)$ as the *standard model*.

One may regard a game $\Gamma_T(\beta, \varepsilon)$ as similar to $\Gamma_T(0, 0)$ if equilibrium play in the static game $\Gamma_1(\beta, \varepsilon)$ is similar to equilibrium play in the static game $\Gamma_1(0, 0)$. That is, in the one-shot case (or, equivalently, in the repeated case with myopic players), the effects of non-material payoffs or satisficing behavior would be small. The results in this paper concern cases where models that are similar to the standard one, in this static sense, have largely different dynamic properties. Taking this perspective hence allows us to separate the strategic (non-myopic) incentives to cooperate in dynamic games from any direct (myopic) incentives.

The static standard model $\Gamma_1(0, 0)$ resembles a prisoners' dilemma: Contributing zero to the public good is the dominant strategy for each player i . In $\Gamma_1(\beta, 0)$, almost the same happens if β is small:

Lemma 1. *If $\beta > 0$ there exist two pure strategy Nash-equilibria in $\Gamma_1(\beta, 0)$. One of these equilibria is $\mathbf{s} = (0, \dots, 0)$, and the other is $\mathbf{s} = (\hat{s}, \dots, \hat{s})$, where*

$$\hat{s} = \min\{K, [\gamma_1\beta/(1-\alpha)]^{1/(1-\gamma_1-\gamma_2)}\}.$$

Observe that $\hat{s} \rightarrow 0$ as $\beta \rightarrow 0$. Thus, scaling the non-material component of players' utility functions toward zero, the "cooperative equilibrium" moves closer to the "selfish equilibrium," and when $\beta = 0$ they collapse into the unique selfish equilibrium. To avoid trivial full contribution outcomes, we shall assume that $\hat{s} < K$.

A similar result holds if $\beta = 0$ and $\varepsilon > 0$:

Lemma 2. *If $\varepsilon > 0$ there exists a continuum of pure strategy ε -equilibria in $\Gamma_1(0, \varepsilon)$. In particular, every profile $\mathbf{s} = (s_1, \dots, s_N)$ such that $s_i \in [0, \varepsilon/(1-\alpha)] \forall i$ is a pure strategy ε -equilibrium.*

⁵This assumption is not crucial for the results, and the same qualitative predictions can be obtained if we applied ε -optimality to total payoffs.

Thus, the smaller ε the closer are players to full rationality, and the smaller the contributions which can be sustained in equilibrium. Again, we shall assume that $\varepsilon/(1-\alpha) < K$, so that the full contribution profile is never an equilibrium in the static game.

Thus, we talk of a game being “close” to the standard model if equilibrium behaviors in the corresponding static games are similar. This will be the case if β or ε is small. Of course, we cannot make a statement as to what the exact (i.e. numerical) range of behaviors is that we would consider close to the standard one. Our concept of closeness simply rests on the standard notion of continuity, guaranteeing that we can make play in the static game as similar to the standard model as we want, by choosing β or ε sufficiently small.

4 Dynamic Equilibrium

This section characterizes the equilibrium of the repeated VCM game with either social preferences and satisficing behavior. Only symmetric equilibria will be considered. A vector of T numbers $\sigma = (\sigma^1, \dots, \sigma^T)$ is called an *equilibrium path* if there exists a ε -subgame perfect equilibrium in $\Gamma_T(\beta, 0)$, or an ε -subgame perfect equilibrium in $\Gamma_T(0, \varepsilon)$, in which σ^t is the contribution made by every player in round t of the game. We will further focus on maximal equilibrium paths, meaning that no other equilibrium path $\tilde{\sigma}$ exists for which $\sum_{t=1}^T \tilde{\sigma}^t > \sum_{t=1}^T \sigma^t$.

The main result is stated below:

Theorem 3. *Consider either a game $\Gamma_T(\beta, 0)$ with $\beta > 0$, or a game $\Gamma_T(0, \varepsilon)$ with $\varepsilon > 0$. There is a unique maximal symmetric equilibrium path $\sigma = (\sigma^1, \dots, \sigma^T)$. This path is independent of T (except for its length) and has the following property: There exists $\bar{T} \in \mathbb{N}$ such that $\sigma^t = K$ for all $t \geq \bar{T}$, and $\sigma^t < \sigma^{t+1}$ for all $1 \leq t < \bar{T}$. Furthermore, σ is increasing in β and ε , and \bar{T} is decreasing in β and ε .*

Note that if the actual length of the game is $T < \bar{T}$, then full contributions will not be observed. However, by increasing the length of the game to \bar{T} or larger, one can generate full contributions in the first $T - \bar{T} + 1$ periods of play.

To prove Theorem 3, recall that in $\Gamma_1(\beta, 0)$ and $\Gamma_1(0, \varepsilon)$ there exist several Pareto-ranked equilibria. One can therefore construct subgame perfect equilibria in $\Gamma_T(\beta, 0)$, and ε -subgame perfect equilibria in $\Gamma_T(0, \varepsilon)$, that involve actions which are not stage-game equilibria; see Radner (1980) and Benoit and Krishna (1985). For ease of exposition, we focus on simple trigger strategies to provide unrelenting punishments. That is, profitable deviations from a prescribed path will be punished by switching to zero contributions, which is always an equilibrium.⁶ Notice that, by the nature of the voluntary contribution mechanism, it is not possible to target punishments or rewards at individual players

⁶These unrelenting punishments are sufficient, but usually not necessary, to sustain the maximal path. That is, a small under-contribution by one player usually does not have to result in the zero contribution

without punishing or rewarding all other players at the same time. The maximal path that can be supported by these punishments is constructed below. This will be done for the game $\Gamma_T(\beta, 0)$, with $\beta > 0$. The proof for the game $\Gamma_T(0, \varepsilon)$ is similar and is in the Appendix.

4.1 Proof of the result

To describe the trigger strategy, let $\sigma = (\sigma^1, \dots, \sigma^T)$ be a contribution path. Define a collection of functions $w = (w^1, \dots, w^T)$, prescribing symmetric contributions to be made by the players in each period as a function of previous play, as follows:

1. In round T , each player's contribution is $w^T = \sigma^T$.
2. In round $t < T$, each player's contribution is

$$w^t(s^{t+1}, \dots, s^T) = \begin{cases} \sigma^t & \text{if } s_i^\tau \geq \sigma^\tau \ \forall i = 1, \dots, N, \ \forall \tau = t+1, \dots, T, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

We will construct the maximal path σ for which w is a subgame perfect equilibrium in $\Gamma_T(\beta, 0)$.

To simplify notation, let $u(s)$ and $\rho(s)$ be the material and non-material payoffs a player obtains (in a single stage) if all players contribute s , and let $u(s', s)$ and $\rho(s', s)$ be payoff the player obtains if she contributes s' while all other players contribute s . Define $v(s)$ and $v(s', s)$ accordingly. Call $b(s)$ a player's best-response to the opponent profile $\mathbf{s}_{-i} = (s, \dots, s)$ in the stage game. As shown in Section 3, b is strictly increasing and strictly concave and satisfies $b(0) = 0$. Define a function $\delta : [0, K] \rightarrow \mathbb{R}$ as follows:

$$\delta(s) = v(b(s), s) - v(s) = (1 - \alpha)[s - b(s)] + \beta[\rho(b(s), s) - \rho(s)]. \quad (4)$$

$\delta(s)$ is the one-shot gain for a player who deviates in some period from s to $b(s)$, while all others hold their contributions constant at s . Also define a function $\mu : [0, K] \rightarrow \mathbb{R}$ as follows:

$$\mu(s) = v(s) - v(0) = (\alpha N - 1)s + \beta \rho(s) > 0. \quad (5)$$

$\mu(s)$ is the utility loss incurred by a player when the symmetric contribution profile s is replaced by the selfish equilibrium $s = 0$ in a given round of the game.

Now construct σ as follows. At stage 1 a Nash equilibrium of Γ_1 must be played. By Lemma 1 there are two candidate contributions, \hat{s} and zero. Set $\sigma^1 = \hat{s}$ and let

$$B_1 = \mu(\sigma^1) > 0$$

profile being played from that period forward, but could instead be punished by a relatively small reduction in the contributions of the other players.

be the payoff difference for a single player between these two equilibria. Next, we will construct the contribution σ^2 to be made in round 2. Suppose $\sigma^2 > \sigma^1$ and consider a deviation to $b(\sigma^2) < \sigma^2$. Given the trigger strategy (3), the deviating player gains $\delta(\sigma^2) > 0$ at stage 2 and loses B_1 at stage 1. In order to support σ^2 at stage 2, we therefore need $\delta(\sigma^2) \leq B_1$. Thus, set

$$\sigma_2 = \max\{s \in [0, K] : \delta(s) \leq B_1\}.$$

Observe that δ is continuous, $\delta(s) \geq 0$, and $\delta(s) = 0$ if and only if (s, \dots, s) is a Nash equilibrium of Γ_1 . Thus, a well defined σ^2 exists and $\sigma^2 > \sigma^1$.

Proceed now in similar fashion and set $B_2 = B_1 + \mu(\sigma^2) > B_1$ and $\sigma^3 = \max\{s \in [0, K] : \delta(s) \leq B_2\}$. Since $B_2 > B_1$ and $d(\sigma^2) \leq B_1$, σ^3 exists and $\sigma^3 > \sigma^2$ if $\sigma^2 < K$ (otherwise $\sigma^3 = \sigma^2 = K$). In general, using the starting values $B_0 = 0$ and $\sigma^1 = \hat{s}$, one can compute B_t and σ^t recursively as follows:

$$B_t = B_{t-1} + \mu(\sigma^t) > B_{t-1}$$

and

$$\sigma^t = \max\{s \in [0, K] : \delta(s) \leq B_{t-1}\}.$$

It is easily seen that $\sigma^t > \sigma^{t-1}$ if $\sigma^{t-1} < K$, and $\sigma^t = \sigma^{t-1} = K$ otherwise. The resulting path σ is then the maximal path that is supported by the trigger strategy w . Since w involves maximal punishments, σ is the maximal symmetric equilibrium path of Γ_T . Finally, the value \bar{T} is defined as $\bar{T} = \min\{t : \sigma^t = K\}$.⁷

This proves Theorem 3 for the game $\Gamma_T(\beta, 0)$, except for the comparative properties with respect to β . These are established in the Appendix (as is the result for $\Gamma_T(0, \varepsilon)$).

4.2 Relation to experimental observations

In Section 2, several empirical observations regarding experimental studies of the repeated VCM game were discussed. Theorem 3 in the previous section is, on a qualitative level, in line with several of these observations. Most importantly, the fact that the maximal contributions σ^t are positive but decline over time is consistent with Observation 1. Furthermore, increasing T to, say T' does not alter $\sigma^1, \dots, \sigma^T$ (not that these are the *final* T stages of the game). This means that, in longer games, high contributions can be sustained for a larger number of rounds, a result consistent with Observation 5.

The restart effect (Observation 6) also does not contradict the equilibrium constructed in the Section 4: When a second set of T stages is added, without this being known to the players beforehand, the same equilibrium can clearly be played in the second set. This means, after contributions have declined toward the end of the first set, restarting the game will result in an increase in contributions from σ^1 to σ^T .

⁷Note that if $\sigma^t < K$ then $\sigma^t \geq t\hat{s}$, by construction. Thus, \bar{T} must be finite.

We now examine how our framework relates to the other observations discussed in Section 2. Observations 2–4 concern the effects of changes in group size (N) and MPCR (α) on observed contribution patterns. Theorem 4 states comparative results when either α or N is changed, and the other parameter is held constant. The proof is in the Appendix.

Theorem 4. *Consider either a game $\Gamma_T(\beta, 0)$ with $\beta > 0$, or a game $\Gamma_T(0, \varepsilon)$ with $\varepsilon > 0$. The maximal equilibrium path σ is increasing in α and N , and the stage after which contributions decrease (\bar{T}) is decreasing in α and N .*

This result is consistent with Observation 2 and Observation 3. Increases in N or α will make groups more successful, in that the maximal path of contributions that can be achieved in equilibrium increases. The reason is that a larger group size (higher N) or a higher efficiency of public good production (higher α) make it easier to punish deviators even late in the game. (This effect can be easily seen when inspecting (5)).

In the following numerical example we illustrate the results of Theorem 4. We consider the VMC game with $T = 15$ and $K = 1$. We take the preference-based approach and set $\beta = 0.04$ and $\varepsilon = 0$. The parameters in the non-material utility function are $\gamma_1 = 1/2$ and $\gamma_2 = 1/4$. Figure 1 depicts the maximal subgame perfect contribution paths for different values of α and N . In the top panel of Figure 1, $\alpha = 0.5$ is held fixed and N is varied. In the middle panel, $N = 4$ is held fixed but α is varied. In the benchmark case $\alpha = 0.5$ and $N = 4$ (the thick curve), full contributions are possible until 8 periods prior to the end of the game. Notice also that in the last few periods the incentives to free-ride on others' contributions are almost entirely preserved and equilibrium contributions are therefore negligible.

In an infinitely repeated game with standard preferences and rationality assumptions, Pecorino (1999) shows that the discount factor required to achieve a desired level of contributions goes to zero as group size increases. A similar result can be shown in a finitely repeated setting as a corollary to Theorem 4: Holding everything else fixed, the amount of non-selfishness (β) or satisficing behavior (ε) required to achieve a desired total level of contributions goes to zero as N increases.

It is also illuminating to examine how equilibrium contributions vary across groups with the same feasibility sets. Consider a change in group size which is accompanied by a inversely proportional change in MPCR, so that under full cooperation the group can produce the same amount of the public good. In this case, smaller groups will provide more of the public good than larger ones, as the following result states. Furthermore, regardless of group size, there exists a lower bound on the best possible equilibrium outcome:

Theorem 5. *Consider either a game $\Gamma_T(\beta, 0)$ with $\beta > 0$, or a game $\Gamma_T(0, \varepsilon)$ with $\varepsilon > 0$. Consider a sequence $(\alpha_m, N_m) \rightarrow (0, \infty)$ as $m \rightarrow \infty$, such that $\alpha_m N_m = c > 1 \forall m$. Let σ_m be the associated sequence of maximal contribution paths. Then the following holds:*

- (a) *For all $m < m'$, $\sigma_m > \sigma_{m'}$.*

- (b) *There exists a positive and increasing limit path $\omega = (\omega^1, \dots, \omega^T)$ such that $\sigma_m > \omega$ $\forall m$ and $\sigma_m \rightarrow \omega$.*

This result is consistent with Observation 4. To illustrate the effect of changing group size and MPCR simultaneously, we use the same example as above, but we let α and N vary while holding $\alpha N = 2$ fixed. The contribution paths for five different values of N are depicted in the bottom panel in Figure 1 (in each case $\alpha = 2/N$).

4.3 Remarks

We now offer a few additional remarks concerning our theoretical results and their interpretation in light of the experimental data.

First, it may seem odd that in the ε -model we assume that players do not fully maximize their objective function but at the same time possess the capacity to follow the trigger strategy posited in (3). Under this strategy, current actions are dependent on the history of past play. The strategy is thus more complex than history-independent strategies such as, say, the strategy to always contribute the same amount to the public good. However, within the set of all history dependent strategies, the strategy (3) is relatively simple: All punishments are implemented by players reverting to the zero-contribution profile once the first deviation is observed. Furthermore, even if not fully maximizing pay-offs, players may still understand the incentives imposed on them by the threat of others withholding their contributions in response to free-riding. In our view, the experimental data are consistent with the idea that players are afraid of being punished in this way if they under-contribute. For example, late in the game the possibilities for punishing deviators by withholding future contributions are necessarily limited; hence it is harder to sustain large contributions in later rounds. The slight irrationality present in the ε -model is needed to generate multiple equilibria in the very last stage, which in turn is necessary for the threat not to unravel if players can foresee the final period. However, ε need not be large to generate this overall incentive structure.

Second, even though there are many qualitative similarities between the predictions of this model and the data, there are also some differences. The graphs in Figure 1 show a very pronounced decline of contributions, from the full level to virtually nothing, in only a few periods. In typical VCM experiments, however, early contributions are not full, late contributions are not zero or almost zero, and the decline is much smoother than what the model here predicts.⁸ Our results should hence be interpreted as a statement about the general nature and direction of the strategic effects in VCM games. As with any theoretical model, other effects are likely to be present in reality. In fact, simple random

⁸The fact that late contributions are significantly above zero can obviously be explained by a sufficiently large value for β , but then it seems surprising that early contributions are not at the maximum possible level. Conversely, early contributions that are below the maximum level can be explained by a sufficiently small β , but then late-round contributions should be almost zero.




Figure 1: Contribution paths for different N and α

Top graph: N varies, $\alpha = 0.5$ constant. Middle graph: α varies, $N = 4$ constant.
 Bottom graph: N varies, $\alpha N = 2$ constant.

errors can account for early under-contributions as well as late over-contributions, relative to the equilibrium path. Furthermore, early under-contributions by mistake would tend to be persistent in our framework, as they cannot be distinguished from intentional free-riding and would have to be punished by withholding at least some contributions in other rounds. Keep in mind that the severe punishments which are used in our equilibrium are not necessary to sustain the maximal path σ . It is straightforward to construct sufficient punishments which sanction a small deviation or mistake made by one player by a small reduction in contributions during the next period. The observed contribution profile over time would then be smoother than what is depicted in the figures above.

Finally, while the focus here is on symmetric equilibria in which all players make the same contributions, a significant degree of heterogeneity is observed in experimental studies. One possible explanation is that subjects play an asymmetric equilibrium. In our framework, an infinity of such equilibria exists. One class of asymmetric equilibria is for some players to play the dynamic equilibrium among themselves, while others free-ride. The first type would contribute to their own benefit as well as the benefit of all other subjects. The free-riding type, on the other hand, could be a player who understands the incentives of the static game but fails to grasp the dynamic equilibrium. Coalitions that can sustain the dynamic equilibrium described in Section 4 need to consist of at least $1/\alpha$ players. For example, in the $4H$ case of Issac and Walker (1988) it would be possible that only two subjects play the dynamic equilibrium. In the early stages of the repeated game, contributions of about 50% of total endowments can then be predicted.

5 Random Reassignment of Players

We have so far considered a repeated VCM game where the same group of participants interacts in each period. A number of experimental papers examine the effects of randomly reassigning participants into groups in every period, precisely to eliminate any strategic effects across rounds. Such a design is sometimes called a *strangers treatment*, or a *contagion-free design*. The experimental evidence in this regard is somewhat inconclusive: Andreoni (1988) makes the observation that randomly rematched subjects contribute more than in a fixed group, while Croson (1996) finds that they contribute less (a survey of these and other results is Andreoni and Croson (2008).)

What effect would random rematching of players have on our theoretical predictions? To answer this question, consider an environment where in each round the VCM game is played by $M \geq 2$ groups of $N \geq 2$ players each. (The model examined so far had $M = 1$.) After each round, the composition of the M groups is determined anew through a random, uniform reassignment in which each player has the same chance of being placed in any one of the M groups. These draws are independent across time. Further, a player does not observe the identity of the other players who participate in her own group, at any time.

5.1 Contagion

Our goal is to examine whether the trigger strategy used previously can generate contributions above the myopic level. Note that, due to the random reassignment procedure, a defection by one player will not immediately spread through the entire population of players as it did for $M = 1$. Instead, it will contaminate the population more slowly. This implies that punishment via withholding contributions works in a different way than before, for two reasons. First, it does not necessarily reduce the defector's payoff, as the defector may be assigned to a new group (consisting of players who still make their equilibrium contributions). Second, it may reduce the payoffs for players who were not in the same group as the defector, and if these players use the same trigger strategy they themselves will start withholding their own contributions from the next period onward. By choosing *not* to punish a defector, a player can thus retard the speed at which defection spreads across the population of players.⁹ The first effect makes it more difficult to deter free-riding by the threat of reducing future contributions, while the second reduces the players' incentives to punish free-riders in the first place.

A full treatment of these contagion dynamics, and how they would be affected if a player decided not to punish a defector, is beyond the scope of this section. Instead, we will restrict our attention to the limiting case explored in Theorem 5 (for $M = 1$) above: Fix the number of groups at $M > 1$ and suppose that $N \rightarrow \infty$ and $\alpha \rightarrow 0$, such that $\alpha N = c > 1$ stays constant. In the limit as $N \rightarrow \infty$, the contagion dynamics become easy to characterize: If every player follows the trigger strategy and a defection occurs in round t , then there will be a fraction of $1/M$ defecting players in every group in round $t - 1$ (almost surely, by the law of large numbers). This, in turn, implies that in all rounds $\tau < t - 1$, *every* player defects (almost surely). Furthermore, no single player can retard these dynamics by choosing not to punish a defecting player.

5.2 Solution concept

A slight difficulty arises due to the fact that the VCM game now possesses only one subgame, namely the entire game itself. The solution concept of subgame perfect equilibrium (resp. ε -SGPE)—now identical to that of Nash equilibrium (resp. ε -equilibrium)—therefore no longer entails any notion of sequential rationality that seems reasonable in this multi-stage game. One way to incorporate sequential rationality is to introduce players' beliefs about their opponents' behavior in the model, and make these beliefs consistent with (a) the random-matching structure of the game and (b) the strategies of the players; for example, by way of a sequential equilibrium. Sequential rationality then boils down to the requirement that a player's strategy be optimal (or ε -optimal) given the player's belief, at every stage.

⁹See Ellison (1994) for a discussion of this retardation effect in a repeated prisoner's dilemma game.

An alternative that bypasses belief formation is the concept of extended subgame perfection (see Kreps and Wilson, 1982). Loosely speaking, a strategy profile is extended subgame perfect if (i) each player's strategy is optimal in each subform of the game (a subform is a subgame with more than one initial node), and (ii) the conditional probability distribution over initial nodes in each subform ("beliefs" so to speak) is derived from Bayes rule when this is possible, including subforms that have a zero probability overall but a positive probability conditional on a larger subform being reached.¹⁰ In the context of our game, condition (ii) implies that, if a deviation occurs (a null event in equilibrium), the distribution over opponent's contribution following this deviation is generated by the contagion dynamics described above. In the following, we use extended subgame perfection to impose sequential rationality on the players. As usual, for the rationality-based model the optimality criterion will be that of ε -optimality.

5.3 The equilibrium contribution path

Let us first look at the preference-based model ($\beta > 0$, $\varepsilon = 0$). Consider the last stage of this game (round 1). Provided no previous defections have been observed by the members of a group, they may either play the zero-contribution profile $(0, \dots, 0)$ or the positive-contribution profile $(\hat{s}, \dots, \hat{s})$. As before, set $\sigma^1 = \hat{s}$, and suppose that $\sigma^2 > \hat{s}$ can be supported in equilibrium. If a player contributes less than σ^2 in round 3, then (under the trigger strategy) all players who are in the same group as the defector in round 2 will reduce their final-round contribution from σ^1 to zero. Due to random rematching with a very large number of players, in round 1 every group must contain a fraction $1/M$ of opponents that contribute zero. The average contribution in the final round is then

$$\bar{s}_{-i} = \frac{1}{M} \cdot 0 + \frac{M-1}{M} \cdot \hat{s} = \frac{M-1}{M} \hat{s} \text{ almost surely.}$$

Thus, for the trigger strategy to be used in equilibrium, a zero contribution must be a best response to an average opponent contribution of $\frac{M-1}{M} \hat{s}$. Whether this is the case depends on the specification of non-material payoffs. If we use the specification from before, given in (1), then it is clear from the associated best response (2) that a zero contribution is not optimal if the opponents contribute a positive amount on average.¹¹ In fact, given that $b(\bar{s}_{-i}) > \bar{s}_{-i}$ for all $\bar{s}_{-i} < \hat{s}$, no reduction in contributions (either to zero or to some small positive level below \hat{s}) can be part of an equilibrium punishment. Thus, the best contribution path in extended SGPE of $\Gamma_T(\beta, 0)$ with random rematching is given by the myopic path $\sigma^t = \hat{s} \forall t$.

¹⁰The precise formal definition of extended subgame perfection is in Kreps and Wilson (1982), p. 877.

¹¹Note that using the best response (2) is appropriate in the limit as $N \rightarrow \infty$. For finite N , on the other hand, the variable \bar{s}_{-i} is random (governed by a binomial distribution) and each player would choose a contribution that maximizes expected payoffs, yielding a different and more complicated best response.

On the other hand, suppose non-material payoffs were changed in such a way that a player wants to contribute a small amount \hat{s} if and only if *all* opponents contribute the *same* small amount \hat{s} , and zero otherwise. This would not alter any of our previous results, but zero would indeed be a best response to an average opponent contribution of $\frac{M-1}{M}\hat{s}$.

Now turn to the rationality-based model. We can show that it is possible to generate large contributions in an extended ε -SGPE of $\Gamma_T(0, \varepsilon)$, even if the players are randomly reassigned, supported by the trigger strategy with unrelenting punishments. Define

$$\lambda(s) = u(s) - u\left(\frac{M-1}{M}s\right) = \lim_{\substack{N \rightarrow \infty, \\ \alpha N = c}} \alpha N \frac{1}{M} s = c \frac{1}{M} s. \quad (6)$$

$\lambda(s)$ is the utility loss incurred by a player when the symmetric contribution profile s is replaced by an asymmetric profile in which a fraction $1/M$ of a player's opponents contributes nothing, and the remaining fraction contributes s . Define $\mu(s)$ and $\delta(s)$ as in (4)–(5), for the limit case:

$$\begin{aligned} \mu(s) &= u(s) - u(0) = \lim_{\substack{N \rightarrow \infty, \\ \alpha N = c}} (\alpha N - 1)s = (c - 1)s, \\ \delta(s) &= u(0, s) - u(s) = \lim_{\substack{N \rightarrow \infty, \\ \alpha N = c}} (1 - \alpha)s = s. \end{aligned}$$

To construct the path σ , begin by setting $\sigma^1 = \lim_m \varepsilon / (1 - \alpha_m) = \varepsilon$. Let $\tilde{B}_1 = \lambda(\sigma^1)$ and set $\sigma_2 = \max\{s \in [0, K] : \delta(s) \leq \tilde{B}_1 + \varepsilon\}$. This gives $\sigma^2 = (c/M + 1) > \varepsilon = \sigma^1$. Now set $\tilde{B}_2 = \lambda(\sigma^2) + B_1$, where B_1 is defined as before (i.e. $B_1 = \mu(\sigma^1)$). \tilde{B}_2 is the cumulative utility loss a defector in round 3 can expect. We can then set $\sigma_3 = \max\{s \in [0, K] : \delta(s) \leq \tilde{B}_2 + \varepsilon\}$. In general, σ^t can be constructed recursively, by letting $\sigma^1 = \varepsilon$ and $B_0 = 0$, and setting

$$\begin{aligned} B_{t-1} &= \mu(\sigma^{t-1}) + B_{t-2}, \\ \tilde{B}_{t-1} &= \lambda(\sigma^{t-1}) + B_{t-2}, \\ \sigma^t &= \max\{s \in [0, K] : \delta(s) \leq \tilde{B}_{t-1} + \varepsilon\}, \end{aligned}$$

for $t \geq 2$. To explain the formula for \tilde{B}_t , recall that one round after the initial defection, a fraction $1/M$ of the players have stopped contributing. And exactly two rounds after the initial defection, all players have stopped contributing.

Comparing the definitions of \tilde{B}_t and B_t , it is clear that the punishment possibilities are more limited when $M > 1$ than when $M = 1$. Thus, maximal equilibrium contributions will be lower in the random reassignment model. The “Strangers” treatment is hence successful at reducing the strategic effects. However, for large N and a slight amount of satisficing behavior, it is still not enough to fully eliminate them.

6 Conclusion

This paper examined strategic behavior in repeated VCM game which is consistent with several stylized facts from the experimental literature, concerning the temporal pattern of contributions, the restart effect, as well as various effects of changes in group size and technology. To get such strategic behavior started, of course, we needed to make assumptions on either preferences or rationality. The main point, however, was to demonstrate that even when these assumption are mild—in the sense that they would not significantly alter the behavior of myopic players—their effects on the dynamic path of contributions is large. We here conclude with a few remarks.

First, the idea that withholding future contributions can punish free-riders effectively has also been explored in Marx and Matthews (2000) in the context of raising funds to build a single public project: Instead of trying to raise all required funds to complete a project at once, a “piece-meal approach” in which smaller amounts are contributed over time can be more successful. To prevent an unravelling of contributions, one needs to make the assumption that the benefit function for the public project exhibits a jump at some point. The existence of this jump serves a role similar to the multiplicity of stage-game equilibria used in this paper.

Second, in order to achieve significant positive contributions, players must understand the opportunity they have to reward cooperative behavior by coordinating on the public spirit equilibrium at the last stage of the dynamic game, and sanction free-riding behavior by coordinating on the selfish equilibrium. These rewards and punishments are *implicit* in that they arise endogenously through the multiplicity of stage game equilibria. An independent experimental literature on *explicit* rewards and punishments has emerged recently (e.g., Fehr and Gaechter [2000]). There, subjects are given a mechanism to induce costs on others; however such punishment also incurs a cost on the person who administers the sanction. In this sense, the public good problem is only deferred to a higher level, but not solved. Nevertheless, it is shown that such higher-order mechanisms tend to increase cooperative behavior throughout all stages of the game.¹²

¹²This raises the question of whether an experimental design is available that would allow one to test if the implicit punishments provided in this model are actually what drives cooperation. To this end, one could provide the subject group with a simple mechanism for inflicting a *credible* punishment after the last stage of a conventional VCM game. This could be achieved, for example, by letting them play a coordination game after the VCM. The coordination game would have two equilibria, one with high payoff the other with low payoff, serving the same role as the equilibria identified in Lemma 1 and Lemma 2.

Appendix

Proof of Lemma 1 and Lemma 2

Proof of Lemma 1. A player's best response to an average opponent contribution of \bar{s}_{-i} is given by

$$b(\bar{s}_{-i}) = \min \left\{ K, \left[\frac{\gamma_1 \beta}{1 - \alpha} (\bar{s}_{-i})^{\gamma_2} \right]^{1/(1-\gamma_1)} \right\}.$$

By the assumptions made on γ_1 and γ_2 , b is non-decreasing. Since $b(0) = 0$, the zero-contribution profile $s = (0, \dots, 0)$ is a Nash equilibrium in $\Gamma_1(\beta, 0)$. To arrive at a second symmetric equilibrium where $s_i = \hat{s} > 0 \forall i$, we need $0 < \hat{s} = \min\{K, \beta(\hat{s})\}$, which gives $\hat{s} = \min\{K, [\gamma_1 \beta / (1 - \alpha)]^{1/(1-\gamma_1-\gamma_2)}\}$. Thus, as $\beta \rightarrow 0$ we have $\hat{s} \rightarrow 0$. These are the only symmetric equilibria. To see that there can be no asymmetric equilibria, suppose the profile s is an equilibrium with $s_i < s_j$, for some i, j . This implies $\bar{s}_{-i} > \bar{s}_{-j}$, and since b is non-decreasing we have $s_i = b(\bar{s}_{-i}) \geq b(\bar{s}_{-j}) = s_j$, a contradiction. \square

Proof of Lemma 2. Fix any player i and let s_{-i} be any profile of contributions among i 's opponents. i 's utility is

$$u_i(s_i, s_{-i}) = 1 - (1 - \alpha)s_i + \alpha \sum_{j \neq i} s_j.$$

The maximum utility player i can attain is $1 + \alpha \sum_{j \neq i} s_j$ (if $s_i = 0$). Thus the difference between this maximum utility and i 's actual utility is

$$\max_{s'_i \in [0, K]} u_i(s'_i, s_{-i}) - u_i(s_i, s_{-i}) = (1 - \alpha)s_i,$$

which is less or equal to ε if and only if $s_i \leq \varepsilon / (1 - \alpha)$. \square

Proof of Theorem 3

Rest of the proof for $\Gamma_T(\beta, 0)$. Most of the proof is in the main text. Here we show that σ increases and \bar{T} decreases in β . Since $\sigma^1 = \hat{s}$ and \hat{s} increases in β , the result is true for σ^1 . Now consider $\sigma^2 = \max\{s \in [0, K] : \delta(s) \leq B_1\}$. We will show that $\delta(s)$ decreases and B_1 increases in β . By definition, $\delta(s)$ is the maximal gain a player can have by deviating from the symmetric profile s (namely to $b(s)$). If $s > 0$ then $b(s) > 0$, so the envelope theorem implies $d\delta(s)/d\beta = \rho(b(s), s) - \rho(s) < 0$. Thus, $\delta(s)$ decreases in β . To see that B_1 increases in β recall that $B_1 = \mu(\hat{s}) = (\alpha N - 1)\hat{s} + \beta\rho(\hat{s})$. Since \hat{s} increases in β and $\alpha N > 1$, B_1 increases in β . Thus, σ^2 must increase in β . The same argument can now be repeatedly applied for σ^t , $t > 2$, which establishes that σ is increasing in β . Since the maximal possible contribution remains fixed at K , \bar{T} must be decreasing in β . \square

Proof for $\Gamma_T(0, \varepsilon)$. We will use the trigger strategy with unrelenting punishments described in (3). The argument will be similar to the one given in the main text, with the exception that an additional step is required in the end (the additional step concerns the choice of which period to allocate the permitted ε -mistakes to).

Define functions $\delta : [0, K] \rightarrow \mathbb{R}$ and $\mu : [0, K] \rightarrow \mathbb{R}$ along the same lines as in the text:

$$\begin{aligned}\delta(s) &= u(0, s) - u(s) = (1 - \alpha)s, \\ \mu(s) &= u(s) - u(0) = (\alpha N - 1)s > 0.\end{aligned}$$

At stage 1, an ε -equilibrium of $\Gamma_1(0, \varepsilon)$ must be played, regardless of the history that preceeded the last stage. By Lemma 2 there is a continuum of candidate contribution levels, ranging from 0 to $\varepsilon/(1 - \alpha)$. Set $\sigma^1 = \varepsilon/(1 - \alpha)$ and let

$$B_1 = \mu(\sigma^1) > 0$$

be the payoff difference for a single player between this ε -equilibrium and the alternative zero-contributions equilibrium.

Now move to period $t = 2$. Suppose that $\sigma^2 > \sigma^1$, and consider a deviation to 0. For the same reason as before, it is sufficient to focus on deviations to a player's best response in the stage game. The deviating player gains $\delta(\sigma^2) > 0$ at stage 2, and loses B_1 at stage 1. In order to support σ^2 as an ε -optimal contribution at stage 2, we need $\delta(\sigma^2) \leq B_1 + \varepsilon$. Therefore

$$\sigma^2 = \max\{s \in [0, K] : \delta(s) \leq B_1 + \varepsilon\}.$$

A well defined σ^2 exists, and $\sigma^2 > \sigma^1$ as long as $\sigma^1 < K$. (Otherwise $\sigma^2 = \sigma^1 = K$). Proceed now in exactly the same fashion as for $\Gamma_T(\beta, 0)$. That is, using the starting values $B_0 = 0$ and $\sigma^1 = \varepsilon/(1 - \alpha) > 0$, recursively compute

$$B_t = B_{t-1} + \mu(\sigma^t) > B_{t-1}$$

and

$$\sigma^t = \max\{s \in [0, K] : \delta(s) \leq B_{t-1} + \varepsilon\}.$$

Observe that $\sigma^t > \sigma^{t-1}$ if $\sigma^{t-1} < K$, and $\sigma^t = \sigma^{t-1} = K$ otherwise. The period \bar{T} is now defined as before: $\bar{T} = \min\{t : \sigma^t = K\}$.

The comparative properties of σ and \bar{T} with respect to ε can also be established in the same manner as before. Since $\sigma^1 = \varepsilon/(1 - \alpha)$, σ^1 clearly increases in ε . Now consider $\sigma^2 = \max\{s \in [0, K] : \delta(s) \leq B_1\}$. $\delta(s)$ does not depend on ε , and $B_1 = \mu(\sigma^1)$ since σ^1 increases in ε and $\alpha N > 1$ implies μ is increasing. Thus, σ^2 must increase in ε . The same argument can now be repeatedly applied for σ^t , $t > 2$, which establishes that σ is increasing in ε . Since the maximal possible contribution remains fixed at K , \bar{T} must be decreasing in ε .

Finally, note that in every round of this equilibrium, each player makes an ε -mistake. In principle, it is possible to shift this ε -mistake from round t to round $t' > t$ and still have an ε -SGPE. However, this will decrease the total amount of public good provided. To see this, consider decreasing σ^1 by a small amount. The definition of an ε -SGPE now permits us to increase σ^t by exactly the same amount, for some $t > 1$. At the same time, however, the decrease in σ^1 will decrease B_1 , which diminishes σ^2 , which in turn decreases B_2 and thus σ^3 , and so on. Thus, to construct the maximal path σ , it is necessary and sufficient to set σ^1 to the largest level possible under the ε -optimality criterion. The same argument can be made to show that the ε -mistake contained in σ^2 should not be shifted to σ^t ($t > 2$), and so on. The resulting path $\sigma = (\sigma^1, \dots, \sigma^T)$ is therefore the maximal path supported by the trigger strategy w . \square

Proof of Theorem 4 and Theorem 5

Proof of Theorem 4. Consider the game $\Gamma_T(\beta, 0)$, for $\beta > 0$. Suppose first that α increases. Then $\sigma^1 = \hat{s} = [\gamma_1\beta/(1-\alpha)]^{1/(1-\gamma_1-\gamma_2)}$ increases. Now consider $\sigma^2 = \max\{s \in [0, K] : \delta(s) \leq B_1\}$. We show that $\delta(s)$ decreases and B_1 increases in α , for $s \geq \sigma^1$. To see that $\delta(s)$ decreases in α , note that by the envelope theorem, $d\delta(s)/d\alpha = -(s - b(s)) \leq 0$ if $s \geq b(s)$. Since we are considering the case $s \geq \sigma^1$ (and $\sigma^1 = \hat{s} = b(\hat{s})$), we have that $\delta(s)$ decreases in α . To see that B_1 increases in α recall that $B_1 = \mu(\sigma^1) = (\alpha N - 1)\sigma^1 + \beta\rho(\sigma^1)$. Hence $dB_1/d\alpha = N\sigma^1 > 0$. Thus, σ^2 must increase in α . The same argument can be repeatedly applied for σ^t ($t > 2$). Suppose next that N increases. This does not change $\sigma^1 = \hat{s}$. For σ^2 , note that $dB_1/dN = \alpha\sigma^1 > 0$ and $d\delta(s)/dN = 0$. Hence σ^2 increases in N . Repeat this argument for σ^t ($t > 2$); the result follows. Finally, since the maximal possible contribution remains fixed at K , \bar{T} must be decreasing in both α and N .

Now consider the game $\Gamma_T(0, \varepsilon)$. Suppose first that α increases. Then $\sigma^1 = \varepsilon/(1-\alpha)$ increases. Now consider $\sigma^2 = \max\{s \in [\sigma^1, K] : \delta(s) \leq B_1\}$. We show that $\delta(s)$ decreases and B_1 increases in α . To see that $\delta(s)$ decreases in α , note that by the envelope theorem, $d\delta(s)/d\alpha = -s < 0$. To see that B_1 increases in α recall that $B_1 = \mu(\sigma^1) = (\alpha N - 1)\sigma^1$. Hence $dB_1/d\alpha = N\sigma^1 > 0$. Thus, σ^2 must increase in α . Repeating the argument for σ^t ($t > 2$) establishes that the path σ is increasing in α . Suppose next that N increases. This does not change $\sigma^1 = \varepsilon/(1-\alpha)$. For σ^2 , note that $dB_1/dN = \alpha\sigma^1 > 0$ and $d\delta(s)/dN = 0$. Hence σ^2 increases in N . Repeat this argument for σ^t ($t > 2$); the result follows. Again, since the maximal possible contribution remains fixed at K , \bar{T} must be decreasing in both α and N . \square

Proof of Theorem 5. Consider the game $\Gamma_T(\beta, 0)$, for $\beta > 0$. To prove the first part, we follow the same steps as in the proof of Theorem 4. Suppose α decreases, N increases, but αN remains constant. Clearly $b(s)$ and \hat{s} decrease; hence $\sigma^1 = \hat{s}$ decreases. Consider $\sigma^2 = \max\{s \in [0, K] : \delta(s) \leq B_1\}$ next. We show that $\delta(s)$ increases and B_1 decreases.

$B_1 = \mu(\sigma^1) = (\alpha N - 1)\sigma^1 + \beta\rho(\sigma^1)$ clearly decreases, as σ^1 decreases and αN remains constant. As for $\delta(s)$, by the envelope theorem $d\delta(s)/d\alpha = -(s - b(s)) < 0$ (while $d\delta(s)/dN = 0$). Thus, $\delta(s)$ increases, implying that σ_2 decreases. Repeating the argument for all σ_t ($t > 2$) establishes the first part of the theorem. To construct the limit path, notice that $(1 - \alpha_m) \rightarrow 1$ as $m \rightarrow \infty$. Hence we have

$$\omega^1 = \lim_{m \rightarrow \infty} \hat{s} = (\gamma_1 \beta)^{1/(1-\gamma_1-\gamma_2)}.$$

Now let

$$\bar{B}_1 = \lim_{m \rightarrow \infty} B_1 = \mu(\omega^1),$$

and set $\omega^2 = \max\{s \in [0, K] : \delta(s) \leq \bar{B}_1\}$. This is well defined with $\omega^2 > \omega^1$. Proceeding in this fashion, one can construct the entire limit path ω . As in the proofs before, entirely analogous steps apply for $\Gamma_T(0, \varepsilon)$, which are therefore omitted here. \square

References

- [1] Ahn, T., E. Ostrom, D. Schmidt, R. Shupp, and J. Walker (2001): “Cooperation in PD Games: Fear, Greed, and History of Play,” *Public Choice*, **106**, 137–155.
- [2] Ambrus, A. and P. Pathak (2009): “Cooperation over Finite Horizons: A Theory and Experiments,” working paper, Harvard University.
- [3] 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**, 297–323.
- [4] Andreoni, J. (1988): “Why Free Ride? Strategies and Learning in Public Goods Experiments,” *Journal of Public Economics*, **37**, 291–304.
- [5] Andreoni, J. (1990): “Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving,” *Economic Journal*, **100**, 464–477.
- [6] Andreoni, J. (1995): “Cooperation in Public Goods Experiments: Kindness or Confusion?,” *American Economic Review*, **85**, 891–904.
- [7] Andreoni, J. and R. Croson (2008): “Partners vs. Strangers: Random Matching in Public Goods Experiments,” in C. Plott and V. Smith, eds.: *Handbook of Experimental Economics Results*, Amsterdam: Elsevier.
- [8] Andreoni, J. and J. Miller (1993): “Rational Cooperation in the Finitely Repeated Prisoners’ Dilemma: Experimental Evidence,” *Economic Journal*, **103**, 570–585.
- [9] Benoit, J. and V. Krishna (1985): “Finitely Repeated Games,” *Econometrica*, **53**, 905–922.

- [10] Brandts, J. and N. Figueras (2003): “An Exploration of Reputation Formation in Experimental Games,” *Journal of Economic Behavior and Organization*, **50**, 89–115.
- [11] 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.
- [12] Charness, G. and M. Rabin (2002): “Understanding Social Preferences with Simple Tests,” *Quarterly Journal of Economics*, **117**, 817–869.
- [13] Croson, R. (1996): “Partners and Strangers Revisited,” *Economics Letters*, **53**, 25–32.
- [14] Eshel, I., L. Samuelson, and A. Shaked (1998): “Altruists, Egoists, and Hooligans in a Local Interaction Model,” *American Economic Review*, **88**, 157–179.
- [15] Ellison, G. (1994): “Cooperation in the Prisoner’s Dilemma with Anonymous Random Rematching,” *Review of Economic Studies*, **61**, 567–588.
- [16] Fehr, E. and S. Gaechter (2000): “Cooperation and Punishment in Public Goods Experiments,” *American Economic Review*, **90**, 980–994.
- [17] Isaac, M., J. Walker, and S. Thomas (1984): “Divergent Evidence on Free Riding: An Experimental Examination of Possible Explanations,” *Public Choice*, **43**, 113–149.
- [18] Isaac, M. and J. Walker (1988): “Group Size Effects in Public Goods Provision: The Voluntary Contribution Mechanism,” *Quarterly Journal of Economics*, **53**, 179–200.
- [19] 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.
- [20] Isaac, M. and J. Walker (1998): “Nash as an Organizing Principle in the Voluntary Provision of Public Goods: Experimental Evidence,” *Experimental Economics*, **1**, 191–206.
- [21] 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.
- [22] Kreps, D. and R. Wilson (1982): “Sequential Equilibrium,” *Econometrica*, **50**, 863–894.

- [23] Ledyard, J. (1995): “Public Goods: A Survey of Experimental Research,” in: J. Kagel and R. Roth (eds.), *Handbook of Experimental Economics*, Princeton University Press, Princeton, NJ.
- [24] Marx, L. and S. Matthews (2000): “Dynamic Voluntary Contribution to a Public Project,” *Review of Economic Studies*, **67**, 327–358.
- [25] Palfrey, T. and J. Prisbrey (1996): “Altruism, Reputation and Noise in Linear Public Goods Experiments,” *Journal of Public Economics*, **61**, 409–427.
- [26] Palfrey, T. and H. Rosenthal (1988): “Private Incentives in Social Dilemmas: The Effects of Incomplete Information and Altruism,” *Journal of Public Economics*, **35**, 309–332.
- [27] Pecorino, P. (1999): “The Effect of Group Size on Public Good Provision in a Repeated Game Setting,” *Journal of Public Economics*, **72**, 121–134.
- [28] Radner, R. (1980): “Collusive Behavior in Noncooperative Epsilon-Equilibria of Oligopolies with Long but Finite Lives,” *Journal of Economic Theory*, **22**, 136–154.